

Robust bilinear-noise-optimal control for gravitational-wave detectors: A mixed LQG and $\mathcal{H}_\infty$ approach

Ian A. O. MacMillan^{*} and Lee P. McCuller[✉]

*California Institute of Technology, PMA, 1200 East California Boulevard,
Pasadena, California 91125, USA*



(Received 12 January 2026; accepted 2 March 2026; published 13 April 2026)

At its lowest frequencies, LIGO is limited by noise in its many degrees of freedom of suspended optics, which, in turn, introduce noise in the interferometer through their feedback control systems. Nonlinear interactions are a dominant source of low-frequency noise, mixing noise from multiple degrees of freedom. The lowest-order form is bilinear noise, in which the noise from two feedback-controlled subsystems multiplies to mask gravitational waves. Bilinear couplings require control trade-offs that simultaneously balance high- and low-frequency noise. Currently, there is no known lower limit to bilinear control noise. Here, we develop benchmark cost functions for bilinear noise and associated figures of merit. Linear-quadratic-Gaussian (LQG) control then establishes aggressive feedback that saturates the lower bounds on the cost functions. We then develop a mixed LQG and $\mathcal{H}_\infty$ approach to directly compute stable, robust, and optimal feedback, using the LIGO's alignment control system as an example. Direct computations are fast while ensuring a global optimum. By calculating optimal robust control, it is possible to construct the lower bound on control noise along the Pareto front of practical controllers for LIGO. This method can be used to drastically improve control noise in existing observatories as well as to set subsystem control noise requirements for next-generation detectors with parametrized design.

DOI: [10.1103/wm9p-zmxc](https://doi.org/10.1103/wm9p-zmxc)

I. INTRODUCTION

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is a ground-based gravitational-wave interferometer that first detected gravitational waves in 2015 [1]. Since its first detection, the LIGO Scientific Collaboration has been working to reduce noise in the instrument to increase its sensitivity and, therefore, detect even fainter gravitational-wave signals [2,3]. These significant improvements in LIGO's sensitivity, as well as the inclusion of additional detectors, have led to an increased detection rate that expands astrophysical study [2–7].

There is considerable interest in reducing low-frequency noise in LIGO (< 30 Hz). There are many low-frequency waveforms, like high-mass binary black hole inspirals, whose detection would benefit from an increased sensitivity in the low-frequency band [8]. Other waveforms that extend to higher frequencies will benefit from higher signal-to-noise detections, improved parameter estimation, and potentially early warning of ongoing events during the

low-frequency portion of long-lived waveforms. Currently, at these low frequencies, both LIGO detectors are limited by noise introduced by the feedback control loops that stabilize auxiliary degrees of freedom needed to maintain interferometer operation. Due to the use of resonant cavities, even small variations in the auxiliary degrees can result in highly nonlinear responses that can produce noise. The lowest meaningful order of nonlinearity results in two channels, one a control signal and one a drifting parameter, which multiply to produce bilinear noise that is indistinguishable from an astrophysical strain signal. Improving the controllers used in LIGO's feedback loops and generally improving the processes for creating these controllers could significantly reduce the low-frequency noise in the detectors, as well as allowing for resources to be spent on other areas of improvement.

A. Motivation

The LIGO instrumentation community widely uses classical control design methods. While an entire system has many degrees of freedom, they are almost all designed and engineered as individual single-input and single-output (SISO) signal paths using transfer functions, Bode plots, and frequency-resolved budgets of noise contributions. These subsystems and control loops are designed depending on how they interconnect to affect the principal measurement

^{*}Contact author: macmillan@caltech.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

channels. The software interfaces of LIGO are optimized for this approach, and it is largely successful; however, it has several weaknesses. First, the loops are all hand designed, which can get the system operating fast, but compounds negatively for global optimization. In particular, without a lower bound on performance, a halting criterion for optimizing is established either by experience or by exhaustion. Furthermore, the design of coupled multiple-input and multiple-output (MIMO) systems must often be iterative, as loop shapes and noise consequences are interdependent. More modern state-space control approaches are an attractive research topic, promising guaranteed existence, stability, and optimality properties even for complicated, intercoupled MIMO systems.

So far, modern control approaches have not been widely adopted due to the unique requirements of gravitational-wave detectors. First, these approaches require very specific cost functions to form well-posed optimization problems. Second, these approaches usually assume white noise inputs; thus, system models must be augmented with noise-shaping filters. The fitting of such filters is nontrivial but now reliable with modern algorithms. The third issue is that interferometers rarely have actuator cost functions (unlike, e.g., aerospace); thus, typical examples of regulator cost do not immediately apply. The fourth issue is that interferometer loops are limited by transmission delays, which create non-minimum-phase zeros that must be incorporated into the control loop design. The final and highly significant hurdle is dynamic range. The range of shaped noise in gravitational-wave detectors spans 8–12 orders of magnitude, from low-frequency seismic motion to interferometric sensors. The combination of these factors causes optimal solutions to sit on the edge of stability regions, where even the best-known numerical algorithms struggle.

B. Approach

This manuscript adopts an intermediate and transitional approach in an effort to supplant hand tuning while managing these challenges. It uses modern methods to pose and solve a SISO control problem. Future work can extend its examples to more intrinsically MIMO problems. A major focus of this work is the problem definition, to pose the modern control problem using figure of merit (FOM) that are relevant to optimizing gravitational-wave detectors. Furthermore, the mixed-sensitivity approach we develop is necessary to find robust solvers that have acceptable gain and phase margins. This approach is mathematically complex to describe, but has largely been solved in previous control literature. By finding solutions with acceptable phase margins, much like our hand-designed controllers, the numerical solutions are generally better behaved and do not live “near the edge” of admissible solution regions, where small perturbations, numerically or in the instrument plants, push them outside. Our mixed-sensitivity solver rapidly arrives at controllers qualitatively similar to hand

designs but guaranteed optimal for our defined figures of merit that are tied to the astrophysics objectives.

The organization of this paper is to be a bridge between classical and modern control approaches. The first half provides the problem setup and results using the language of transfer functions and noise projections. Our results are included in these sections, showing how to form a modern control optimization problem to minimize bilinear control noise and the resulting sets of controllers. The optimal controllers set bounds on the best possible noise performance, useful for determining a halting point for hand designs. We then transform the problem description to use a mixed-sensitivity approach, yielding controllers with acceptable and adjustable stability margins. The second half of the paper translates the problem and solutions into state-space form and summarizes the literature and approach required to build our custom mixed-sensitivity numerical solver. This combination of approaches establishes a framework to utilize modern control systems to solve open control challenges in gravitational-wave detectors.

II. BACKGROUND

LIGO’s sensitivity is limited by the sum of a number of individual noise sources in the interferometer. The extent to which these noise sources obscure astrophysical signals is expressed using the signal-to-noise ratio (SNR). The SNR of a signal can be calculated from the Fourier transform of the event waveform $h_{\text{gw}}(f)$ and the power spectral density (PSD) of the interferometer’s noise $S_h(f)$, using

$$\text{SNR}^2 = 4 \int_0^\infty \frac{|h_{\text{gw}}(f)|^2}{S_h(f)} df. \quad (1)$$

The detection range,

$$d_{\text{range}} = \frac{d_{\text{gw}}}{2.26} \frac{\text{SNR}}{\rho_0}, \quad (2)$$

is the average effective distance to which events louder than a threshold SNR called ρ_0 can be detected. The range can be calculated using the SNR formula and the distance to the fiducial source. For an $h_{\text{gw}}(f)$ waveform modeled from a source at distance d_{gw} , the detection range averaged over the sky is included above. The threshold SNR is usually taken to be $\rho_0 = 8$, and the factor of 2.26 accounts for the average of the interferometer antenna pattern over the sky [9].

The astronomical detection range is critical to the interferometer’s sensitivity because it determines the detection rate $\mathcal{R}_{\text{det}}$. The astronomical detection range, detection rate, and SNR are related through

$$\mathcal{R}_{\text{det}} \propto d_{\text{range}}^3 \propto \text{SNR}^3. \quad (3)$$

Even minor decreases in the interferometer’s noise can have a significant impact on the detection rate. One of LIGO’s

long-standing goals is increasing the detection rate for binary neutron star (BNS) inspirals. Binary neutron star waveforms are taken to be the standard fiducial for computing SNR as a FOM for the interferometer performance in this work.

A. Noise from control systems

At the low-frequency end of its detection band, LIGO's sensitivity is limited by noise injected by control systems. At frequencies below LIGO's detection band, the control systems stabilize the operating points of otherwise highly nonlinear optical cavity resonators and interferometers. Noise coming from the control actuation of cavity alignment and auxiliary length are two of the largest of all known noise sources in the detector below 10 Hz [2,3]. Interferometric control noise is any noise that arises from the addition of a stabilizing or noise-suppressing controller to a degree of freedom of the interferometer system. When a controller is added in a feedback loop to a degree of freedom, some of the noise inherent to that degree of freedom, along with new measurement noise, passes through the controller and is (re)injected into the system at the actuator. While a controller can reduce the overall noise of a degree of freedom, it also introduces a new noise source from the controller's sensor. This occurs when the loop bandwidth is higher than the frequency where (input-referred) plant noise becomes lower than the sensing noise. For complex systems with many degree of freedom and many control systems, the control noise can become a significant noise source. From optics controllers to seismic isolation, each LIGO detector has more than 300 controllers that all maintain its operating point but add noise to the principal readout: the differential-arm length (DARM). This is the measurement of arm-length differences caused by gravitational waves and thus the output of interest for astrophysics.

Here, a model is constructed for how a control system for a single degree of freedom affects the signal and introduces noise. We will later express bilinear noise in terms of this model. The plant noise of some degree of freedom under feedback control N_c affects the signal seen at the output of the interferometer through

$$h_{\text{signal}}(f) = h_{\text{gw}}(f) + N_{\text{det}}(f) + C(f)N_c(f), \quad (4)$$

where $h_{\text{gw}}(f)$ is a potential signal, N_{det} is all the detector noise from sources other than the degree of freedom being modeled, and $C(f)$ is a frequency-dependent coupling of the degree of freedom into DARM. This coupling is often indirect and sourced from bilinear effects. For simple estimates, it is often possible to assume $C(f)$ is constant across all frequencies, but in the general case, measurement data necessitate that it be frequency dependent. An example measurement of $C(f)$ and physical motivation are described in Sec. III B. The PSD of the total noise in the detector is thus

$$S_h(f) = S_{\text{det}}(f) + |C(f)|^2 S_c(f), \quad (5)$$

where $S_{\text{det}}(f)$ is the PSD of all the detector noise from sources other than the degree of freedom being modeled and $S_c(f)$ is the PSD of the noise of interest, in this case a definition of control noise. By linearizing Eq. (1), and including Eq. (5)'s expression of $S_h(f)$, the control-noise-dependent SNR becomes

$$\text{SNR}^2 = \text{SNR}_{\text{det}}^2 - \text{SNR}_{\text{lost}}^2, \quad (6)$$

where

$$\text{SNR}_{\text{det}}^2 = 4 \int_0^\infty \frac{|h_{\text{gw}}(f)|^2}{S_{\text{det}}(f)} df \quad (7)$$

and

$$\text{SNR}_{\text{lost}}^2 = 4 \int_0^\infty \frac{|C|^2 |h_{\text{gw}}(f)|^2}{S_{\text{det}}(f)^2} S_c(f) df. \quad (8)$$

Adapting Eq. (2) gives the control-noise-dependent expression for the detection range lost as

$$d_{\text{lost}} = \frac{d_{\text{gw}}}{2.26\rho_0} \frac{\text{SNR}_{\text{lost}}^2}{2\text{SNR}_{\text{det}}}. \quad (9)$$

The lost detection range can also be computed using the linearization. The point of linearizing the equation is to split it into terms that show the specific impact of a particular noise, S_c , on the SNR and BNS range. In this equation, S_c can be any noise of interest, assuming it is small enough to keep the linearization accurate. The S_c term is inside an integral that includes a frequency-dependent weighting. That weighting term will later be useful to define a shaping filter for the control optimization problem in terms of the lost SNR or range cost.

The root mean square (rms) of a noise from channel N_c is indicated by $\tilde{R}_c$. In order to keep the interferometer locked, there is a maximum allowable rms. If the noise in the interferometer exceeds that limit, the interferometer is unable to maintain operation, e.g., sufficient alignment,

$$\tilde{R}_c \equiv \sqrt{\int_0^\infty S_c(f) df} < \tilde{R}_{c, \text{unlock}}. \quad (10)$$

A fundamental requirement of control systems is to ensure that no BNS signal exceeds its lock-loss threshold. Thus, there is a trade-off between operational requirements and noise injection.

There are two primary sources of control noise: environmental noise and measurement noise. Environmental noise is any noise that acts on the states of the suspended interferometer plant. For example, seismic noise is an

environmental noise that is of great concern for ground-based interferometers. Measurement noise is the difference between the measured states of a system and their true value, generally due to an imperfect intermediate sensor used to probe the states. Both the measurement and environmental noise pass through a control filter and then feed back into the system, but measurement noise would not otherwise affect plant states except through the control system actuating it into real plant noise.

Considering only two noise sources is a useful simplification. In reality, there are many noise sources that all contribute differently to every degree of freedom in the interferometer. However, these noise sources can only enter the system in two places: either before the plant (environmental noise) or before the controller (measurement noise). This distinction also aligns with the assumptions and model of a typical optimization using classical control theory to shape the open-loop gain.

Without physically decreasing environmental or measurement noise, which much of the engineering resources of LIGO are dedicated to, the best way to decrease control noise is to find an optimal controller K , such that the minimum amount of noise is fed through into DARM while achieving a total rms goal.

Here we distinguish two channels N_p and N_a . The former, which must meet a noise threshold, is the sum of environmental noise and controller actuation, filtered by the plant's transfer function. This noise, N_p , is the noise directly seen by the interferometer. The latter, N_a , is the

signal sent by an actuator via its control system and shaped by the plant's transfer function, excluding any direct contributions from environmental noise. These two noises are soon presented visually in Fig. 1. The optimal controller must ensure that the plant's noise $\tilde{R}_p$ will meet requirements by shaping its spectrum S_p through the environmental and measurement noise contributions,

$$S_p(f) = \left| \frac{P(f)}{1 - G(f)} \right|^2 S_{\text{env}}(f) + \left| \frac{G(f)}{1 - G(f)} \right|^2 S_{\text{meas}}(f), \quad (11)$$

where

$$G(f) = K(f)P(f) \quad (12)$$

is the open-loop gain of the full system, $P(f)$ is the transfer function of the plant component, either the input or output of $P(f)$ represents the degree of freedom to be controlled, $K(f)$ is the transfer function of the controller, $S_{\text{env}}(f)$ is the PSD of the environmental noise, and $S_{\text{meas}}(f)$ is the PSD of the measurement noise. Here, it is explicitly indicated that each of these expressions is a function of frequency; this dependence is continued but will not be explicitly stated in the following formulas.

