

Optimizing Active Vibration Isolation Systems in
Ground-based Interferometric Gravitational-wave
Detectors

優化地上重力波干涉儀的主動隔震系統

TSANG, Terrence Tak Lun

曾德麟

A Thesis Submitted in Partial Fulfilment
of the Requirements for the Degree of
Doctor of Philosophy
in
Physics

The Chinese University of Hong Kong
September 2022

Thesis Assessment Committee

Professor LI Hua Bai (Chair)

Professor LI Tjonnie Guang Feng (Thesis Supervisor)

Professor CHEN Benmei (Committee Member)

Professor DOOLEY Katherine (External Examiner)

摘要

重力波是由大質量加速物體產生的時空漣漪。重力波可以被檢測為應變的波動。由於其幅度極小，重力波應變很難被探測到。最強的重力波源是黑洞和超新星碰撞等天文事件。這些大約10 到幾千赫茲的重力波會導致大約 10^{-21} 的應變，是一個極其微小的影響。為了檢測如此微量的應變，科學家們建造了世界上最大和最複雜的激光干涉儀。並且，在2015 年9 月14 日，激光干涉重力波天文台（LIGO）的兩個探測器首次探測到重力，開啓了宇宙的新窗口。

為了能夠探測到重力波，降噪在重力波探測器中尤為重要。重力波探測器中有很多噪聲源，例如量子噪聲和熱噪聲。對於基於地面的重力波探測器，地震噪聲，即地面的振動，特別有趣。在重力波檢測頻帶，地震噪聲會在干涉儀的光學系統之間引起不必要的震動，這對應於比重力波引起的應變高幾個數量級。因此，激光干涉儀的核心光學器件需要由彈簧質量系統或多擺錘懸掛，以作於地面和光學器件之間的低通濾波器。在較低頻率下，地震噪聲會激發反射鏡懸架的共振。此外，在0.1 - 0.5 Hz 左右地面有較大幅度的振動，這被稱為二次微震，是海浪引起的地面振動。共振和微震會導致光學器件過度運動，並會使干涉儀錯位。因此，需要對較低頻率的運動進行減震，這是通過主動防震來完成的。主動隔振是指使用位移傳感器和線圈磁鐵執行器等主動元件通過控制系統實現運動抑制。然而，有源組件並不完美，並且包含可能對隔離系統的性能產生不利影響甚至損害動力波探測器靈敏度的自噪聲。因此，它們需要對主動控制器進行適當的設計和優化，以在重力波探測器中實現主動隔離。

神岡重力波探測器(KAGRA) 是日本最近建造的地下低溫重力波探測器。在KAGRA 中，干涉儀的主要光學系統使用了三種類型的主動隔振系統，即A 型、B 型和B_p型懸架。它們為干涉儀中的不同光學元件提供不同尺度的地震隔離。調試這些懸架可能是一項艱鉅的任務。為了正確實施主動隔離系統，確定了幾個子任務。這些包括傳感器校準、傳感器和執行器對齊、懸架的頻率響應建模、設計反饋控制系統以實現光學元件的主動阻尼和粗對準。為了解決這些任務，使用KAGRA 懸架開發並演示了幾種懸架調試技術。本論文中提供的這些方法是穩健且易於使用的基線方法。這些方法的彙編可作為當前和未來專員處理調試任務的基本指南。作為第四次觀測運行(O4) 之前調試的一部分，這些方法已經在一些KAGRA 隔振系統中實施，例如B 型懸架。

A 型懸架和B 型懸架的預隔離階段利用多種類型的傳感器進行主動隔離，包括線性可變差動變壓器(LVDT)、地震檢波器和地面地震儀。它們可以在稱為傳感器融合和傳感器校正的兩種特殊控制方案中相互協調使用。在傳感器融合中，相對位移傳感器(例如LVDT) 和地震檢波器被混合成一個具有整體噪聲性能的“超級傳感器”。並且，在傳

感器校正中，因為相對位移傳感器測量懸掛平台與地面之間的相對位移，放置在懸架附近的地震儀可用於消除相對位移傳感器中的地震噪聲耦合。這兩種方案可以單獨使用（使用傳感器融合或傳感器校正）或一起使用（傳感器校正的相對位移傳感器和地震檢波器之間的傳感器融合）。這些控制策略的成功實施可以降低控制系統中的有效傳感器噪聲，從而有可能實現更好的主動隔離性能。這需要優化這些方案中使用的控制調節器，即互補濾波器和傳感器校正濾波器。這兩個控制優化問題是相似的，因為它們對應於在兩個相互衝突的目標之間尋求最佳權衡。此外，有源隔離系統中單自由度的反饋控制問題具有相似的性質，它尋求干擾抑制和噪聲衰減之間的最佳折衷。因此，這可以使用類似的方法來解決。

為了解決控制調節器設計問題，提出了使用H-infinity (\mathcal{H}_∞) 方法的特殊技術。 \mathcal{H}_∞ 方法是一種尋求閉環廣義系統的 \mathcal{H}_∞ 範數最小化的優化方法。由於 \mathcal{H}_∞ 範數的特性，可以設計廣義對象來指定控制系統中需要最小化的信號的頻率相關上限。為傳感器融合、傳感器校正和反饋控制問題指定了特殊的廣義系統。並且，通過正確選擇與頻率相關的上界，問題的解決方案產生了控制調節器，其最小化目標信號（傳感器噪聲或受控位移），其頻譜是乘法偏移（顯示等於 \mathcal{H}_∞ 範數）。換句話說，傳感器噪聲或反饋控制的位移在某種意義上被最小化，因為它的頻譜在對數刻度的所有頻率上都同樣接近下限。

\mathcal{H}_∞ 方法在KAGRA 中使用信號回收鏡(SRM) 懸架進行測試，這是一種B 型懸架。結果與傳感器融合和傳感器校正的原始KAGRA 實現進行了比較。實驗結果顯示 \mathcal{H}_∞ 方法導致更好的傳感器融合和傳感器校正噪聲性能。特別是，與KAGRA 傳感器校正相比， \mathcal{H}_∞ 傳感器校正對地震噪聲耦合的抑製程度與KAGRA 傳感器校正相似，但在低頻下幾乎沒有地震計噪聲注入。相比之下，KAGRA 傳感器校正顯著的放大了低頻噪聲，使其無法使用。同樣，KAGRA 傳感器融合也引入了來自檢波器的低頻噪聲，但程度較低。相比之下， \mathcal{H}_∞ 傳感器融合與 \mathcal{H}_∞ 傳感器校正相結合，對檢波器噪聲有更好的抑制，同時實現了更好的地震噪聲解耦。這使得 \mathcal{H}_∞ 方法相對適用於傳感器校正和傳感器融合。

至於反饋控制，預隔震器隔震控制的 \mathcal{H}_∞ 優化給出了一個高階控制器，在KAGRA的控制軟件Foton使用出現錯誤，因此不能進行測試。但模擬結果顯示， \mathcal{H}_∞ 控制可以抑制預隔震器的主導運動，同時不會因注入噪聲而破壞高頻下的被動隔震性能。這正是引力波探測器中的主動隔離系統所需要的。 \mathcal{H}_∞ 方法也被用來解決一個更簡單的阻尼控制問題，從而產生了一個更簡單的控制器，可以在KAGRA 控制系統中實現。它與使用臨界阻尼標準設計的控制濾波器進行比較。實驗結果顯示，臨界阻尼濾波器在較高頻率處注入了過多的噪聲，而 \mathcal{H}_∞ 控制在沒有過多噪聲的情況下實現了類似的阻尼水平。 \mathcal{H}_∞ 控制的結果是有希望的，因此它應該進一步開發用於引力波探測器中的主動隔離系統。

Abstract

Gravitational waves are ripples of spacetime produced by massive accelerated objects. Gravitational waves can be detected as a fluctuation in strain, which is a fractional change in distance between two objects. The gravitational-wave strain is very hard to be detected due to its extremely small amplitude. The strongest sources of gravitational waves are astronomical events such as the collision of black holes and supernovae. These events emit gravitational waves around 10 to a few kHz. Detectable gravitational wave strains are around 10^{-21} , which is extremely tiny. To detect such a tiny amount of strain, scientists built the world's largest and most complex laser interferometer. And, on September 14th 2015, the two detectors of the Laser Interferometer Gravitational-Wave Observatory (LIGO) have made the first detection of gravitational waves, which has opened a new window on the universe.