The total plant noise N_p , which is presented in Eq. (11), is the typical minimization metric for modern control. It is the measure of all the noise coming out of a system, including all coherent cancellations and interactions. In a system with no cross-coupling to any other degree of freedom, minimizing S_p would minimize the sum of both




FIG. 1. The layout of the modeled system with a controller attached. Here, the input white noise u_{n1} is being shaped to appear as environmental noise seen by the plant. The system has two figures of merit, which are simply transfer functions that weigh the output noise: a broadband FOM outputs the unshaped noise, i.e., every frequency has the same weight, and a BNS inspiral FOM, which represents the band of detection for BNS signals. In this model, the two plant blocks are identical and serve as a way to represent S_a for actuation-point control noise minimization. These figures of merit are discussed further in Sec. III B. The input associated with u_∞ , which has a dotted line, is used for the $\mathcal{H}_\infty$ minimization discussed in Sec. VA.

measurement noise and environment noise injected into the system.

The definition of N_p is often presented as control noise in modern control. However, this differs from most previous work in gravitational-wave interferometer sensing and control, where control noise is defined as any noise that passes through the controller and becomes a new contribution through actuation on the degree of freedom. The definition of actuation-point control noise S_a is

$$S_a = \left| \frac{PG}{1-G} \right|^2 S_{\text{env}} + \left| \frac{G}{1-G} \right|^2 S_{\text{meas}}, \quad (13)$$

and

$$N_a = \frac{PG}{1-G} N_{\text{env}} + \frac{G}{1-G} N_{\text{meas}}. \quad (14)$$

The actuation point that control noise definition is useful because optimizing to reduce it is more directly a means of penalizing actuation, as K is in the numerator of both environmental and measurement noise contributions to Eqs. (13) and (14). In large multi-degree of freedom coupled systems like LIGO, it is often true that actuation on one degree of freedom can inject noise into another. The PSD S_c , from Eq. (6), could be set equal to any noise of interest like S_p ; however, in this paper $S_c = S_a$ as is appropriate for the alignment control degree of freedom we investigate and as is typical for gravitational-wave (GW) control. The noise from this metric N_a is then used as the noise N_a in Eq. (4). An approach more widely used in modern control that only uses S_p as an optimization metric is discussed in Appendix B.

In this work, both S_a and S_p will be used: the former for the control effect on SNR and BNS range, and the latter for the minimization of the total rms $\tilde{R}_p$ from Eq. (10), needed to maintain the operating point. Finding a “best” controller is difficult because of the trade-off between suppressing environmental noise in S_p with a large K and suppressing measurement noise in S_a with a small K . Finding the optimal controller requires a multiobjective optimization: the rms control noise $\tilde{R}_p$ and the BNS detection range.

B. Application to alignment sensing control system

Control noise associated with the alignment sensing control (ASC) system is one of the largest sources of control noise at low frequencies [2,3]. The ASC system is responsible for keeping the LIGO test masses aligned in their angular degrees of freedom. Each LIGO interferometer has eight principal control filters dedicated to maintaining angular stability in the arm cavities. These eight control filters are the product of three distinguishing degrees. Each system is either in pitch or yaw (P/Y) rigid-body rotation, differential or common motion between the two Michelson

arms (C/D), and HARD and SOFT terms in the labels indicate the type of spring reaction to beam motion.

HARD modes are when the two test masses that make up LIGO’s Fabry-Pérot arm cavities rotate in opposite directions, e.g., one rotates very slightly up, and the other rotates very slightly down. Since LIGO’s Fabry-Pérot cavities have a high g-factor, the differential rotation causes a rotation of the optical axis, and radiation pressure pushes the arms back toward their nominal angle. Thus, HARD modes increase suspension stiffness and occur at higher frequencies. SOFT modes are the opposite, where the mirrors move at a common angle, which translates the cavity axis. The translation moves the beam off center and produces a negative spring constant (antispring).

The name of each mode is the composition of these distinguishing degrees; i.e., common-soft-yaw is the common soft mode in the yaw angular tilt. In this study, only the differential hard (DHARD) yaw controller is calculated. The choice of actuation-point noise optimization for the DHARD yaw degree of freedom presented in this paper was motivated by the strong coupling between angular actuation and length in LIGO Hanford’s arms.

III. PROBLEM FORMULATION

The heart of the control noise problem is straightforward to outline: too much measurement noise gets passed through the controller and is actuated spuriously into the DARM degree of freedom, injecting noise. As with all LIGO systems, the goal is to minimize the noise that is fed into the readouts of the interferometer that contain information of astrophysical importance, mainly the DARM readout. The solution to this control problem is more difficult, as it must account for many different factors, including the plant’s response, the measurement noise passed through the control system, and the system’s stability and robustness.

The modern control optimization method that is well suited to this problem is linear-quadratic-Gaussian (LQG) control. An LQG controller is the optimal linear controller for a system that is disturbed by Gaussian white noise and has a quadratic (rms-like) cost function. Described in Secs. VIB and VIC, LQG controllers minimize a system’s $\mathcal{H}_2$ norm, which is a measure of the rms at system outputs from all of its noise inputs. The “ $\mathcal{H}$ ” stands for solutions within a Hardy space, indicating that the system must be closed-loop stable. This is necessary to have well-posed noise in the time domain and appropriate for implementable feedback controllers. The typical quadratic cost function for LQG is based on the total noise rms S_p of a degree of freedom LIGO’s performance relies on multiple criteria, so we will use a weighted rms that includes both the total noise and the lost detection range as figures of merit.

So far, the noise model of Sec. II A has not included the nonlinear/bilinear aspects. It has simply been a model of direct, linear noise coupling, which should naively only be optimized for the BNS range while ignoring the degree of

freedom rms. The following sections investigate how to incorporate a weighted rms to balance noise suppression against control noise injection.

A. System overview

When creating a SISO controller for a single degree of freedom, a MIMO system is constructed that takes into account the multiple inputs that can affect the degree of freedom, as well as multiple outputs to measure its performance. This system, whose block diagram is presented in Fig. 1, includes all of the information that is used to find the optimal controller, mainly of three types: the noise in the system, the plant of the system, and the FOM of the system. The first two of these pieces are rather straightforward and regularly measured at the LIGO sites, while the motivation for the FOM is given in Sec. III B.

The full system model, outlined by the dashed gray box in Fig. 1, includes four inputs, a control input u_c , a measurement noise input u_{n1} , a state noise input (environmental noise) u_{n2} , and an $\mathcal{H}_\infty$ disturbance input. Both noise inputs are white noise inputs that pass through shaping filters to model the measurement and environmental noise. The control input u_c is connected to the output of the controller and is not a white noise input. The $\mathcal{H}_\infty$ disturbance input u_∞ serves as a port for the solver to inject unshaped signals into the system and is a necessity for optimal control, as analyzed later. The system has two types of outputs: the figure of merit outputs and the measured output y_{meas} . y_{meas} gives the state of the plant with added measurement noise. Measurement noise is only observed by the controller and represents an imperfect estimate of the plant's true state.

The system layout includes multiple blocks, each representing a transfer function. The plant P represents the physical system being simulated; in this case, it is the mechanism by which the actuator converts force/torque inputs into displacements/tilts. The measurement noise M is the shaping filter that turns the input Gaussian white noise into the appropriately shaped sensing noise, while the environmental noise E serves the same purpose for the state noise on the plant, both shown in Fig. 2. The transfer functions relate to Eqs. (11), (13), and (14) through

$$S_{\text{env}} = |E|^2 S_{n1}^u \quad (15)$$

and

$$S_{\text{meas}} = |M|^2 S_{n2}^u, \quad (16)$$

where

$$S_{n1}^u, \quad S_{n2}^u \equiv 1/\sqrt{\text{Hz}} \quad (17)$$

and are uncorrelated. This indicates that these inputs are driven by Gaussian white noise that is then shaped by the relevant noise-shaping filters. The two figures of merit filters, F_{BNS} and F_{flat} , represent the transfer function that weights the frequency domain signal properly for two different measures of the controller's performance described in Sec. III B.




FIG. 2. The noise spectrum of the DHARD yaw degree of freedom at the LIGO Hanford Observatory decomposed into fits of measurement and environmental contributions. The spectra are fit to filters M and E , see relations of Eqs. (15)–(17). The measured open-loop (OL) amplitude spectral density (ASD) of y_{meas} is the measured spectrum of the cavity axis (see Fig. 1) and is equal to $|EP + M|$. This spectrum assumes no control, $K = 0$ (open loop). After fitting a ZPK to $|EP|$, the contribution from E was isolated by dividing by the model for the plant (P). The filters are using the Python package WIELD.IRRATIONAL to refine fits with gradient descent after making initial guesses using the AAA algorithm [10,11].

The two figures of merit outputs are free of any direct measurement noise, and they are each connected to a point representing the N_p and N_a plant and actuator noise terms of Eqs. (13), (14), and (11). The frequency response of the plant must have two copies in our model to produce the different contributions of K (to be designed) in the expressions of N_a and N_p .

The FOM block includes a shaping filter for each path that weights the relative contribution of S_a or S_p in the optimization problem. The construction and weighting of the FOM filters is described in Sec. III B and is fundamental to finding the Pareto front of optimal controllers for a given system. The absence of measurement noise indicates that the FOM outputs report the true value (up to weighting) of the degree of freedom from which it is constructed and not what would be observed through an imperfect physical measurement.

For this alignment control application, the state of the plant at N_p represents the total angular motion of the mirror or cavity in its degree of freedom (differential, HARD, yaw). A flat filter representing the first FOM then has y_{FOM1} , which integrates into the total variance of the mirror motion. This is the typical objective that is minimized in LQG/ $\mathcal{H}_2$ optimal control. The N_a connection to y_{FOM2} represents how the control actuation pushes on the mirror in pitch or yaw. If the beam is miscentered, the alignment actuation is converted into length noise. The BNS FOM then contains a weighing filter for how important this length noise conversion is for the SNR or range figure of merit.




FIG. 3. The magnitude of two figures of merit filters used in the augmented system, as shown in Fig. 1. Both of these figures of merit shape either N_p or N_a to create metrics for the controller to minimize. The broadband FOM results in an unshaped signal that must be below a specified threshold to keep the interferometer locked. The BNS FOM shapes S_a allowing for the direct optimization of LIGO's BNS detection range.

B. Figures of merit

As discussed in the previous sections, figures of merit are measures of the performance of a closed-loop system. In the example presented in this paper, an essential new complexity is added: multiple figures of merit.

In the model of the system, shown in Fig. 1, two figures of merit are shown, which have been indicated by Eqs. (6) and (10). Shown inside the red dashed outline in Fig. 1 and represented by transfer functions, these two figures of merit, F_{BNS} and F_{flat} , allow for different frequency weightings of the output signal. Figure 3 shows the magnitude of the two figures of merit used in the augmented ASC system. The first broadband FOM, F_{flat} , is required to minimize $\tilde{R}_p$ which is necessary to preserve the inequality in Eq. (10). The broadband FOM represents the unshaped total noise in the degree of freedom seen after the plant, which affects the operational locking of the interferometer. This FOM is also necessary in order for the problem to be well posed for optimal control, ensuring assumption 4 in Sec. VIC.

The second FOM, passing through F_{BNS} , is a measure of how much the alignment actuation influences the SNR or range of the LIGO detector to an inspiral of two $1.4M_\odot$ neutron stars at a distance of 40 Mpc, using Eqs. (6)–(8). The inspiral determines the function $h_{\text{gw}}(f)$. The factor $C(f)$ is a coupling factor for actuation-point control noise into DARM, computed using noise budget data from [3]. S_{det} is the detector's noise PSD. We fit the weighting filter so that

$$|F_{\text{BNS}}(f)| = 2 \frac{|C(f)||h_{\text{gw}}(f)|}{S_{\text{det}}(f)} \cdot \frac{d_{\text{gw}}}{16 \cdot 2.26 \cdot \text{SNR}_{\text{det}}}, \quad (18)$$

where the first term creates a weight filter to compute the SNR_{lost} , and the second term converts to lost range.

By connecting to the N_a actuation-point noise, the total variance of y_{FOM2} then produces the lost SNR or lost range. Practically, this filter (seen in Fig. 3) has a dynamic range going from low to high frequencies that spans more than the 16 orders of magnitude of a double-precision float, creating issues that we detail in the following sections.

The frequency-dependent coupling function $C(f)$ is measured through an excitation injected into the control loop [3]. When excited, S_a from Eq. (13) becomes

$$S_{a,\text{exc}} = S_a + \left| \frac{G}{1-G} \right|^2 S_{\text{exc}}, \quad (19)$$

where S_{exc} is the PSD of the excited signal summed in before the controller. When this excitement is added, S_h from Eq. (5) becomes

$$S_{h,\text{exc}} = S_{\text{det}} + |C|^2 \left(S_a + \left| \frac{G}{1-G} \right|^2 S_{\text{exc}} \right). \quad (20)$$

From Eqs. (19) and (20), $|C|^2$ can be calculated and fit from

$$|C|^2 = \frac{S_{h,\text{exc}} - S_h}{S_{a,\text{exc}} - S_a}. \quad (21)$$

1. Pareto front of multiple FOM for bilinear noise

To fully motivate the use of these two figures of merit in the bilinear noise case, Eq. (6) can be rewritten to directly depend on both figures of merit. Consider the noise coupling factor C from Eq. (8). For the multiplication effect that occurs for bilinear noise on a single degree of freedom, a model is assumed where $C \propto \tilde{R}_p$ represents an intermodulation effect. The effect of an intermodulation can be seen by rewriting Eq. (4) in the time domain, with the control noise term becoming bilinear

$$h_{\text{signal}}(t) = h_{\text{gw}}(t) + N_{\text{det}}(t) + \chi N_a(t) N_p(t), \quad (22)$$

where χ is a scale factor for the cross-coupling. For our alignment example, N_p is the residual angular motion of the mirror, while N_a is the actuated tilt. If the residual angular motion is nonzero, then the applied tilt causes a coupling to arm length.

When the N_p noise is dominated by low-frequency components, this term becomes $\chi \tilde{R}_p N_a(t)$ by linearizing. See Appendix A for a formal derivation. This substitution can then be used in Eq. (8) via $C(f) = \chi(f) \cdot \tilde{R}_p$. In the alignment control example problem, this relation for constructing $C(f)$ yields the coupling factor called “A2L” for alignment to length. This factor is measured [3] and translates angular motion in the alignment system into length motion in the DARM gravitational-wave channel. This form

for C and similar models of rms-dependent coupling factors indicates that one aspect of the range depends on the rms as computed from the flat FOM, while the other depends on the rms as computed from the shaped BNS FOM.

Assuming this form of the C coupling function, the SNR loss of Eq. (8) becomes

$$\text{SNR}_{\text{lost}}^2 = 4 \cdot \tilde{R}_p^2 \cdot \int_0^\infty \frac{|\chi|^2 |h_{\text{gw}}(f)|^2}{S_{\text{det}}^2(f)} S_c(f) df \quad (23)$$

$$= \underbrace{\int_0^\infty |F_{\text{flat}}|^2 S_p df}_{\tilde{R}_{\text{flat}}^2} \cdot \underbrace{\int_0^\infty |F_{\text{BNS}}|^2 S_c df}_{\tilde{R}_{\text{BNS}}^2 / \tilde{R}_{\text{flat}}^2}, \quad (24)$$

where $\tilde{R}_{\text{BNS}}^2$ denotes the rms of the noise that comes out of the BNS FOM, which is output y_{FOM2} in Fig. 1, $\tilde{R}_{\text{flat}}^2$ is the rms of the total plant motion during control, and $\tilde{R}_{\text{flat}}^2$ is a constant for rms at the time of the measurement of the A2L coupling factor $C(f)$.