To be able to detect gravitational waves, noise mitigation is particularly important in gravitational-wave detectors. There are many noise sources in a gravitational-wave detector such as quantum noise and thermal noise. For ground-based gravitational-wave detectors, seismic noise, which is the vibration of the ground, is particularly interesting. At the observation band, the seismic noise can cause unwanted motion between the optics of the interferometer. This adds strain noise to the detector that is orders of magnitude higher than that caused by gravitational waves. Therefore, the core optics of the laser interferometer need to be suspended by spring-mass systems or multiple-pendulums, which behave like low-pass filters between the ground and the optics. At lower frequencies, the seismic noise can excite the resonances of the mirror suspensions. Also, the ground hums around 0.1 - 0.5 Hz, which is known as the secondary microseism and is the vibration of the ground caused by ocean waves. The resonances and the microseism can cause excessive motion in the optics and eventually knock the interferometer out of alignment. Therefore, the motion at lower frequencies needs to be damped and this is done by active isolation. Active vibration isolation refers to the use of active components such as displacement sensors and coil-magnet actuators to achieve motion suppression via control systems. However, the active components are not perfect and contain self-noise that can adversely affect the performance of the isolation system or even compromise the sensitivity of the detector at higher frequencies. Hence, they require proper design and optimization of the control laws and digital filters to achieve active isolation in a gravitational-wave detector.

KAGRA is an underground cryogenic gravitational-wave detector recently constructed in Japan. In KAGRA, three types of active vibration isolation systems are used for the main optics of the interferometer, namely, Type-A, Type-B, and Type-Bp suspensions. They provide seismic isolation of different scales for different optics in the interferometer. Commissioning these suspensions can be a daunting task. For proper implementation of an active isolation system, several sub-tasks are identified. These include calibration of sensors, sensors and actuators alignment, frequency response modeling of the suspension, and designing feedback control systems to achieve active damping and coarse alignment of the optics. To tackle these tasks, several suspension commissioning techniques are developed and demonstrated using the KAGRA suspensions. These methods provided in this thesis are conservative approaches that are robust and easy to use. The compilation of the methods serves as a guideline for current and future commissioners for tackling daily commissioning tasks. Nevertheless, the methods have already been implemented in some KAGRA vibration isolation systems, such as the Type-B suspensions, as part of the commissioning for the upcoming fourth observation run (O4).

The pre-isolation stage of the Type-A and Type-B utilizes several types of sensors for active isolation, including linear variable differential transformers (LVDT), geophones, and ground seismometers. They can be used in coordination with each other in two special control schemes called sensor fusion and sensor correction. In sensor fusion, the relative displacement sensor (LVDT) and geophone are blended into a “super sensor” with better overall noise performance. And, in sensor correction, a seismometer placed in the proximity of the suspension can be used to remove the seismic noise coupling in the relative displacement sensors, which measure the relative displacement between a suspended platform and the ground. The two schemes can be used individually (either using sensor fusion or sensor correction) or together (sensor fusion between the sensor-corrected LVDT and the geophone). successful implementation of these control strategies could reduce the effective sensor noise in the control system and hence potentially achieve better active isolation performance. This calls for an optimization of the control filters, i.e. complementary filters and sensor correction filter, that are used in these schemes. The two control optimization problems are similar as they correspond to one that seeks an optimal trade-off between two conflicting objectives. Furthermore, the feedback control problem for a single degree of freedom in an active isolation system is of a similar nature, which seeks the optimal trade-off between disturbance rejection and noise attenuation. Hence, that can be solved using similar approaches.

To solve the control filter design problems, special techniques using the H-infinity (\mathcal{H}_∞) method are proposed. The \mathcal{H}_∞ method is an optimization approach that seeks the minimization of the \mathcal{H}_∞ norm of a closed-loop generalized plant. Due to the property of the \mathcal{H}_∞ norm, the generalized plant can be designed to specify frequency-dependent upper bounds for the signals that need to be minimized in a control system. Special

generalized plants are specified for sensor fusion, sensor correction, and feedback control problems. And, with a proper choice of frequency-dependent upper bounds, the solution to the problems yields control filters that minimize the target signal (sensor noise or controlled displacement) in a way that its spectrum is a multiplicative offset (which is shown to be equal to the \mathcal{H}_∞ norm) from the lower bound. In other words, the sensor noise or feedback-controlled displacement is minimized in the sense that its spectrum is equally close to the lower bound at all frequencies in a logarithmic scale.