As a minimization problem, a necessary condition for optimality can be established by taking the derivative with respect to the controller parameters $\mathbf{K}$, in bold to represent parameters in a state-space matrix form. Setting these derivatives to zero finds an extremal point of the noise as

$$0 = \frac{1}{4|\chi|^2} \frac{d}{d\mathbf{K}} \text{SNR}_{\text{lost}}^2 \quad (25)$$

$$= \tilde{R}_{\text{flat}}^2 \frac{d}{d\mathbf{K}} \tilde{R}_{\text{BNS}}^2 + \tilde{R}_{\text{BNS}}^2 \frac{d}{d\mathbf{K}} \tilde{R}_{\text{flat}}^2. \quad (26)$$

This relation can be satisfied through the equivalence

$$0 = \frac{d}{d\mathbf{K}} \tilde{R}_{\text{syn}}^2, \quad (27)$$

when

$$\zeta \equiv \frac{\tilde{R}_{\text{BNS}}}{\tilde{R}_{\text{flat}}}. \quad (28)$$

The LQG and mixed-sensitivity approaches below find control systems that meet this necessary condition. By scanning over the weight parameter ζ , the Pareto front of optimal controllers is produced. The above derivation indicates that the bilinear-optimal controller exists along this Pareto front.

There are multiple reasons why the parameter ζ should be abstracted as an independent term for Eq. (29). The main reason is that, prior to optimization, the ratio of the rms values of the two figures of merit cannot be known. Thus, there is no presumptive ζ to target. Since LQG can only optimize over a single parameter, all of a system's FOM must be combined. With ζ , we can define a synthetic weighted rms of the two figures of merit, which meaningfully relates to the above condition on the derivative of the SNR as

$$\tilde{R}_{\text{syn}}^2 = \zeta^2 \tilde{R}_{\text{BNS}}^2 + \tilde{R}_{\text{flat}}^2. \quad (29)$$

The weighted sum of the two figures of merit outputs creates a single synthetic output with its own rms, $\tilde{R}_{\text{syn}}$. The problem now boils down to a problem to minimize $\tilde{R}_{\text{syn}}$ over many given weightings of ζ . It is advantageous to only weight one of the figures of merit to achieve the desired relative weighting because leaving the weighting of the flat FOM untouched allows for consistent $\mathcal{H}_\infty$ norms across different weightings, which will be important in Sec. VA.

For values of ζ that are zero, or very large, the optimization of $\tilde{R}_{\text{syn}}$ will find the minimum rms in one of the respective figures of merit. Because the rms values satisfy a minimization constraint, the optimized values of both $\tilde{R}_{\text{BNS}}$ and $\tilde{R}_{\text{flat}}$ are expected to be monotonic but not necessarily smooth or convex in ζ .

Together, this suggests that scanning ζ will allow one to find the global minimum of the lost SNR, but does not establish that the control-optimal SNR loss function is a convex optimization problem in ζ . For any more complex dependence of C on χ or $\tilde{R}_c$, the derivative argument above is expected to still apply, but the relationship of Eq. (27) will become more complex. Finally, the above derivation and example based on the C coupling parameter assume only a single degree of freedom. More elaborate models of the C parameter are likely dependent on the rms of multiple degree of freedom. Thus, the decision point of where to balance flat rms vs the BNS-weighted rms is, in general, more complex. These are major reasons that ζ should be treated abstractly to produce a Pareto front. The following plots show that the ζ parametrization clearly shows the boundary of the optimal space for the two major rms contributions. Thus, scanning ζ to find a boundary that shows the specific trade-off between the flat-weighted rms and the BNS-weighted rms is the most generally useful expression and is achieved through the ζ parameter. We believe the above argument can apply to MIMO systems using a vector of ζ scales, but leave this for future work.

IV. LQG OPTIMIZATION

The LQG problem is to find an optimal controller that yields the lowest $\tilde{R}_{\text{syn}}$ for a given ζ . This is equivalent to finding a controller that minimizes the $\mathcal{H}_2$ norm of a “transfer function matrix” $\mathbf{T}_2(f)$ from the u_{n1}, u_{n2} vector of inputs to the $y_{\text{FOM1}}, y_{\text{FOM2}}$ vector of outputs, with a ζ scaling first applied to the y_{FOM2} output. Here $\mathbf{T}$ is a set of transfer functions that define the system from Fig. 1. A mathematical description of this system using state-space systems and a more rigorous definition of $\mathbf{T}$ can be found in Sec. VIA.

By contrast, LQG is often taught and presented starting from the separation principle, combining a Kalman filter as a noise-optimal state linear-quadratic estimator with a linear-quadratic regulator that minimizes an actuator cost function. The actuator cost is often described as fuel consumption, but when using the FOM weighting filters, it can also be interpreted as part of a noise-minimization procedure. This interpretation of LQG as an $\mathcal{H}_2$ noise-minimization

procedure can be established intrinsically through minimization [12], as well as through the case analysis of the separation principle [13].

The $\mathcal{H}_2$ optimization can prove difficult by hand or even with a gradient-descent numerical optimizer because of design trade-offs. Finding the perfect balance between suppressing environmental noise and introducing minimal measurement noise, all while weighting the output by two figures of merit, is a job that optimal control can perform in a single algebraic solution step. Here, LQG is used to find the optimal controller to give the lowest possible $\tilde{R}_{\text{syn}}$. The control transfer function K is determined by the state-space description of the control $\mathbf{K}_{\text{LQG}}$, using the method described in Sec. VIC. The method of the numerical solver is not essential to understanding the results of this paper, but an important point is that LQG optimization is non-iterative. This is essential for speed, as our method requires a scan over the ζ parameter that indicates the trade-off between two figures of merit. Because of dynamic range and numerical challenges, an understanding of the numerical solver has been essential to implement this work with a reliable software implementation, but ideally, it should not be required for end users.

A. LQG optimal controllers

Using the plant system, whose layout is presented in Fig. 1, a set of optimal LQG controller filters was calculated

using the method described in Sec. VIC. These controllers, corresponding to the filter K in Fig. 1, were calculated for different weights of the two figures of merit corresponding to the BNS range and total rms for a range of ζ values. The change in the weighting ζ causes the BNS FOM to increase its contribution to the $\tilde{R}_{\text{syn}}$. By calculating a number of the ζ values, it is possible to show the performance of those controllers in relation to the LIGO's early observing run O4 differential-hard-yaw open-loop gain with respect to any trade-off between total noise and lost performance (observing range). Since the solver is always finding the lowest noise across different weightings, the results form a Pareto front, showing the boundary of all optimal linear controllers for this cavity-alignment degree of freedom.

Figure 4 shows a comparison between several LQG optimal controllers for several ζ values. To calculate the controllers, a system scaled by ζ is used. During the solve, if ζ is too large, the solver will either fail or create suboptimal controllers due to numerical error in eigenvalue computation and loss of the ability to order deflating subspaces. Once the controller is calculated, it is connected to an unscaled system, where $\zeta = 1$. The unscaled system models the controllers' performance, as a Lyapunov equation can compute the lost BNS range by computing the rms of the BNS FOM output. The two performance metrics are then plotted. Lower is better, representing either suppressed noise in the alignment plant or less loss of detection range. The LQG open-loop gain of the controllers is shown in




FIG. 4. The rms noise at the two figures of merit outputs given the white noise disturbances for the closed-loop system, including the system shown in Fig. 1 and optimal LQG controllers calculated for specified ζ values. Each point represents the performance of a single controller that is the LQG optimal solution for the plant with a different FOM weighting. The rms value is shown on the axis of the corresponding FOM. The BNS FOM's rms is calibrated to represent BNS range lost by the noise from a given controller, while the flat FOM represents the rms noise in radians for ASC DHARD yaw. The number next to each point represents the ζ value used to acquire the desired weighting. The color of each point corresponds to the phase margin of the controller. The triangle represents the current LIGO hand-tuned controller's performance and phase margin. All of the LQG optimal controllers whose performance beats the hand-tuned controller in both figures of merit have poor phase margins, and they get worse as the BNS range FOM is emphasized using a higher ζ factor. The margins can be improved by the LQG optimal solutions with an $\mathcal{H}_\infty$ performance bound discussed in Sec. V and presented in Fig. 7.




FIG. 5. The open-loop gain G of the unscaled plant and controller system shown for different values of ζ . As ζ gets lower, the optimizer gives less weight to the $\hat{R}_{\text{BNS}}$. A hand-designed LIGO DHARD yaw controller is included for comparison. Each of these LQG controllers was optimized using an LQG solver with no enforced stability margins. The hand-tuned controller meets the stability requirements necessary for implementation in the interferometer.

comparison to a hand-tuned LIGO loop design for DHARD yaw used during LIGO's fourth observing run.

The open-loop gain of the plant and controller system for the same LQG controllers is shown in Fig. 5. There are several features and similarities demonstrated in both the optimal LQG controllers and the LIGO hand-tuned controller.

The closed-loop gain of $G/(1 - G)$, shown in Fig. 6, is the filter through which the noise S_{meas} is seen in S_a and S_h as described by Eq. (11). When the peak of $G/(1 - G)$ is

high, the stability margins are poor, which is visible in Fig. 6. This suggests that if one can constrain the height, then one can enforce more acceptable margins.

The stability margins are computed for the open-loop gain G of the LQG controllers. The phase margin, which is related to the peak of $G/(1 - G)$, is presented as the color of the points in Fig. 4. All of the gain and phase stability margins for the LQG controllers are unacceptably low, below 10° , where 30° is usually taken to be a lower limit. LIGO's control systems must be robust and remain stable




FIG. 6. The closed-loop gain $G/(1 - G)$ of the plant and controller system shown for different values of the relative FOM weighting parameter ζ . These are the same LQG controllers presented in Figs. 7 and 5. The current hand-tuned LIGO DHARD yaw controller is included for comparison. The large peaks seen around 10 Hz are gain peaking, which causes low stability margins by inflating the $\mathcal{H}_\infty$ norm of the closed-loop system. This is further discussed in Sec. V.

when dealing with drifting operating conditions of the interferometer. If the gain of the optical or mechanical components of the interferometer plant is increased or decreased by a small margin, then the gain of G will also be affected. If the margins are smaller than the change in the gain or phase of the plant, then the system will become unstable.

B. Properties and limitations apparent in LQG controllers

All of the calculated controllers aim to reduce the noise of the current hand-tuned LIGO DHARD yaw controller. Figures 4–6 show the rms and open- and closed-loop transfer functions for the LQG controllers and for the hand-tuned LIGO controller. The properties exhibited by the LQG controllers should be related to the hand-tuned controller, but might also inform future design.

The pieces of information that motivate the hand tuning of the controllers are generally the rms of S_p and the open- and closed-loop transfer function to see the stability margins, as well as the noise spectrum presented in Fig. 2. The hand-tuned controller has an integrator under 1 Hz to suppress the most significant contribution to the total rms: tilt noise. The unity-gain frequency of the open-loop gain is below 10 Hz, and the hand-tuned controller rolls off rapidly above 10 Hz to decrease the impact on the observing range. The roll-off enforces a hand-tuned version of the F_{BNS} weighting ζ , but costs a lot of phase at the unity-gain frequency. The hand-tuned controllers sometimes are designed with multiple unity-gain frequency crossings to target noise peaks where needed, but these are challenging to keep stable, so hand design often just suppresses the lowest-frequency noise.

The LQG controllers have a much more aggressive open-loop gain. There is a strong roll-off after 10 Hz to minimize the impact on the BNS range, meaning the rate of the roll-off is highly dependent on ζ . When ζ is very low and thus the F_{BNS} does not contribute much to the total rms of $\tilde{R}_{\text{syn}}$, the roll-off is almost nonexistent. The limit on when the roll-off becomes non-negligible is when $\zeta\tilde{R}_{\text{BNS}}$ becomes larger than the residual environmental noise $\tilde{R}_p$. One of the most noticeable characteristics of the LQG controllers is that they are aggressive. The closed-loop gain, shown in Fig. 6, shows that the unchecked gain peaking can get to about 70 times its dc gain, which corresponds to an unacceptable phase margin of around a few degrees.

V. SOLUTION WITH CUSTOM STABILITY MARGINS

While the goal of many of LIGO’s control loops is to minimize the $\mathcal{H}_2$ norm of the closed-loop system, optimally, for bilinear noise, the controllers must also be robust and stable. When hand-designing control loops, the measures of this robustness are the gain and phase margins. The gain margin is the smallest gain shift that makes the system

unstable, and the phase margin is the smallest amount of phase lag that causes the system to become unstable.

The pure LQG solutions for the system, presented in Sec. IV A, show unacceptable stability margins and gain peaking. The best controller for many of LIGO’s applications would be an LQG controller if there were a way to guarantee its stability margins. Unfortunately, to quote the entire abstract of Doyle’s paper on guaranteed margins for LQG controllers, “There are none” [14].

We must instead employ a more complicated strategy. The phase margin for the system is directly linked to the peak of the closed-loop gain. This peak can be measured using an $\mathcal{H}_\infty$ norm and can be bounded. The following work mixes such a bound within the LQG solution and presents robust and noise-optimal controllers for our example alignment (ASC) system and can be used for many other LIGO control systems.

A. $\mathcal{H}_\infty$ optimization problems

Before describing the mixed-sensitivity approach, it is worth first describing the pure $\mathcal{H}_\infty$ optimization approach and contrasting it to the $\mathcal{H}_2$ minimization performed by LQG. In changing an LQG problem to utilize an $\mathcal{H}_\infty$ -optimal solver, the same transfer function $\mathbf{T}_2(f)$ of noise inputs to FOM outputs can be used. The optimization objective is then to find the controller that minimizes the maximum value of $\|\mathbf{T}_2(f)\|_2$ over all frequencies f (cf. Sec. VI F). $\mathcal{H}_\infty$ optimization is typically claimed to find controllers more robust to model changes or adversarial (optimally bad) inputs; however, it requires its own precise and uniquely defined FOM to be applicable or relatable to astrophysics metrics of gravitational-wave detectors. In particular, $\mathcal{H}_\infty$ makes no guarantees about minimizing rms noise levels and is thus fundamentally unrelated to SNR and detection range metrics.

In effect, when using noise-shaping filters and FOM weights, a simple reinterpretation of an $\mathcal{H}_2$ model into $\mathcal{H}_\infty$ neither improves robustness nor helps with the true noise optimization problem. One must instead use a different set of weights and shapes to define a $\mathbf{T}_\infty$ transfer function matrix that is dedicated to the $\mathcal{H}_\infty$ optimization. The following sections show the mixed-sensitivity approach that retains the $\mathbf{T}_2(f)$ of the $\mathcal{H}_2$ optimization while including a dedicated weighting to apply an $\mathcal{H}_\infty$ bound. In the following, $\mathbf{T}_\infty$ is implemented by utilizing the additional input: u_∞ of Fig. 1.