The \mathcal{H}_∞ methods are tested with the signal recycling mirror (SRM) suspension, which is a Type-B suspension, in KAGRA. The results are compared to the original KAGRA implementation of sensor fusion and sensor correction. And, results show that the \mathcal{H}_∞ approaches result in better sensor fusion and sensor correction noise performance. In particular, the \mathcal{H}_∞ sensor correction suppresses the seismic noise coupling to a similar extent compared to the KAGRA sensor correction but with virtually no seismometer noise injection at low frequency. In contrast, the KAGRA sensor correction introduces a significant low-frequency noise amplification, which makes it unusable. Similarly, the KAGRA sensor fusion also introduces low-frequency noise from the geophone but to a lower extent. In comparison, the \mathcal{H}_∞ sensor fusion, combined with \mathcal{H}_∞ sensor correction, has a better suppression of the geophone noise while achieving a better seismic noise decoupling. This makes the \mathcal{H}_∞ method relatively suitable for sensor correction and sensor fusion.

As for feedback control, the \mathcal{H}_∞ optimization for seismic isolation control of the pre-isolator gave a high-order control filter that has been misinterpreted by the control software, Foton, used in KAGRA and therefore cannot be tested. However, the simulation result shows that the \mathcal{H}_∞ control can suppress the dominant motion of the pre-isolator while not ruining the passive seismic isolation performance at high frequency by injecting noise. This is exactly what is needed for active isolation systems in gravitational-wave detectors. The \mathcal{H}_∞ method was also used to solve a simpler damping control problem, which resulted in a simpler controller that can be implemented in the KAGRA control system. It is compared to a control filter designed using a critical damping criterion. The result shows that the critical damping filter injects excessive noise at higher frequencies whereas the \mathcal{H}_∞ control achieves a similar level of damping without excessive noise. The results from \mathcal{H}_∞ control are promising and hence it should be further developed for active isolation systems in gravitational-wave detectors.

Acknowledgments

I came from a mechanical engineering background and had zero experience in gravitational-wave physics. Yet, my PhD advisor, Tjonnie Guang Feng Li, chose me as his PhD student, which came as a surprise to me. So, Tjonnie, I am very grateful to you for trusting my capabilities and seeing potential in me. I must also thank you for realizing there was an opportunity for me to work in experimental physics at KAGRA, a facility that I had never heard of. And, thank you for accepting me as the only experimentalist (and as a freak) in your research group. Thank you for supporting me and my research, both mentally and financially, for the last four years.

During my PhD studies, I spent the majority of my time in KAGRA, which is a gravitational-wave detector in Japan. When I first arrived at KAGRA, I was assigned to work with the Type-B team, which worked on installing the Type-B suspensions that suspend the beamsplitter and signal-recycling mirrors. At the time, the group consists of five people: Mark Barton, Enzo Nicolas Tapia San Martin, Ryohei Kozu, Naoatsu Hirata, and Fabián Erasmo Peña Arellano. I was provided the chance to participate in the installation work and I have learned a lot about the physical hardware from these people. So, Mark, Enzo, Kozu-Kun, Hirata-san, and Fabian, I must thank you for showing me the inner workings of KAGRA. Fabian, you have been one of my closest friends and you deserve special thanks. Thank you for the endless discussions about various things. You are an extremely humble yet knowledgeable person and I have learned a lot from you. Your extreme optimism is infectious and it has helped me a lot during stressful times. You have been extremely supportive but are also capable of giving constructive criticism. You have been a great friend, or perhaps, a great mentor to me and I would not have done this without you. Thank you very much.

Many people have worked on vibration isolation research at KAGRA before me. My work is based on the previous results and I am standing on the shoulders of giants, so to speak. So, I would like to thank the giants including Takanori Sekiguchi (Although I have not met you in person, your thesis on KAGRA suspensions has been very helpful to my research.), Koki Okutomi, and Yoshinori Fujii.

Thanks should also go to all people that I have worked with in KAGRA. This includes Yoichi Aso, Lucia Trozzo, Eleonora Capocasa, Ryutaro Takahashi, Ayaka Shoda, Takafumi Ushiba, Keiko Kokeyama, Masayuki Nakano, Yutaro Enomoto, Takahiro Ya-

mamoto, Tomo Akutsu, Simon Zeidler, Matteo Leonardi, Kouseki Miyo, Dan Chen, Darkhan Tuyenbayev, Naoki Aritomi, Marc Eisemann, Michael Page, Yuhang Zhao, Yohei Nishino, and Masahide Tamaki.