What, then, is the difference between an $\mathcal{H}_\infty$ optimization and a bound? Both applications of $\mathcal{H}_\infty$ use an additional single parameter, γ , to apply the bounded real lemma [15] and produce algebraic Riccati equations that only have an admissible solution when $\|\mathbf{T}_\infty(f)\|_2 < \gamma$ for all f . The $\mathcal{H}_\infty$ optimization problem is to find the smallest γ , and associated controller, such that these equations hold. In finding the minimal admissible γ , there is no remaining freedom in the controller to also optimize an LQG $\mathcal{H}_2$ problem and minimize noise rms due to the uniqueness of the optimal $\mathcal{H}_\infty$ controller. When a γ above the minimum

value is chosen, there is a large class of $\mathcal{H}_\infty$ -suboptimal controllers [16], all related through a γ bounded linear fraction transformation [15], that can be considered as the admissible set for an $\mathcal{H}_2$ noise optimization problem. In this way, it is possible to find a solution with a minimal $\mathcal{H}_2$ norm within the set of $\mathcal{H}_\infty$ -suboptimal controllers.

B. LQG optimal solutions given an $\mathcal{H}_\infty$ performance bound

As discussed in Sec. VA, placing a limit on a suitably defined $\mathcal{H}_\infty$ norm can enforce stability margins of a system. However, $\mathcal{H}_\infty$ solvers do not return controllers with acceptable noise. The ideal controller for most of LIGO's control loops would be the optimal LQG filter that has an $\mathcal{H}_\infty$ norm enforcing robustness. Here, a bound on the phase margin is derived, and results from the alignment control model are shown. Section VII describes the numerical method for calculating the optimal solution [17,18].

The phase margin ϕ_m of the system is expressed as the phase when the magnitude crosses the unity-gain frequency or expressed as an equation,

$$\phi_m = \angle G(f_u), \quad (30)$$

where f_u is the unity-gain frequency. As the unity-gain frequency is approached

$$\lim_{f \rightarrow f_u} |G(f)| = 1, \quad (31)$$

meaning, at f_u , the closed-loop gain is at or near its maximum. If a bound is added to the closed-loop gain and

the phase margin is considered in radians, then

$$\left| \frac{G(f)}{1 - G(f)} \right| \leq \gamma \quad (32)$$

and

$$\lim_{f \rightarrow f_u} \left| \frac{G(f)}{1 - G(f)} \right| \approx \frac{1}{|\phi_m|}, \quad (33)$$

yielding

$$|\phi_m| \geq \frac{1}{\gamma}. \quad (34)$$

By limiting how large γ can be through the controller, the phase margin is also limited. For example, $\gamma \approx 1.27$ limit enforces a $\pi/4 = 45^\circ$ phase margin.

To implement this, the additional u_∞ input of Fig. 1 is employed. By using the transfer function matrix from the input u_∞ to the y_{FOMI} output to define $\mathbf{T}_\infty(f)$, the ideal $\mathcal{H}_\infty$ bound can be established in the definitions above.

The results, presented in Figs. 7–9, show a number of controllers of varying ζ and γ values. Shifting these two parameters allows a group of controllers to be calculated. A single preferred controller can then be postselected from the group to meet the desired outcome. These plots look similar to those in Sec. IV, but include a varying γ bound as a parameter.

Interestingly, our configuration of $\mathbf{T}_\infty(f)$ to measure the closed-loop gain allows even $\gamma = 0$ to be admissible, which




FIG. 7. The rms plot that shows the rms motion of the optic in radians and the range lost from added noise, which are the outputs from the two figures of merit. The plot shows the phase margin of all the controllers as a color-coded representation. The phase margin is related to the gain margin monotonically, and their relation as well as their relation to γ is presented in Fig. 10. Each branch of the plot, shown by the light dotted line, has the same ζ value with different γ values. The number at the top of each branch is the ζ value for that branch. The dark dashed line shows the $\mathcal{H}_2$ Pareto front, which is the limit to the controller's noise performance, as it represents the performance of the LQG controllers from Fig. 4, which are rms optimal.




FIG. 8. The open-loop gain G of controllers from Fig. 7 with a constant ζ value but with varying γ values. As γ decreases, the robustness of the controller, as measured by the stability margins, increases. These calculated controllers are compared to the current LIGO controller for DHARD yaw.




FIG. 9. The closed-loop gain $G/(1-G)$ of controllers from Fig. 7 with a constant ζ value but with varying γ values. As γ decreases, the robustness of the controllers increases. The phase margin of the controller is directly linked to the height of the maximum peak of the closed-loop gain. As γ decreases, this peak is suppressed, as shown here, and thus, the stability margins increase. This relation is shown in Fig. 10. These calculated controllers are compared to the closed-loop gain of the current LIGO controller for DHARD yaw.

occurs when $G = 0$, everywhere. In practice, no interesting control loops that influence or optimize the rms can be created as $\gamma \rightarrow 1$. Numerically, our solver becomes increasingly fragile in this limit. Our procedure is to first choose a moderately large γ value, and then solve for a controller. If this solution is successful, then it chooses a smaller value and tries again. γ is successively reduced toward $\gamma = 1$ until it cannot reliably converge to a solution. For this ASC example, phase margins of around 45° are achieved.

The controllers' rms performance, presented in Fig. 7, shows that the current LIGO controller for DHARD yaw

can be improved in both the BNS range as well as the overall rms in the loop while keeping the same phase margin. In the lost range due to the added noise, an order-of-magnitude improvement can be made. The total rms can only be improved by about four times for controllers that would be implemented while maintaining the controller's stability.

Each group of controllers in Fig. 7 that is connected by a light dotted line represents controllers with the same ζ , displayed at the top of each branch, for a varying γ . The phase margin is shown by the color of each point. It is clear

that, as the phase margin increases, so does the rms noise. In effect, every choice of phase margin produces a Pareto front.

The open-loop gain of a single branch of these constant ζ controllers is presented in Fig. 8 and compared to the current LIGO DHARD yaw controller. By plotting only one ζ value, it is possible to show the effect that γ has on the controllers. As γ is decreased, the controller becomes more robust, and the overall amplitude decreases, especially at low frequencies, making the controllers less aggressive. This is also true for the other controllers at other ζ values. However, at around 7 Hz, the response stays almost constant for the varying γ values. The fact that the controller is trying to stay at or above unity gain shows that it is not merely pushing down the gain everywhere but actually trying to maintain the noise performance while improving stability. In comparison to the non- $\mathcal{H}_\infty$ norm limited controllers in Fig. 5, these are much less aggressive at low frequencies. Between 0.1 and 1 Hz, the gain of the optimal controllers is still sometimes significantly larger than the current LIGO controller. In Fig. 2, the environmental noise shaped by the plant is shown to peak around 0.15 Hz. The controller's aggressiveness is trying to suppress the environmental noise as seen at the plant output.

The closed-loop transfer function of $G/(1 - G)$ for the system is shown in Fig. 9. The $\mathcal{H}_\infty$ norm for each closed-loop system is the maximum of the closed-loop magnitude. Since the $\mathcal{H}_\infty$ norm must always be less than γ , as γ is decreased, the peak of the magnitude of the closed-loop transfer function is forced down. As shown in Eq. (34), this sets a limit on the phase margin.

While the phase margin is directly related to γ , the gain margin is an important element of the controller's stability. Figure 10 shows that the gain margin and phase margin are monotonic. By putting a bound on γ , it is possible to also create a controller with both acceptable gain and phase margins.

C. Properties and limitations of $\mathcal{H}_\infty$ norm limited LQG controllers

The first and most important change from the pure LQG controllers to the controllers presented in Fig. 7 is the improved phase margin. At the lowest γ values, the phase margin can beat the margins of the current hand-tuned controller.

As seen in the closed-loop transfer function, the $\mathcal{H}_\infty$ norm is being suppressed as the γ value decreases, enforcing an upper limit. The limiting of the $\mathcal{H}_\infty$ norm is not seen in the plot of the closed-loop gain of the pure LQG controllers in Fig. 6, since there is no $\mathcal{H}_\infty$ boundary enforced. The enforcement of the limit directly leads to more robust controllers as measured by the gain margin and phase margin.

While the LQG results in Sec. IV A show controllers that do not show all of the same features of the current hand-tuned LIGO controller, the $\mathcal{H}_\infty$ bounded controllers show many of the same features. In the open-loop transfer function (Fig. 8), the roll-off of the hand-tuned and $\mathcal{H}_\infty$ bounded controllers is very similar in frequency and magnitude. This similarity appears even though $\mathcal{H}_\infty$ bounded controllers' rms is calculated from N_a rather than u_c , as is often used in LIGO's controller optimization. This similarity may suggest that this distinction is not as important. At very low frequencies, the $\mathcal{H}_\infty$ bounded controllers are actually slightly less aggressive than the hand-tuned controller. As γ decreases, the $\mathcal{H}_\infty$ bounded controllers get less aggressive.

A notable difference between the current hand-tuned LIGO controller and the calculated $\mathcal{H}_\infty$ bounded controllers is their behavior around unity gain ($G = 1$). The hand-tuned controller causes many unity-gain crossings in the open-loop gain as the gain peaks over the line to target noise peaks. Only the $\mathcal{H}_\infty$ bounded controllers with the lowest γ values exhibit this behavior.




FIG. 10. The gain margin, phase margin, and γ for the LQG controllers with $\mathcal{H}_\infty$ performance bounds displayed in Fig. 7. The phase margin, gain margin, and γ are monotonic. The stability margins are also shown for the current LIGO controller for DHARD yaw, which has no γ value since it was hand-tuned and not created by an $\mathcal{H}_\infty$ minimization algorithm.

It is important to note that the solver employed does not exactly implement the gain-peaking constraints that enforce phase margin. Our implementation is slightly more restrictive. Current numerical solution methods for the mixed-sensitivity problem cannot fully independently define $\mathbf{T}_2$ and $\mathbf{T}_\infty$. Therefore, our method also has to include the y_{FOM2} output, which instead implements the constraint

$$\left| \frac{G(f)}{1 - G(f)} \right|^2 \leq \frac{\gamma^2}{|F_{\text{flat}}(f)|^2 + \zeta^2 |F_{\text{BNS}}(f)|^2}. \quad (35)$$

This could be corrected using a compensating weighting filter on the u_∞ input, but this increases the number of system states, requires more computation, and requires a more advanced set of numerical equations. The existing results are satisfactory, though slightly more constrained than necessary to enforce a stability requirement.

VI. RICCATI EQUATION APPROACH TO MIXED-SENSITIVITY CONTROL

This paper, so far, presents results and explores the LQG, $\mathcal{H}_\infty$, and mixed LQG/ $\mathcal{H}_\infty$ control problem from the view of classical control. To numerically solve these problems, the tools and methods of modern control must be employed. This section explains how the above results were obtained in a way that is useful to a reader trying to replicate or further develop the techniques. This section is not essential to understand the results given above. Notably, additional transformations on the system in Fig. 1 are required to meet the structure and conditions of an LQG problem.

We start with the definitions and rules for establishing a well-posed LQG problem, relating them to the physical requirements and necessary arrangements on the state diagrams. Additional considerations are required to enforce and solve the $\mathcal{H}_\infty$ bounds. The system is then extended to the mixed-sensitivity problem.

A. State-space system description

To define an LQG problem, the “system model” component of Fig. 1 must be assembled into a monolithic state space. All of the component subsystems are individually modeled or fit from data and converted to state-space form. They are then connected according to the diagram using composition rules, and the entire system, excluding the controller, is put into the form

$$\dot{\vec{x}} = \mathbf{A}\vec{x} + \mathbf{B}_1\vec{u}_i + \mathbf{B}_2\vec{u}_c \quad (36)$$

$$\vec{y}_{\text{FOM}} = \mathbf{C}_1\vec{x} + \mathbf{D}_{12}\vec{u}_c \quad (37)$$

$$\vec{y}_{\text{meas}} = \mathbf{C}_2\vec{x} + \mathbf{D}_{21}\vec{u}_i + \mathbf{D}_{22}\vec{u}_c \quad (38)$$

where $\mathbf{A}$ is the $n \times n$ state matrix, $\mathbf{B}_j$ is the $n \times m_j$ input matrix, $\mathbf{C}_i$ is the $p_i \times n$ sensing matrix, $\mathbf{D}_{ij}$ is the $p_i \times m_j$

feedthrough matrix, $\vec{x}$ is the state vector, $\vec{u}_c$ is the control input, $\vec{u}_i$ is the vector of inputs for simulating noise in your system, $\vec{y}_{\text{FOM}}$ is the vector of FOM outputs, and $\vec{y}_{\text{meas}}$ is the measured output. Using $s = 2\pi jf$, the system transfer function matrix $\mathbf{T}$ is given by

$$\mathbf{T}(f) = \begin{bmatrix} \mathbf{A} & \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{C}_1 & 0 & \mathbf{D}_{12} \\ \mathbf{C}_2 & \mathbf{D}_{21} & \mathbf{D}_{22} \end{bmatrix} \quad (39)$$

$$\equiv \begin{bmatrix} \mathbf{C}_1 \\ \mathbf{C}_2 \end{bmatrix} (s - \mathbf{A})^{-1} \begin{bmatrix} \mathbf{B}_1 & \mathbf{B}_2 \end{bmatrix} + \begin{bmatrix} 0 & \mathbf{D}_{12} \\ \mathbf{D}_{21} & \mathbf{D}_{22} \end{bmatrix}. \quad (40)$$

It is important to note that while the system described by $\mathbf{T}$ is often called the “plant” in control literature, in this work it will continue to be called the system model so as to differentiate it from the plant in Fig. 1. Given a $\mathbf{T}$ in the proper form, it is possible to calculate a controller $\mathbf{K}$ that minimizes the $\mathcal{H}_2$ norm of the combined system shown in Fig. 11 [13,19].

In the following, the inputs and outputs of the $\mathbf{T}_2$ and $\mathbf{T}_\infty$ components of this system will need to be distinguished. Those can be distinguished by segmenting the noise inputs into the white noise and infinity contributions, using

$$\mathbf{B}_1 = \begin{bmatrix} \mathbf{B}_{1n} & \mathbf{B}_{1\infty} \end{bmatrix}, \quad (41)$$

$$\mathbf{D}_{21} = \begin{bmatrix} \mathbf{D}_{21w} & \mathbf{D}_{21\infty} \end{bmatrix}, \quad (42)$$

and

$$\vec{u}_i = \begin{bmatrix} \vec{u}_n \\ u_\infty \end{bmatrix}. \quad (43)$$




FIG. 11. A block diagram depicting the connection of the controller to the plant and abstracting away the internal structure of other block diagrams, such as Fig. 1. The transfer function from $\vec{u}_n$ to $\vec{y}_{\text{FOM}}$ is $\mathbf{T}_2$, discussed in Sec. VI A. The transfer function from $\vec{u}_\infty$ to $\vec{y}_{\text{FOM}}$ is referred to as $\mathbf{T}_\infty$. We depict this abstraction to clarify that transformations which preserve the T_2 or T_∞ response do not affect the optimal closed-loop response, but may influence the optimal $\mathbf{K}$, requiring it to be mapped through inverse transformations.

B. Calculating the $\mathcal{H}_2$ norm

With the state-space representation of the system, we can use it to calculate the total rms noise given any controller. Along with evaluating the transfer function equations of Eq. (40) to construct the frequency-resolved noise budget, we can use matrix mathematics from numerical Lyapunov solvers to efficiently calculate the total rms noise without numerically evaluating integrals [20]. This technique is used to construct the plots of Figs. 4 and 7.

The $\mathcal{H}_2$ norm of the system is a measure of the rms at $\vec{y}_{\text{FOM}}$ output from the white noise input $\vec{u}_n$. For a continuous system with a transfer function $\mathbf{T}_2(s)$, the $\mathcal{H}_2$ norm is given by

$$\|\mathcal{H}_2\| = \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{trace}[\mathbf{T}_2(j\omega)^\dagger \mathbf{T}_2(j\omega)] d\omega}. \quad (44)$$

The $\mathcal{H}_2$ norm is calculated by solving the Lyapunov equation given by

$$\mathbf{A}^T L_o + L_o \mathbf{A} + \mathbf{C}^T \mathbf{C} = \mathbf{A} L_c + L_c \mathbf{A}^T + \mathbf{B} \mathbf{B}^T = 0, \quad (45)$$

where L_o is the observability Gramian of $(\mathbf{C}, \mathbf{A})$, and L_c is the controllability Gramian of $(\mathbf{A}, \mathbf{B})$. Both are determined by numerically solving the Lyapunov equations. In this expression, the matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ are the matrices of the $\mathbf{T}_2$ system whose $\mathcal{H}_2$ norm is being measured (usually with a controller solved and attached). The $\mathcal{H}_2$ norm can then be calculated by

$$\|\mathcal{H}_2\| = \sqrt{\text{trace}[\mathbf{B}^T L_o \mathbf{B}]} = \sqrt{\text{trace}[\mathbf{C} L_c \mathbf{C}^T]}. \quad (46)$$

Note that these Lyapunov equation methods apply only when $\mathbf{A}$ is stable, but that condition is guaranteed for the closed-loop systems returned by LQG-based solvers. Allowing $\mathbf{A}$ to be unstable changes the interpretation to an $\mathcal{L}_2$ norm calculation and can be solved using the SLICOT function AB13BD [21]. We have noticed that standard Lyapunov equation solvers sometimes fail to produce accurate noise estimates due to the dynamical range of the transfer functions involved. More advanced Lyapunov solvers that compute the “square root” of the noise covariance matrices using Hammarling’s method and return Cholesky decompositions are more numerically reliable [22,23].

The use of Lyapunov equations to compute the $\mathcal{H}_2$ norm is foundational to LQG theory and can be used to derive the following Riccati equations for optimal control [12]. While this is widely known in the control community and textbooks, we include this here for a physics audience.

C. Optimal LQG controllers

Here, we outline some of the major points of the LQG numerical solution process, as it will inform the mixed-sensitivity method to follow.

The optimal LQG controller, $\mathbf{K}_{\text{LQG}}$, is found by solving two continuous algebraic Riccati equations,

$$\begin{aligned} \mathbf{A}^T \bar{\mathbf{X}} + \bar{\mathbf{X}} \mathbf{A} - (\bar{\mathbf{X}} \mathbf{B}_2 + \mathbf{C}_1^T \mathbf{D}_{12})(\mathbf{D}_{12}^T \mathbf{D}_{12})^{-1} \\ (\mathbf{B}_2^T \bar{\mathbf{X}} + \mathbf{D}_{12}^T \mathbf{C}_1) + \mathbf{C}_1^T \mathbf{C}_1 = 0, \end{aligned} \quad (47)$$

and

$$\begin{aligned} \mathbf{A} \bar{\mathbf{Y}} + \bar{\mathbf{Y}} \mathbf{A}^T - (\bar{\mathbf{Y}} \mathbf{C}_2^T + \mathbf{B}_1 \mathbf{D}_{21}^T)(\mathbf{D}_{21} \mathbf{D}_{21}^T)^{-1} \\ (\mathbf{C}_2 \bar{\mathbf{Y}} + \mathbf{D}_{21} \mathbf{B}_1^T) + \mathbf{B}_1 \mathbf{B}_1^T = 0. \end{aligned} \quad (48)$$

Each of these Riccati equations corresponds to the Lyapunov equations of the prior section. The $\bar{\mathbf{Y}}$ is the observability Gramian for the optimally controlled full-state feedback system, where the estimator/controller is unknown but assumed to be a Kalman filter. It is also a positive definite symmetric matrix. As the sensing noise becomes smaller, the Kalman gain increases and $\bar{\mathbf{Y}}$ becomes increasingly positive semidefinite, eventually unable to be resolved numerically. Analogously, $\bar{\mathbf{X}}$ is a positive (semi)definite matrix and the controllability Gramian of the closed-loop system.

The stabilizing solutions, defined below, $\bar{\mathbf{X}} \geq 0$ and $\bar{\mathbf{Y}} \geq 0$, are then used to calculate the Kalman gain K_s and the regulator feedback F_s , using

$$\mathbf{F}_s \equiv -(\mathbf{D}_{12}^T \mathbf{D}_{12})^{-1} (\mathbf{B}_2^T \bar{\mathbf{X}} + \mathbf{D}_{12}^T \mathbf{C}_1) \quad (49)$$

and

$$\mathbf{K}_s \equiv -(\mathbf{D}_{21} \mathbf{D}_{21}^T)^{-1} (\mathbf{C}_2 \bar{\mathbf{Y}} + \mathbf{D}_{21} \mathbf{B}_1^T). \quad (50)$$

The strictly proper controller that solves the $\mathcal{H}_2$ -control problem is given by

$$\bar{\mathbf{K}} = \left[\begin{array}{c|c} \mathbf{A} + \mathbf{B}_2 \mathbf{F}_s + \mathbf{K}_s^T \mathbf{C}_2 & \mathbf{K}_s^T \\ \hline -\mathbf{F}_s & 0 \end{array} \right]. \quad (51)$$

Note that the above equations do not use $\mathbf{D}_{22}$, and assume it to be zero. One can remove or add the D_{22} using a feedforward term on $\mathbf{L}$, which can be reinterpreted (and numerically implemented) as a feedback term on $\bar{\mathbf{K}}$. This amounts to a linear canonical transformation to find the final optimal LQG controller $\mathbf{K}_{\text{LQG}}$ of the as-given system,

$$\mathbf{K}_{\text{LQG}} = \bar{\mathbf{K}} (\mathbf{I}_{p_2} + \mathbf{D}_{22} \bar{\mathbf{K}})^{-1}, \quad (52)$$

where $\mathbf{I}$ is the identity matrix of dimension p_2 .

The existence and numerical solvability of these equations depend on many results in control theory and mathematical analysis. We emphasize two that are relevant to discussing our mixed-sensitivity approach.

The first is the separability of the solutions. The controller $\bar{\mathbf{K}}$ is computed from two independent subproblems. The two independent problems are the linear-quadratic

observer (a Kalman filter) and the linear-quadratic regulator. Each problem has stability guarantees for the limiting case of state feedback or state estimation, yet these combine to form a single control system that is itself stable and that stabilizes the system under feedback. This separability is a major reason that the solution is noniterative and can be proven to be globally optimal.

The second key result we emphasize is the development of deflating subspace methods to solve algebraic Riccati equations [24,25]. In short, these methods transform either of the above equations into a single matrix representing a particular system of equations. The matrix is transformed into a Schur form, upper triangular with its eigenvalues along the diagonal. During the conversion to Schur form, a modified QR or QZ algorithm reorders the stable eigenvalues first along the diagonal. Once in this form, the system of equations can be solved using a column of the orthogonal transformation matrix, the deflating subspace. In the original formulation of Laub, the matrix is “Hamiltonian,” a form that ensures that half of the eigenvalues are stable. That form can be generalized [26] to solve directly for the feedback control as well. The choice of eigenvalues represents the poles of the closed-loop system, ensuring stability under optimal control. Under the conditions of the following section, the Gramians $\bar{X}$ and $\bar{Y}$ are ensured to exist and to be positive definite, which in turn ensures that the controller, as an independent system, is itself stable.

We mention these two as they are essential for the solution of the mixed-sensitivity problem and for further improving numerical robustness for high dynamic range gravitational-wave system models. For example, the `ordqz` algorithm that is essential to the SciPy implementation of the deflating subspace method, from which we base our implementations, greatly benefits from matrix rescaling [27,28]. This has led us toward some custom (re)implementations of these solvers and a desire to test against quad-precision linear algebra implementations, where the physical dynamic range is a substantially smaller fraction of the floating-point epsilon. We leave this for future work and note that several libraries exist in the Julia language, which may make such a future effort possible. In our testing, we have not found that iterative refinement improves upon the reliability of a deflating subspace solver [29].

D. Mathematical assumptions and their physical interpretation

In order to calculate the optimal controllers for the system given in Eqs. (36)–(40), the following assumptions must be met to ensure that the problem is well posed. These are presented mathematically here, followed by the physical requirements and manipulations of our system to generically enforce them in the list that follows, as well as Sec. VI E.

- (1) The inputs $\vec{u}_n$ must be unit white noise.
- (2) $(\mathbf{A}, \mathbf{B}_2)$ is stabilizable and $(\mathbf{C}_2, \mathbf{A})$ is detectable. This assumption is necessary for an internally stable

closed-loop system. More restrictive assumptions are given in [13].

- (3) $\mathbf{D}_{12}$ and $\mathbf{D}_{21}$ both have full column rank.
- (4) $\mathbf{D}_{12}^T \mathbf{D}_{12}$ is invertible and numerically nonsingular.
- (5) $\mathbf{D}_{21} \mathbf{D}_{21}^T$ is invertible and numerically nonsingular.
- (6) $\mathbf{D}_{11} = 0$. Note that $\mathbf{D}_{22}$ is often also assumed to be 0 for use separable solutions, but does not have to be. If $\mathbf{D}_{22} \neq 0$, the optimal controller can be computed as though it is zero and then transformed as shown in Eq. (52).
- (7) Rank $\begin{bmatrix} \mathbf{A} - j\omega \mathbf{I}_n & \mathbf{C}_1 \\ \mathbf{B}_2 & \mathbf{D}_{12} \end{bmatrix} = n + m_2$, $\forall \omega \in \mathbb{R}$, where n is the order of $\mathbf{A}$, m_2 is dimension of $\vec{u}_c$, and $\mathbf{I}$ is the identity matrix also of order n .
- (8) Rank $\begin{bmatrix} \mathbf{A} - j\omega \mathbf{I}_n & \mathbf{C}_2 \\ \mathbf{B}_1 & \mathbf{D}_{21} \end{bmatrix} = n + p_2$, $\forall \omega \in \mathbb{R}$, where p_2 is the dimension of $\vec{y}_{\text{meas}}$.

These formal assumptions restrict the types of systems that can be solved for using this method. The reasoning behind each assumption, as well as a translation into more useful language, is presented as follows:

Assumption 1. This assumption requires that the noise entering the system is Gaussian white noise. This requirement is ensured for systems that require non-white noise by using noise-shaping filters, E and M , on the two noise inputs, which unwhiten the noise to admit a physical source. Note that at least one of the two noise shapes must have a nonsingular D matrix, indicating that the noise becomes asymptotically white at sufficiently high frequency. This is typically the measurement noise M . This is required formally to enforce assumption 3, but is physically necessary to make a well-posed optimization trade-off between low-frequency system noise and measurement noise.

Assumption 2. This requirement is required for the control and estimation subproblems of the separation principle to be individually numerically well posed. Physically, an observable but unstabilizable state contributes to rms but cannot be controlled to suppress its contribution. Thus, it is either a fundamental problem from its real noise contribution or should simply be ignored. If ignored, model reduction should be used. Because of well-known issues in choosing tolerances, structural algorithms to produce minimal realizations [30] are not preferred. Instead, Gramian-based truncation algorithms both reduce systems and improve other numerical properties. We employ singular perturbation approximation-based model reduction [31] using AB09ND from SLICOT [32].

Assumptions 3–5. These requirements are necessary to make a well-posed minimization problem where the control system couples white noise into its states through $\mathbf{D}_{21}$ and then that noise is coupled to the output through $\mathbf{D}_{12}$. Without those requirements, there is excessive freedom to choose how the noise is rolled off. Along with requirements on the E and M shaping

filters, this also requires that the “plant” component has a nonsingular $\mathbf{D}$ matrix and thus cannot be strictly proper. Additionally, at least one of the FOM filters must have a nonsingular $\mathbf{D}$ matrix. Accommodating more general strictly proper plant systems is non-trivial, but is generically possible with the augmentation methods presented in the following section.

Assumption 6. The $\mathbf{D}_{11}$ condition is necessary to prevent an infinite noise power from contaminating the problem. The $\mathbf{D}_{22}$ conditions are a mathematical convenience in the problem statement but can be compensated for by advanced problem restructuring [33].

Assumptions 7–8. If either of these rank conditions is insufficiently met, then the open-loop gain may have a zero on the imaginary line, and the existence of solutions is not guaranteed, or at least substantially more nuanced. It can often be solved anyway with the numerical reliability of generalized eigenvalue and descriptor-system techniques. See remark 14.1 of [15].

E. System augmentation

In order to accommodate the assumptions of Sec. VID, some augmentations and manipulations are made to make the solver more applicable to general interferometer systems. One example of such augmentations, which has already been discussed, is the inclusion of noise shaping and the FOM in Sec. III B. Further manipulations change the system laid out in Fig. 1 into a new modified system of Fig. 12.

To make the new system equivalent to the original, its transfer matrix $\mathbf{T}(f)$ from the $\vec{u}_n$ inputs to the $\vec{y}_{\text{FOM}}$ outputs must be preserved. The control inputs and outputs can also be modified during this process, but only in ways that can be stably inverted. For a SISO plant and control systems, there is a general method for meeting the assumptions necessary for the solver, as long as the noise-shaping filters and figures of merit meet the physical conditions outlined at the end of Sec. VID.

The following manipulations ensure that assumption 4 is met. To do so, much of the plant must be included in the environmental noise block. Most physical plants, such as optical systems or suspensions, have some roll-off at high frequency. This makes them strictly proper systems and implies that the plant subsystem’s state-space representation has $\mathbf{D} = 0$. By moving most of the plant to the environmental noise block, it can be assured that the remainder has a balanced number of poles and zeros, so that $\mathbf{D}_{12} \neq 0$. If the plant is stable, then the only part of the plant that is left in the plant block is the delay or other non-minimum-phase zeros, which can not be moved without changing the physical interpretation. Unstable poles must also remain in the plant section to preserve the problem formulation. The requirement is that the reformulated plant system P' has $\lim_{f \rightarrow \infty} P'(f) > 0$ while containing the same non-minimum-phase zeros and unstable poles as P .

Under this redefinition, P' becomes the new plant subsystem, K' is the new calculated controller, and $G' = K'P'$ is the new open-loop gain. Because the $\mathbf{T}_2$ and $\mathbf{T}_\infty$ transfer functions are preserved, the open-loop gain must be the same as the original formulation, thus $G = G'$. After computing K' , the controller can be compensated using $K \equiv K'P'/P$ to find the controller of the original formulation. Note that this representation may not be a proper transfer function and thus may not admit a state-space representation. A simple reason the aforementioned solvers fail is that their outputs are in state-space form, which may not be able to represent optimal controllers, even if the process can properly pose the optimization problem. For SISO systems, the final controller may be computed from K' using a zero-pole-gain (ZPK) form or using a descriptor state-space representation.

After transformation, Eq. (11) becomes

$$S_p = \left| \frac{1}{1 - G'} \right|^2 S'_{\text{env}} + \left| \frac{G'}{1 - G'} \right|^2 S_{\text{meas}}, \quad (53)$$




FIG. 12. The layout of the augmented and modified plant system with a controller attached. This alteration of Fig. 1 shows the modifications that preserve $\mathbf{T}$ while better ensuring the reliability conditions of Sec. VID.

where

$$S'_{\text{env}} = |P/P'|^2 S_{\text{env}}, \quad (54)$$

which is equivalent to the original formulation.