I am very grateful to Aso-san, who has been the host researcher for my visits. Thank you for discussing vibration isolation research with me and entrusting me with the KAGRA suspensions. And, thank you for dealing with all the tedious administrative stuff to get me from Hong Kong to KAGRA, especially for the one during the pandemic. I would like to extend my sincere thank to Lucia Trozzo, who has kick-started my research by teaching me the concept of sensor fusion. I must also thank you for your support in my research and for all the control theory and vibration isolation discussions. Special thanks should also be given to Kokeyama-san. Thank you for appreciating me and being supportive to me. I would like to express my deepest appreciation to unsung heroes who have worked endless hours in the KAGRA control room with the simple goal of getting KAGRA to work. Nakano-san, Enomoto-san, Yamamoto-san, Ushiba-san, and others who I might have left out, thank you for your hard work. KAGRA would not have functioned without you guys. I gratefully acknowledge the help from Eri Sakamoto, Mihoko Okinaka, Masafumi Kurachi, Satoru Ikeda, and Yasui Hiromi, who are the administrative/technical staff in KAGRA. You have made my life in Japan so much happier and easier.

I cannot leave the Chinese University of Hong Kong without thanking Otto Akseli Hannuksela, Isaac Chun Fung Wong, Adrian Ka Wai Chung, Patrick Chi Kit Cheong, Alan Tsz Lok Lam, Harry Ho Yin Ng, Juno Chun Lung Chan, Eungwang Seo, Liiyung Yeow, Ania Liu, Hemantakumar Phurailatpam, Anson Yip Ka Long, Amit Jit Singh, Nidhi Pant, and others, who are people that I have met in the gravitational wave research group lead by Tjonnje (and recently Otto). You guys are all very knowledgeable and I wish I could have learned more from you. I would also like to acknowledge the help from Jason Yee Ching Lam and Angus Ho Ting Wong, who are undergraduate students of CUHK. Thank you for helping me to build my suspension prototype at CUHK.

At last, I would like to thank Katherine Dooley for inviting me for giving a talk to the Einstein Telescope instrumentation board about my optimal sensor fusion method. This would be my first time communicating my research with the experts in the gravitational-wave community. Giving a talk has surely helped me realize the importance of my work and has helped me regain motivation during the pandemic.

Contents

Abstract	iv
Acknowledgments	x
Table of Contents	xiii

I Introduction to experimental gravitational-wave physics

1 Gravitational Waves	3
1.1 Gravitational waves	4
1.2 Gravitational wave sources	7
1.3 A brief history of gravitational wave detection	8
2 Gravitational-Wave Detector	11
2.1 Michelson interferometer	11
2.1.1 Photon shot noise	14
2.1.2 Radiation pressure noise	16
2.1.3 Optimal arm length	18
2.2 Optical cavity	20
2.2.1 Folded Interferometer	20
2.2.2 Fabry-Perot cavity	21
2.2.3 Power and signal recycling cavities	24
2.3 Noises in gravitational-wave detectors	25
2.3.1 Quantum noise	25
2.3.2 Thermal noise	26
2.3.3 Seismic noise	26
3 Vibration Isolation	29
3.1 Passive isolation	30
3.1.1 Oscillators: spring-mass and pendulums	30
3.1.2 Anti-springs	32
3.1.3 Cascaded spring-mass systems	35
3.2 Active isolation	38

3.2.1	Interlude: Transfer functions	38
3.2.2	Feedforward control	39
3.2.3	Feedback control	41

II Active vibration isolation systems in KAGRA

4	KAGRA Vibration Isolation Systems	47
4.1	Types of vibration isolation systems	47
4.2	Suspension stages	50
4.2.1	Pre-isolator	50
4.2.2	Standard filters	52
4.2.3	Marionette and intermediate stages	53
4.2.4	Test mass	54
4.3	Sensing and actuation technologies	55
4.3.1	Linear variable differential transformer (LVDT)	55
4.3.2	Geophone and inertial sensors	57
4.3.3	Optical sensing and electromagnetic actuator (OSEM)	58
4.3.4	Photocoupler position sensor (photosensor)	59
4.3.5	Optical lever	60
5	Control Topology	65
5.1	Control basis	69
5.2	Control problem formulation	70
6	Suspension Commissioning Techniques	77
6.1	Sensor calibration	78
6.1.1	Sigmoidal calibration curve	78
6.1.2	Geophone calibration	86
6.1.3	Inter-calibration of sensors	92
6.2	Control basis alignment	92
6.2.1	Sensing matrices	93
6.2.2	Actuation matrices	108
6.3	Modeling and fitting	112
6.3.1	Modeling frequency response data with transfer functions	113
6.3.2	Frequency spectrum modeling	123
6.4	Coarse alignment and local damping control	127
6.4.1	Damping control	128
6.4.2	Coarse alignment control	134
6.4.3	Post-filtering	140

III \mathcal{H}_∞ optimization of control filters in active vibration isolation systems