A useful convention for simple plant systems, such as with our ASC example, is to make $|P'| = 1$ (an inner factor), so that it captures delay and stability effects, without any additional shaping. For this convenient convention, the computed $K' = G$ when the system has no delay, $P' = 1$. This convention always has the property $|K'| = |G|$.

If the plant is unstable, all the unstable poles must remain in the plant. To ensure that assumption 4 is still met, the stable version of the unstable poles, i.e., the pole swapped to the left-hand plane, must be added to the plant as zeros. Then those same stabilized poles must be added to the environmental noise system as poles. Since there will be the same number of poles and zeros in P , $\mathbf{D}_{22} \neq 0$. When executed correctly, the stable poles added to E' will cancel the zeros added to P' .

These modifications are essential and relatively simple to implement for SISO models. We leave for future work the task of refactoring MIMO systems in an analogous manner. We reiterate that the composition of P' uses only the unstable poles and non-minimum-phase zeros of P . Making P' an “inner coprime” factor of P will satisfy the conditions. We implemented the factorization described above using ZPK representations of SISO filters due to the unavailability of inner/outer factorization functions in Python’s wrappers of SLICOT (e.g., SB08CD), but anticipate that coprime factorization will be effective to extend this formulation for MIMO plants.

In addition to the structural modifications, we also numerically modify the system representation. Given that many individual control systems in LIGO deal with noise over multiple orders of magnitude and models can contain over a hundred states, numerical problems can quickly become a problem for the solvers. To combat this, the system undergoes a matrix balancing stage [34,35]. This better conditions the problems for generalized eigenvalue solvers. Before solving and after attaching the controller, the system’s order was reduced using square-root singular perturbation approximation with preliminary scaling of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ [31,36]. We find that these conditioning steps can be essential to solving for complex controllers of gravitational-wave interferometer noise models.

F. Calculating the $\mathcal{H}_\infty$ norm

The $\mathcal{H}_\infty$ norm is a measure of the peak gain of a closed-loop system over all frequencies. When used as a reinterpretation of the $\mathbf{T}_2$ transfer function, it can be considered as a measure of a system’s robustness against an environment (or adversary) that directs all input noise power to the worst possible frequency. For specific $\mathbf{T}_\infty$ designs, it can be used to compensate specific, known perturbations to $\mathbf{A}$ of the

plant system. For gravitational-wave detectors, the stationary noise does not actively counter our control systems, but has such a large magnitude that it prevents solutions with acceptable stability margins.

Our design of the $\mathbf{T}_\infty$ transfer function, using u_∞ to the $\tilde{\mathbf{y}}_{\text{FOM}}$ outputs, specifically targets the phase margin, just like hand-tuned classical control designs. The $\mathcal{H}_\infty$ norm represents the maximum possible gain that is achievable by the system given all possible disturbances.

For a continuous system with a transfer function $\mathbf{T}_\infty(s)$, the $\mathcal{H}_\infty$ norm is most formally defined as

$$\|\mathcal{H}_\infty\| = \sup_{\omega} \sigma_{\max}[\mathbf{T}_\infty(j\omega)], \quad (55)$$

where $\sup$ is the supremum over ω and $\sigma_{\max}$ is the maximum singular value of the matrix $\mathbf{T}(j\omega)$.

The means to efficiently find this norm is another highly worthwhile result from mathematical control called the bounded real lemma. Unlike directly solving for the $\mathcal{H}_2$ norm, the $\mathcal{H}_\infty$ norm can be detected by first assuming a bound γ , where $\|\mathcal{H}_\infty\| < \gamma$. Thus, finding the optimal $\mathcal{H}_\infty$ controller requires a search for the minimum value of γ .

The notion to test this comes from finding the zeros of $0 = \mathbf{T}_\infty(f)\mathbf{T}_\infty^\dagger(f) - \gamma^2\mathbf{1}$, which is the cross-spectrum matrix of the system (positive semidefinite), less the γ bound. This equation can be produced from a state space connected to its time-reversed transpose. Finding the zeros of this equation is equivalent to finding the poles of the system $\gamma^{-2}\mathbf{T}_\infty(f)\mathbf{T}_\infty^\dagger(f)$ that has its output fed back to its input. The $\mathbf{A}$ matrix of this feedback system is the following Hamiltonian matrix $\mathbf{H}$, expressed here using $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ of $\mathbf{T}_\infty$,

$$\mathbf{H} \equiv \begin{bmatrix} \mathbf{A} & \gamma^{-2}\mathbf{B}\mathbf{B}^T \\ -\mathbf{C}^T\mathbf{C} & -\mathbf{A}^T \end{bmatrix}. \quad (56)$$

It follows that, if there are no eigenvalues of this matrix along the imaginary line, then $\|\mathcal{H}_\infty\| < \gamma$ for the $\mathbf{T}_\infty$ system. Thus, it is possible to choose a value for γ , construct the Hamiltonian matrix, check the spectrum of the solution, then adjust the γ value and repeat until the lowest γ value is found within the desired precision. The above algorithm has been highly refined, and it can be directly computed using the SLICOT routine AB13DD [37].

We present this result to a physics audience to indicate that $\mathcal{H}_\infty$ bounds can be clearly and concisely stated and can be related to a physically meaningful feedback relation. The mixed-sensitivity approach to follow was originally derived by applying Eq. (56) to a noise-minimization problem of Eq. (46) using Lagrangian constraints.

Finally, the use of singular values in the statement of the norm, Eq. (55), may not be intuitive. For our single input of u_∞ , there can be only one singular value of this matrix, which must be the 2-norm of the vector. This leads to the

simple γ -constraint of Eq. (35) without further concern about the complexity of solving for singular values.

G. Optimal $\mathcal{H}_\infty$ controllers

Here, the solution method for $\mathcal{H}_\infty$ -optimal control is presented to relate it to our mixed-sensitivity approach and as background for a physics audience. The primary takeaway is in how $\mathcal{H}_\infty$ utilizes the parameter γ as a bound on its norm, which is also utilized in mixed-sensitivity control.

Similar to the $\mathcal{H}_2$ controller, the $\mathcal{H}_\infty$ controller is also found at the minimum of its respective norm. The $\mathcal{H}_\infty$ norm is a measure of the worst-case scenario for the control system as the maximum gain from an input disturbance to the output over all frequencies. In order to calculate an optimal controller, $\mathbf{K}_{\infty(\text{sub})}$, for the system given in Eqs. (36)–(38), it must follow the assumptions of Sec. VIC [38].

Unlike solving for the optimal $\mathcal{H}_2$ controller, the $\mathcal{H}_\infty$ controller, like its norm, has one free parameter, γ , where $\|\mathcal{H}_\infty\| < \gamma$. Thus, finding the optimal $\mathcal{H}_\infty$ controller requires a search for the minimum value of γ .

A suboptimal $\mathcal{H}_\infty$ controller is found by solving two continuous algebraic Riccati equations,

$$\begin{aligned} \mathbf{A}^T \bar{\mathbf{X}}_\infty + \bar{\mathbf{X}}_\infty \mathbf{A} - \bar{\mathbf{X}}_\infty \\ (\gamma^{-2} \mathbf{B}_1 \mathbf{B}_1^T - \mathbf{B}_2 \mathbf{B}_2^T) \bar{\mathbf{X}}_\infty + \mathbf{C}_1^T \mathbf{C}_1 = 0 \end{aligned} \quad (57)$$

and

$$\begin{aligned} \mathbf{A} \bar{\mathbf{Y}}_\infty + \bar{\mathbf{Y}}_\infty \mathbf{A}^T - \bar{\mathbf{Y}}_\infty \\ (\gamma^{-2} \mathbf{C}_1^T \mathbf{C}_1 - \mathbf{C}_2^T \mathbf{C}_2) \bar{\mathbf{Y}}_\infty + \mathbf{B}_1 \mathbf{B}_1^T = 0, \end{aligned} \quad (58)$$

where $\bar{\mathbf{X}}_\infty \geq 0$ and $\bar{\mathbf{Y}}_\infty \geq 0$ must be stabilizing solutions. They are guaranteed to exist as $\gamma \rightarrow \infty$, but will fail to exist at some minimum $\gamma_{\min}$. The resulting admissible controller is

$$\mathbf{K}_{\infty(\text{sub})} = \left[\begin{array}{c|c} \mathbf{A}_{K_\infty} & \mathbf{Z}_\infty \bar{\mathbf{Y}}_\infty \mathbf{C}_2^T \\ \hline -\mathbf{B}_2^T \bar{\mathbf{X}}_\infty & 0 \end{array} \right], \quad (59)$$

where

$$\begin{aligned} \mathbf{A}_{K_\infty} = \mathbf{A} + (\gamma^{-2} \mathbf{B}_1 \mathbf{B}_1^T - \mathbf{B}_2 \mathbf{B}_2^T) \bar{\mathbf{X}}_\infty \\ - \mathbf{Z}_\infty \bar{\mathbf{Y}}_\infty \mathbf{C}_2^T \mathbf{C}_2 \end{aligned} \quad (60)$$

and $\mathbf{Z}_\infty = (\mathbf{I}_n - \gamma^{-2} \bar{\mathbf{Y}}_\infty^T \bar{\mathbf{X}}_\infty)^{-1}$. The optimal $\mathcal{H}_\infty$ controller $\mathbf{K}_\infty$ is the unique controller where γ is at its minimum possible value $\gamma_{\min}$ for these equations to produce a closed-loop stable system. The failure to produce a solution below this minimum can happen in several ways [40].

One can inspect the form of Eq. (58) and notice that, as $\gamma \rightarrow \infty$, it recovers the equations and solution of the LQG problem. More generally, for the set of suboptimal controllers $\gamma_{\text{opt}} < \gamma < \infty$, this method of $\mathcal{H}_\infty$ solution calculates

the “central” controller that obeys the γ bound while minimizing a $\mathbf{T}_2 = \mathbf{T}_\infty$ problem. Thus, for mixed-sensitivity problems for which $\mathbf{T}_2$ should be the same as $\mathbf{T}_\infty$, a simple computation of a suboptimal controller also solves the mixed-sensitivity problem [13].

For the phase margin constraint and construction of $\mathbf{T}_\infty$ presented in this work, the suboptimal central controllers of the associated $\mathcal{H}_\infty$ -optimal control problem lack all information about the environmental and sensing noise weights. They are thus not useful for noise minimization. The following shows how to find noise-optimal solutions within the set of suboptimal controllers.

VII. OPTIMAL LQG CONTROLLERS WITH A BOUNDED $\mathcal{H}_\infty$ NORM

The previous section introduced control concepts for integration into a mixed-sensitivity solver and defined the relevant state-space representation.

Mixed-sensitivity control offers the best of both LQG and $\mathcal{H}_\infty$ optimization, minimizing noise while ensuring robust performance. Despite its considerable benefits, it is less commonly known or taught, as it is advanced and algebraically cumbersome to describe and solve.

The methodology of this paper shows that it can be incorporated into existing designs in a straightforward manner through the addition of an u_∞ input, such that end users will only need to input the various components of Fig. 1, and software can do the rest, shielding the user from the problem description. This then requires numerical algorithms that can reliably solve such mixed-sensitivity problems despite the aforementioned numerical challenges, without requiring careful user tuning. This is what we largely, but not entirely, achieved through the following methods.

For context, mixed-sensitivity control was studied shortly after the development of a reliable $\mathcal{H}_\infty$ -optimal solution using the separation principle. Interestingly, the more direct method of Lagrangian constraints was used to apply the $\mathcal{H}_\infty$ bound to the LQG noise-minimization problem. This was explored in the two initial papers leading into this subject [17,18]. These two works establish the necessary conditions as a set of coupled Riccati equations to solve and maintain a weaker form of the separation principle. Shortly thereafter, sufficient conditions were determined [41] and the two algebraic approaches were related as duals of each other [42]. A notable aspect of these alternate approaches is that they set up the problem using a similar arrangement of inputs and outputs. Namely, there are three sets of inputs and two sets of outputs, or vice versa. One of each is used for observing and control and the remainder to define the optimization and bound norms. For the approach of this work to ensure phase margin, the arrangement of inputs and outputs is why the $\mathbf{T}_2$ and $\mathbf{T}_\infty$ matrices are related and lead to the constraint of Eq. (35).

Three inputs and outputs are needed to define a fully general problem that decouples the $\mathbf{T}_2$ and $\mathbf{T}_\infty$ transfer

matrices, but this removes all remaining algebraic separability. Since the original algebraic formulations, ongoing work has reframed the mixed-sensitivity problem using linear matrix inequalities (LMIs) on a convex domain [43–45]. Our first attempts to optimize our system used this approach, through the Matlab function `h2hinfsvn`. This function uses an LMI approach. For the systems of this paper, it was only successful for $\zeta = 0$ problems, with no BNS weighting. It was also unacceptably slow. As to why, we can only speculate that the LMI formulation or the specific numerical implementation and internal numerical tolerances are not currently structured or balanced to handle the large dynamic range of gravitational-wave interferometer models. Alternatively, they may be particularly unsuited when the $\mathbf{T}_2$ -weighted optimization problem is defined so differently than the central controllers of our rather dissimilar $\mathbf{T}_\infty$ problem. To produce this work, we utilize the earlier algebraic approach of Bernstein-Haddad (BH) [17,18].

We use the convention of u_{inf} being on the input side, as that is required for our system layout. This follows the conventions of [13] and is dual to [17,18]. We otherwise use BH equations, as they present algebraic solutions for the most general problem. Our formulation uses the following definitions from BH, in our notation, to express the algebraic Riccati equations that follow:

$$\mathbf{R}_1 = \mathbf{B}_{1n} \mathbf{B}_{1n}^T, \quad (61)$$

$$\alpha^2 \hat{\mathbf{R}}_2 = \mathbf{R}_2 = \mathbf{D}_{21n} \mathbf{D}_{21n}^T, \quad (62)$$

$$\mathbf{R}_{12} = \mathbf{B}_{1n} \mathbf{D}_{21n}^T, \quad (63)$$

$$\mathbf{V}_1 = \mathbf{C}_1 \mathbf{C}_1^T, \quad (64)$$

$$\mathbf{V}_2 = \mathbf{D}_{12} \mathbf{D}_{12}^T, \quad (65)$$

and

$$\mathbf{V}_{12} = \mathbf{C}_1^T \mathbf{D}_{12}. \quad (66)$$

These factors are related to the noise inputs and FOM outputs. More relations to map our formulation back to the BH papers are given in Appendix C. The expressions below require structural similarities between the $\mathbf{D}$ matrices of the rms and $\mathcal{H}_\infty$ problems. Namely, they have to be proportional so that the β factor can be defined as the ratio, $\beta^2 = \mathbf{D}_{21\infty} \mathbf{D}_{21\infty}^T / \mathbf{R}_2$. This is easy to ensure for SISO models, but may be more difficult in MIMO extensions.

The factor α is nominally 1. The equations below are all expressed in terms of $\hat{\mathbf{R}}_2$ rather than $\mathbf{R}_2$ or $\mathbf{R}_{2\infty} = \mathbf{D}_{21\infty} \mathbf{D}_{21\infty}^T$, leading to the included factors of α and β . By choosing $\alpha > 1$, we are effectively rescaling the noise matrices in the problem, which is useful in finding solutions.

The set of equations we then must solve is the following three Riccati equations in the positive semidefinite matrices $\mathbf{Q}$, $\mathbf{Z}$, $\hat{\mathbf{Q}}$. Included below are the equations along with some convenience definitions to shorten the Riccati equations and highlight their quadratic form. The first of the equations, solving for $\mathbf{Q}$, is independent of the others, indicating the remaining separability of the problem. The equations in $\mathbf{Z}$, $\hat{\mathbf{Q}}$ are intercoupled and both rely on the solution to $\mathbf{Q}$. Their simultaneous solution requires an iterative method. The equations are

$$0 = \mathbf{A}^T \mathbf{Q} + \mathbf{Q} \mathbf{A} + \mathbf{V}_{1\infty} - (\mathbf{Q} \mathbf{B}_2 + \mathbf{V}_{12\infty}) \mathbf{V}_{2\infty}^{-1} (\mathbf{B}_2^T \mathbf{Q}^T + \mathbf{V}_{12\infty}^T), \quad (67)$$

$$\mathbf{Q}_a \equiv \mathbf{Q} \mathbf{B}_2 + \mathbf{V}_{12\infty}, \quad (68)$$

$$\bar{\mathbf{A}}_Z \equiv \mathbf{A}^T - \frac{\beta^2}{\alpha^2 \gamma^2} \hat{\mathbf{Q}} \mathbf{R}_1, \quad (69)$$

$$0 = \bar{\mathbf{A}}_Z^T \mathbf{Z}' + \mathbf{Z}' \bar{\mathbf{A}}_Z + \mathbf{R}_1 - \mathbf{Z}' \left(\frac{\mathbf{C}_2^T \hat{\mathbf{R}}_2^{-1} \mathbf{C}_2}{\alpha^2} - \frac{\beta^4 \hat{\mathbf{Q}} \mathbf{R}_1 \hat{\mathbf{Q}}}{\alpha^4 \gamma^4} - \frac{\beta^2 \mathbf{Q}_a \mathbf{V}_{2\infty}^{-1} \mathbf{Q}_a^T}{\alpha^2 \gamma^2} \right) \mathbf{Z}', \quad (70)$$

$$\mathbf{Z} \equiv \mathbf{Z}' / \alpha^2, \quad (71)$$

$$\bar{\mathbf{A}}_{\hat{\mathbf{Q}}} \equiv \mathbf{A}^T - \mathbf{C}_2^T \hat{\mathbf{R}}_2^{-1} \mathbf{C}_2 \mathbf{Z}, \quad (72)$$

and

$$0 = \bar{\mathbf{A}}_{\hat{\mathbf{Q}}} \hat{\mathbf{Q}} + \hat{\mathbf{Q}} \bar{\mathbf{A}}_{\hat{\mathbf{Q}}}^T + \mathbf{Q}_a \mathbf{V}_{2\infty}^{-1} \mathbf{Q}_a^T + \hat{\mathbf{Q}} \frac{\beta^2}{\gamma^2} \mathbf{Z} \mathbf{C}_2^T \hat{\mathbf{R}}_2^{-1} \mathbf{C}_2 \mathbf{Z} \hat{\mathbf{Q}}. \quad (73)$$

Using the solutions ($\mathbf{Q}$, $\mathbf{Z}$, and $\hat{\mathbf{Q}}$) to the above equations, the optimal controller is then computed to be

$$\mathbf{K}_{\mathcal{H}_{2/\infty}} = \left[\begin{array}{c|c} \mathbf{A}_K & \mathbf{B}_K \\ \hline \mathbf{C}_K & 0 \end{array} \right], \quad (74)$$

where

$$\mathbf{B}_K = \mathbf{Q}_a \mathbf{V}_{2\infty}^{-1}, \quad (75)$$

$$\mathbf{C}_K = -\hat{\mathbf{R}}_2^{-1} \mathbf{C}_2 \mathbf{Z}, \quad (76)$$

and

$$\mathbf{A}_K = \mathbf{A}^T + \mathbf{C}_2^T \mathbf{C}_K - \mathbf{B}_K \mathbf{B}_2^T - \mathbf{B}_K \mathbf{D}_{22}^T \mathbf{B}_K. \quad (77)$$

The last term on $\mathbf{A}_K$ accounts for the $\mathbf{D}_{22} \neq 0$, so no further adjustment is required.

The challenge in solving these equations is in finding the simultaneous solutions. We note some discussion from [17]. As long as $\hat{\mathbf{Q}}$ is positive semidefinite, then the solution for $\mathbf{Z}$ will be noise optimal and stable, but may not obey the given $\mathcal{H}_\infty$ constraint. Likewise, for any positive semidefinite $\mathbf{Z}$, a solution for $\hat{\mathbf{Q}}$ indicates that the control is stable and the $\mathcal{H}_\infty$ bound is met, although it may not be noise optimal. These equations are solved iteratively using a holonomy method, where parameters are scanned from a parameter space where solutions are reliably found to one where they are more challenging to find. The initial condition of the iteration is for $\hat{\mathbf{Q}} = 0$.

Instead of performing holonomy on solutions where γ is transitioned from ∞ to its final value, we instead transition α from ∞ to the value 1 while alternately solving the coupled equations. This method is physically motivated. It is equivalent to scaling the sensing noise from being large and dominant, where loops do not need aggressive shaping, to the final form where sensing noise is many orders of magnitude below the environmental noise, while maintaining γ at its nominal desired value.

A. Controllers stability and implementation

While the closed-loop system is guaranteed to be stable with either the $\mathcal{H}_2$ or the $\mathcal{H}_\infty$ bounded controller, the controllers alone are not guaranteed to be stable. We have seen that the outputs of the mixed-sensitivity solver produce unstable $\mathbf{K}_{\mathcal{H}_{2/\infty}}$.

In theory, there is nothing concerning about a controller alone being unstable since the overall stability of the system is determined in the closed-loop case. However, it is possible that some may not want to use controllers that are themselves unstable. Especially when gaining lock in one of the interferometer's degree of freedom, a stable controller is beneficial. A controller that is independently stable and whose closed loop is stable is called a "strongly stable" controller. The equations for strongly stable controllers have been computed in [46,47]. However, it is beyond the goals of this paper to implement.

We have found that, in situations where the solved controller is unstable, simply flipping its unstable poles to the stable side still generally results in a closed-loop stable system. This is performed during a reduction stage to lower the order of the controller for implementation in digital control. We have verified that the final system model using an order-reduced controller is still stable and that the phase margin of the order-reduced controller is still acceptable after the reduction step.

VIII. SUMMARY AND FUTURE APPLICATIONS

In this work, we developed three new concepts to incorporate modern optimal methods into gravitational-wave interferometer control.

First, we established a specific figure of merit for optimizing control noise using existing SNR or detection range computations, interpreted as weighted $\mathcal{H}_2$ norms on the control performance, indicating that LQG control is an appropriate framework for optimizing interferometer noise performance.

Second, we compute a Pareto front between a total plant motion and the SNR figures of merit by solving for an optimal controller while adjusting a relative weight ζ between two noise-metric outputs of our system model.

Third, we use mixed sensitivity control to utilize LGQ for GW control without the issue of arbitrarily-low phase margin; in our case, by imposing a specific $\mathcal{H}_\infty$ constraint on the maximum magnitude closed-loop sensitivity function $G(f)/[1 - G(f)]$.

Implementing these methods together poses several challenges due to the dynamic range of weighting filters involved, but we were able to reliably form solutions using algebraic methods to solve Riccati equations and an iterative method to jointly solve the coupled Riccati equations that arise in mixed-sensitivity control.

A. Applications

There are several applications beyond our alignment control example where our method could optimize gravitational-wave detectors.

One of the most apparent applications for this work is to optimize or analyze all classical SISO-designed controllers in LIGO. This becomes an obtainable goal since the majority of the systems that contribute significant control noise are already modeled with their noise contributions budgeted [3]. The primary barrier to this is unmodeled cross-couplings between different control loops. While achieving SNR-optimized control for the whole interferometer is an ambitious goal, it is possible that using SISO optimizations on each individual control loop can approach the MIMO-optimized control noise. Future work to adapt this methodology for noise-optimal MIMO designs could reveal whether a gap exists between the purely SISO design approach and the intrinsically MIMO design in LIGO's systems.

One of the most exciting applications of this work is to automatically update controllers in changing noise environments. The noise sources modeled to produce these controllers are not static. They drift over time. Having a robust controller can help the control system adapt to drifting conditions, but the controller is only optimal against a specific noise model. In this case, it is advantageous to have a control solver that continuously updates its models and solutions to respond to both changing noise and a changing plant. This work constructs automated and optimal designs that eliminate the laborious hand tuning. This opens the possibility of further automating controller design. Commissioning work then moves from designing control loops to maintaining physical and data-driven

models of the interferometer, which are useful for more than just control design.

Currently, it is not known how much control noise can be eliminated from the interferometer as a whole. As shown for our example in Sec. V, there are controllers whose noise can be improved. If, however, it is found that the existing controllers are approaching the minimum possible control noise in the interferometer, it will become imperative to use additional means to improve the overall sensitivity. If the goal remains to reduce control noise past what our optimal approach can deliver, then nontraditional control (artificial intelligence or nonlinear controllers) or hardware upgrades must then be considered. This work enables one to find the limits of control-loop design for SISO systems and better understand how hardware and control loops interact and trade off. With the automated and optimal design we present, more effective hardware and instrumentation requirements can be established more efficiently. Thus, this work could prove useful in parametric designs and upgrades to gravitational-wave detectors, where noise models require auxiliary loop designs, as we now know how to fully saturate performance and treat control-coupled noises as a fundamental noise source.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the U.S. National Science Foundation, as well as Elenna Capote, Matthew Evans, Gabriele Vajente, Jeffrey Wack, Terrence Tsang, and the McCuller Group at Caltech for their helpful comments and discussions. LIGO operates under Cooperative Agreement No. PHY-2309200. Advanced LIGO was built under Grant No. PHY-0823459.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: THE BILINEAR NOISE METRIC AND SPECTRAL CONVOLUTION

This appendix shows how a multiplicative, bilinear coupling affects the spectrum of a signal. Its result is that the spectrum of a convolution is the convolution of a spectrum. Start with two independent noise channels N_A and N_B . Each has a spectral density defined as

$$\langle N_A(f)N_A(-f-f') \rangle \equiv \frac{1}{2}S_A(f)\delta(f') \quad (\text{A1})$$

and

$$\langle N_B(f)N_B(-f-f') \rangle \equiv \frac{1}{2}S_B(f)\delta(f'), \quad (\text{A2})$$

where their independence gives the property $\langle N_A(f) \times N_B(f') \rangle = 0$.

Now assume that you have a channel $N_x(t) \equiv N_A(t)N_B(t)$. By the convolution theorem, then

$$N_x(f) = N_A(f) \otimes N_B(f) \quad (\text{A3})$$

$$\equiv \int_{-\infty}^{\infty} N_A(f')N_B(f-f')df'. \quad (\text{A4})$$

The spectrum of this bilinear product can now be derived through

$$\frac{1}{2}S_x(f)\delta(f') \equiv \langle N_x(f)N_x(-f-f') \rangle \quad (\text{A5})$$

$$= \left\langle \int_{-\infty}^{\infty} N_A(f'')N_B(f-f'')df'' \int_{-\infty}^{\infty} N_A(f''')N_B(-f-f'-f''')df''' \right\rangle \quad (\text{A6})$$

$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \underbrace{\langle N_A(f'')N_A(f''') \rangle}_{\frac{1}{2}S_A(f'')\delta(f''+f''')} \underbrace{\langle N_B(f-f'')N_B(-f-f'-f''') \rangle}_{\frac{1}{2}S_B(f-f'')\delta(f'+f''+f''')} df''df''' \quad (\text{A7})$$

$$= \frac{1}{4}\delta(f') \int_{-\infty}^{\infty} S_A(f'')S_B(f-f'')df''. \quad (\text{A8})$$

Thus, for this type of bilinear noise,

$$S_x(f) = \frac{1}{2}S_A(f) \otimes S_B(f). \quad (\text{A9})$$

This short derivation should now be applied to the bilinear noise problem of Sec. III B. Now we apply it to the bilinear noise term, $N_X = \chi N_p(t) N_a(t)$ seen in Eq. (22).

We first assume that the rms noise power of S_p is concentrated at low frequencies, i.e.,

$$S_p(f) \approx \tilde{R}_p \delta(f + \epsilon) \quad (\text{A10})$$

for some small ϵ .

Applying the convolution from this derivation then gives

$$S_X(f) = \chi \frac{\tilde{R}_p}{2} (S_a(f + \epsilon) + S_a(f - \epsilon)) = \chi \tilde{R}_p S_a(f). \quad (\text{A11})$$

This final result indicates that $|C| = \chi \tilde{R}_p$ as used in Sec. III B.

APPENDIX B: ALTERNATE NOISE ARRANGEMENT

Along with the model presented in the main text, there are two major alternate approaches to where to place the FOM to model bilinear-type noise couplings, and we also suggest a third approach. Each has its own reasons for implementation.

Another alteration present in this system is the apparent bypass of the plant by the environmental noise. In this configuration, the environmental noise-shaping filter now also includes the shape of the plant. This means the new shape of the alternate E' filter can be written as $|EP|$.

Letting this input noise-shaping filter represent the environmental noise passed through the plant allows it to be connected after the plant, since it already has picked up the shape from the plant. This alteration allows the very useful option of fitting the filter to an open-loop spectrum at y_{meas} as is shown in Fig. 2.

1. Actuator-centered form

For the ASC bilinear noise model often assumed by LIGO control commissioners, the control noise arises solely from the u_c feedback contribution to the noise. This type of bilinear noise injection can be modeled by the following coupling:

$$h_{\text{signal}}(t) = h_{\text{gw}}(t) + N_{\text{det}}(t) + \chi N_p(t) u_c(t). \quad (\text{B1})$$

This coupling results when the coupling function is computed from

$$|C|^2 = \frac{S_{h,\text{exc}} - S_h}{S_{K,\text{exc}} - S_K}, \quad (\text{B2})$$

where the excited channel in the denominator is the output of the control filter (for alignment, it is related to units of torque on the mirror).

Figure 13 shows how to set up the augmented plant system to implement this alternate form. Note that the injection of the environmental noise is now after the plant, which puts it into the same units as the measurement noise. This makes the injection into DARM appear as




FIG. 13. Two alternate system layouts to Fig. 1, the noise-augmented plant system with a feedback controller attached. In this configuration, the environmental noise block is the equivalent of $|EP|$ rather than just E as in all previous system layouts. This change does not alter the system but allows the user to fit an open-loop spectrum at point y_{meas} as is shown in Fig. 2. Here, the figures of merit are arranged showing two alternate bilinear noise optimizations. The first, the actuator-centered (AC) form often used by LIGO, is where the BNS FOM is connected to the control signal u_c and treated as a noise N_K with the purple connection. The second is the total-noise (TN) form that can be used for noncoupled degree of freedom and connects the BNS FOM to the plant's state noise N_p with a green connection. Only one of these two alternate connection forms can be used at once.

$$S_K = \left| \frac{G(f)}{1 - G(f)} \right|^2 (S_{\text{env}} + S_{\text{meas}}), \quad (\text{B3})$$

for this arrangement of noises in the model.

This arrangement is primarily a conceptual change from the model used in the main text. In it, the coupling function takes on different units since it is modeled before entering the plant, and so it is in units of DARM length/torque for our alignment control example.

2. Total-noise optimization form

The form and connection point of the FOM that is perhaps the most appropriate for quadratically nonlinear (rather than bilinear) systems requires using the total noise S_p as discussed in Sec. III B. This setup is detailed in Fig. 13 with the connection from N_p to the BNS FOM. The connection of both figures of merit to N_p ensures that both the actuation-point control noise as well as environmental noise are minimized together. This is a more typical arrangement of inputs/outputs and weight filters for LQG problems and does not require two copies of the plant. It corresponds to quadratically nonlinear noise couplings of the form

$$h_{\text{signal}}(t) = h_{\text{gw}}(t) + N_{\text{det}}(t) + \chi N_p^2(t). \quad (\text{B4})$$

In terms of numerical implementation, there are several reasons that the LIGO preferred forms in Appendix B 1 and the main text are less preferable to the total-noise form. One of the major reasons is that, in LIGO's form, most of the low-frequency environmental noise shows up in S_c when $|G|$ is large. The dynamic range of the BNS FOM covers nearly all of the range of a double float. Thus, in computing the noise contribution in this model, there are numerical difficulties multiplying the environmental noise filter with the BNS-weighting filter inside a state space. The second reason is that bilinear noise is well motivated to be driven by N_p , or at least some mixture of N_p and u_c . In some instances, LIGO has found that further accelerating the roll-off of G does not reduce the noise in DARM. The primary noise coupling chosen for this work provides a reason for that result: the bilinear noise can also be driven by residual environmental noise in the plant that is reflected in N_p rather than in N_a or $u_c = N_K$.

3. Measurement-only optimization form

A final alternate approach would be to use a BNS FOM channel with the noise coupling $S_c' = |G/(1 - G)|^2 S_{\text{meas}}$ only, without the environmental noise contribution. This then requires the two figures of merit to be sourced with different balances of the shaped noise filters (one with environmental noise and one entirely without). The authors do not know of a means to implement this through the block diagram without two copies of the unknown

controller K ; however, they believe it should be possible to develop a new set of Riccati equations to solve. This might be done by determining the equivalent noise model using a sum of Lyapunov equations and then solving for the constrained noise-minimization problems using the techniques of [12,17,18,48]. Doing so is likely to further reduce the separability structure expected of LQG and may result in three coupled Riccati equations, rather than the two coupled equations of mixed-sensitivity control. This alternate approach could be primarily useful for numerical implementation reasons and perhaps would be useful to investigate using a convex optimization approach that can more easily represent and solve constrained optimization problems.

APPENDIX C: TRANSLATION OF THE BH EQUATIONS

The notation table of [17], as well as the notation table and Fig. 1 mapping in [18] can be used to translate the notation of this work into the algebraic equations, through the following mappings:

$$A \rightarrow \mathbf{A}_T, \quad B \rightarrow \mathbf{C}_2^T, \quad C \rightarrow \mathbf{B}_2^T, \quad (\text{C1})$$

$$D \rightarrow \mathbf{D}_{22}^T, \quad E_1 \rightarrow \mathbf{B}_{1n}^T, \quad E_{1\infty} \rightarrow \mathbf{B}_{1\infty}^T, \quad (\text{C2})$$

$$E_\infty \rightarrow \mathbf{D}_{11\infty}^T, \quad E_2 \rightarrow \mathbf{D}_{21n}^T, \quad E_{2\infty} \rightarrow \mathbf{D}_{21\infty}^T, \quad (\text{C3})$$

$$D_1 \rightarrow \mathbf{C}_1^T, \quad \text{and} \quad D_2 \rightarrow \mathbf{D}_{12}^T. \quad (\text{C4})$$

Because E_∞ is structurally zero in the problem setup, the following BH parameters are simplified as

$$\mathbf{M} = \mathbf{1}, \quad \mathbf{N} = \mathbf{1}, \quad (\text{C5})$$

$$\mathbf{V}_{1\infty} = \mathbf{V}_1, \quad \mathbf{V}_{2\infty} = \mathbf{V}_2, \quad \mathbf{V}_{12\infty} = \mathbf{V}_{12}, \quad (\text{C6})$$

$$\mathbf{R}_{01\infty} = \mathbf{0}, \quad \mathbf{R}_{02\infty} = \mathbf{0}. \quad (\text{C7})$$

This simplification leads to the reduced expressions

$$\mathbf{R}_{1\infty} = \mathbf{B}_{1\infty} \mathbf{B}_{1\infty}^T, \quad (\text{C8})$$

$$\mathbf{R}_{2\infty} = \mathbf{D}_{21\infty} \mathbf{D}_{21\infty}^T, \quad (\text{C9})$$

$$\mathbf{R}_{12\infty} = \mathbf{B}_{1\infty} \mathbf{D}_{21\infty}^T, \quad (\text{C10})$$

and

$$\hat{\mathbf{R}}_2 \equiv \mathbf{R}_2 / \alpha^2 = \mathbf{R}_{2\infty} / \beta^2, \quad (\text{C11})$$

where

$$\alpha \equiv 1. \quad (\text{C12})$$

$$\mathbf{R}_{1\infty} = \mathbf{0} \quad \text{and} \quad \mathbf{R}_{12\infty} = \mathbf{0}. \quad (\text{C13})$$

We also assume that there is no additional shaping filter on the u_∞ input, causing $E_{1\infty}$ term from BH to be zero. This leads to the additional simplifications

Note that all of these are then applied to the full-form equations of [18] and then simplified into the $\mathbf{Z}$ -form description following Sec. V of [17].

-
- [1] B.P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari *et al.*, Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [2] A. Buikema, C. Cahillane, G. Mansell, C. Blair, R. Abbott, C. Adams, R. Adhikari, A. Ananyeva, S. Appert, K. Arai *et al.*, Sensitivity and performance of the advanced LIGO detectors in the third observing run, *Phys. Rev. D* **102**, 062003 (2020).
 - [3] E. Capote *et al.*, Advanced LIGO detector performance in the fourth observing run, *Phys. Rev. D* **111**, 062002 (2025).
 - [4] R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. Adhikari, V. Adya, C. Affeldt *et al.*, GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
 - [5] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
 - [6] R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. Adhikari, V. Adya, C. Affeldt *et al.*, Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog, *Phys. Rev. D* **103**, 122002 (2021).
 - [7] R. Abbott *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration), Tests of general relativity with GWTC-3, *Phys. Rev. D* **112**, 084080 (2025).
 - [8] E. D. Hall, K. Kuns, J. R. Smith, Y. Bai, C. Wipf, S. Biscans, R. X. Adhikari, K. Arai, S. Ballmer, L. Barsotti *et al.*, Gravitational-wave physics with cosmic explorer: Limits to low-frequency sensitivity, *Phys. Rev. D* **103**, 122004 (2021).
 - [9] L. S. Finn and D. F. Chernoff, Observing binary inspiral in gravitational radiation: One interferometer, *Phys. Rev. D* **47**, 2198 (1993).
 - [10] L. McCuller, Wield irrational, <https://github.com/wieldphysics/wield-irrational> (2023).
 - [11] Y. Nakatsukasa, O. Sète, and L. N. Trefethen, The AAA algorithm for rational approximation, *SIAM J. Sci. Comput.* **40**, A1494 (2018).
 - [12] M. Athans and E. Tse, A direct derivation of the optimal linear filter using the maximum principle, *IEEE Trans. Autom. Control* **12**, 690 (1967).
 - [13] J. Doyle, K. Glover, P. Khargonekar, and B. Francis, State-space solutions to standard H_2 and H_∞ control problems, *IEEE Trans. Autom. Control* **34**, 831 (1989).
 - [14] J. Doyle, Guaranteed margins for LQG regulators, *IEEE Trans. Autom. Control* **23**, 756 (1978).
 - [15] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control* (Prentice Hall, Englewood Cliffs, NJ, 1996).
 - [16] K. Glover and J. C. Doyle, State-space formulae for all stabilizing controllers that satisfy an H_∞ -norm bound and relations to relations to risk sensitivity, *Systems and Control Letters* **11**, 167 (1988).
 - [17] D. Bernstein and W. Haddad, LQG control with an H_∞ performance bound: A Riccati equation approach, *IEEE Trans. Autom. Control* **34**, 293 (1989).
 - [18] W. Haddad and D. Bernstein, Generalized Riccati equations for the full- and reduced-order mixed-norm H_2/H_∞ standard problem, in *Proceedings of the 28th IEEE Conference on Decision and Control*, Vol. 1 (1989), pp. 397–402.
 - [19] V. Ionescu and M. Weiss, The l^2 -control problem for time-varying discrete systems, *Systems and Control Letters* **18**, 371 (1992).
 - [20] This is well known in the control systems community, but less well known in the GW interferometer control commissioning community.
 - [21] A. Varga, On computing 2-norms of transfer function matrices, in *1992 American Control Conference* (1992), pp. 1054–1056, [10.23919/ACC.1992.4792246](https://doi.org/10.23919/ACC.1992.4792246).
 - [22] A. C. Antoulas, 6. Sylvester and Lyapunov equations, in *Approximation of Large-Scale Dynamical Systems*, Advances in Design and Control, (Society for Industrial and Applied Mathematics, 2005), pp. 173–204, [10.1137/1.9780898718713](https://doi.org/10.1137/1.9780898718713).
 - [23] S. J. Hammarling, Numerical solution of the stable non-negative definite Lyapunov equation Lyapunov equation, *IMA J. Numer. Anal.* **2**, 303 (1982).
 - [24] A. Laub, A Schur method for solving algebraic Riccati equations, *IEEE Trans. Autom. Control* **24**, 913 (1979).
 - [25] A. J. Laub, Invariant subspace methods for the numerical solution of Riccati equations, in *The Riccati Equation*, Communications and Control Engineering Series, edited by S. Bittanti, A. J. Laub, and J. C. Willems (Springer, Berlin, Heidelberg, 1991), pp. 163–196, [10.1007/978-3-642-58223-3_7](https://doi.org/10.1007/978-3-642-58223-3_7).
 - [26] P. Van Dooren, A generalized eigenvalue approach for solving Riccati equations, *SIAM J. Sci. Stat. Comput.* **2**, 121 (1981).

- [27] R. James, J. Langou, and B. R. Lowery, On matrix balancing and eigenvector computation, [arXiv:1401.5766](#).
- [28] B.N. Parlett and C. Reinsch, Balancing a matrix for calculation of eigenvalues and eigenvectors, *Numer. Math.* **13**, 293 (1969).
- [29] M. Șuşcă, V. Mihaly, M. Stănese, and P. Dobra, Iterative refinement procedure for solutions to algebraic Riccati equations, in *2020 IEEE International Conference on Automation, Quality and Testing, Robotics* (2020), pp. 1–5, [10.1109/AQTR49680.2020.9130027](#).
- [30] A. Varga, Computation of irreducible generalized state-space realizations, *Kybernetika* **26**, 89 (1990), [https://dml.cz/handle/10338.dmlcz/124686](#).
- [31] A. Varga and K. H. Fasol, A new square-root balancing-free stochastic truncation model reduction algorithm, *IFAC Proceedings Volumes series 12th Triennial World Congress of the International Federation of Automatic Control. Volume 1 Theory, Sydney, Australia, 18–23 July*, 26, 469 (1993), [10.1016/S1474-6670\(17\)49168-3](#).
- [32] P. Benner, V. Mehrmann, V. Sima, S. Van Huffel, and A. Varga, SLICOT: A subroutine library in systems and control theory, in *Applied and Computational Control, Signals, and Circuits: Volume 1*, edited by B. N. Datta (Birkhäuser, Boston, MA, 1999), pp. 499–539, [10.1007/978-1-4612-0571-5_10](#).
- [33] M. G. Safonov, D. J. N. Limebeer, and R. Y. Chiang, Simplifying the H_∞ theory via loop-shifting, matrix-pencil and descriptor concepts, *Int. J. Control* **50**, 2467 (1989).
- [34] D. Lemmonier and P. Van Dooren, Balancing regular matrix pencils, *SIAM J. Matrix Anal. Appl.* **28**, 253 (2006).
- [35] R. C. Ward, Balancing the generalized eigenvalue problem, *SIAM J. Sci. Stat. Comput.* **2**, 141 (1981).
- [36] Y. Liu and B. Anderson, Singular perturbation approximation of balanced systems, in *Proceedings of the 28th IEEE Conference on Decision and Control*, Vol. 2 (1989), pp. 1355–1360, [10.1109/CDC.1989.70360](#).
- [37] N. A. Bruinsma and M. Steinbuch, A fast algorithm to compute the H_∞ -norm of a transfer function matrix, *Systems and Control Letters* **14**, 287 (1990).
- [38] The solution presented here, adapted from [13], is less general than the solution to the LQG problem, given in Sec. VI C. The Glover-Doyle solution is more general, although it goes beyond the needs of most $\mathcal{H}_\infty$ controllers [16,39].
- [39] J.-L. Hong and C.-C. Teng, A derivation of the Glover-Doyle algorithms for general H_∞ control problems, *Automatica* **32**, 581 (1996).
- [40] P. M. Gahinet, Troubleshooting for state-space H_∞ design, in *1992 American Control Conference* (1992), pp. 634–638, [10.23919/ACC.1992.4792146](#).
- [41] J. Doyle, K. Zhou, and B. Bodenheimer, Optimal control with mixed H_2 and H_∞ performance objectives, in *1989 American Control Conference* (1989), pp. 2065–2070, [10.23919/ACC.1989.4790529](#).
- [42] H.-H. Yeh, S. Banda, and B.-C. Chang, Necessary and sufficient conditions for mixed H_2 and H_∞ optimal control, *IEEE Trans. Autom. Control* **37**, 355 (1992).
- [43] P. Khargonekar and M. Rotea, Mixed H_2/H_∞ control: A convex optimization approach, *IEEE Trans. Autom. Control* **36**, 824 (1991).
- [44] M. Rotea and T. Iwasaki, Generalized H_2/H_∞ control: General formulas, in *Proceedings of the 1995 34th IEEE Conference Decision Control*, Vol. 3 (1995), pp. 3158–3163, [10.1109/CDC.1995.478633](#).
- [45] M. A. Rotea and P. P. Khargonekar, Generalized H_2/H_∞ control, in *Robust Control Theory*, The IMA Volumes in Mathematics and Its Applications, edited by B. A. Francis and P. P. Khargonekar (Springer, New York, NY, 1995), pp. 81–103, [10.1007/978-1-4613-8451-9_5](#).
- [46] V. Kapila and W. M. Haddad, H_2 and mixed H_2/H_∞ stable stabilization, in *Proceedings of the 34th Conference on Decision and Control* (1995), Vol. 2, pp. 1911–1916.
- [47] Y. W. Wang, W. M. Haddad, and D. S. Bernstein, Stable stabilization with H_2 and H_∞ performance constraints, *J. Math. Syst. Estim. Control* **6**, 181 (1996), [https://websites.umich.edu/~dsbaero/library/WangH2Hinfinity1.pdf](#).
- [48] M. Athans, The matrix minimum principle, *Inf. Control* **11**, 592 (1967).