7	Active Isolation Control Filters	155
7.1	Sensor fusion using complementary filters	158
7.2	Sensor correction for relative displacement sensors	163
7.3	Feedback control filters	166
8	\mathcal{H}_∞ Optimization of Control Filters	169
8.1	\mathcal{H}_∞ method	169
8.1.1	Formalism	169
8.1.2	Solving the \mathcal{H}_∞ problem	172
8.2	Filter design problem as an \mathcal{H}_∞ problem	175
8.2.1	Sensor fusion problem	175
8.2.2	Sensor correction problem	182
8.2.3	Feedback control problem	184
8.3	\mathcal{H}_∞ optimization of the SRM pre-isolator control filters	192
8.3.1	The SRM pre-isolator	193
8.3.2	\mathcal{H}_∞ sensor fusion	200
8.3.3	\mathcal{H}_∞ sensor correction	205
8.3.4	\mathcal{H}_∞ sensor fusion with sensor correction	209
8.3.5	\mathcal{H}_∞ sensors closed-loop performance and summary	212
8.3.6	\mathcal{H}_∞ control	216
9	Concluding Remarks	225
9.1	Discussions	225
9.1.1	On \mathcal{H}_∞ methods	225
9.1.2	On inertial sensor pre-filtering	226
9.1.3	On the implementation of inertial sensors	230
9.1.4	On the results	231
9.2	Future prospects	231

Part I

Introduction to experimental gravitational-wave physics

Chapter 1

Gravitational Waves

In 1915, Albert Einstein published the general theory of relativity, which redefines Newton’s law of universal gravitation. General relativity describes gravity as a geometry property of the four-dimensional spacetime, i.e. curvature of spacetime. The curvature of spacetime is a result of the presence of matter and radiation. This relationship between spacetime curvature and the energy and momentum of the matter and radiation is given by the Einstein field equations [1], which is a set of ten coupled partial differential equations written in the form

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (1.1)$$

The left-hand side of Eqn. (1.1) is the Einstein tensor $G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R$ and it only depends on terms that describe the geometry of spacetime. The stress-energy tensor $T_{\mu\nu}$ on the right-hand side of Eqn. (1.1) describes the density and flux of energy momentum. While the mathematics of Einstein’s general theory of relativity can be difficult, it can be summarized in just 12 words: “*Spacetime tells matter how to move; matter tells spacetime how to curve.*” [2]

In 1916, Einstein predicted the existence of gravitational waves as a solution to the linearized Einstein field equations under the weak-field limit, i.e. when spacetime is nearly flat [3, 4]. In the weak-field limit, the Einstein field equations have solutions corresponding to transverse waves of spatial strain that travels at the speed of light and they are generated by time variation of the mass quadrupole moment of the source. The gravitational-wave strain has a remarkably small amplitude, thus making them very hard to detect. In this chapter, the nature of the gravitational waves is discussed, to an introductory extent following Ref. [5]. Detailed explanations can be found in Refs. [6, 7]. In this chapter, Sec. 1.1 provides a brief derivation of the gravitational wave under the weak-field limit. And Sec. 1.2 discusses some possible astrophysical sources of gravitational waves. Lastly, Sec. 1.3 briefly introduces the history of gravitational wave detection.

1.1 Gravitational waves

Gravitational waves are tiny ripples of spacetime. Here, spacetime refers to the combination of the three-dimensional space x, y, z , and the one-dimensional time t and it is first introduced in the theory of special relativity. Analogous to the three spatial coordinate (x, y, z) , a point in spacetime coordinate (t, x, y, z) is called an event. The combination of space and time allows one to defined a so-called spacetime interval ds , which is invariant under coordinate transformations, between two neighboring events such that

$$ds^2 = -c^2 dt^2 + dx^2 + dy^2 + dz^2, \quad (1.2)$$

where c is the speed of light. Eqn. (1.2) can be generalized as

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \quad (1.3)$$

where $g_{\mu\nu}$ is the metric tensor, μ and ν are indices that range from 0 to 3 such that $(ct, x, y, z) = (x^0, x^1, x^2, x^3)$ (they are indices, not exponents). Note that the Einstein summation convention is used here where a repeated index implies a summation over all possible values of the index, i.e. $a_i b^i = \sum_i a_i b^i$. It is easy to see that Eqn. (1.2) can be recovered by setting the metric tensor as $g_{\mu\nu} = \eta_{\mu\nu}$, where

$$\eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (1.4)$$

is called the Minkowski metric. The Minkowski metric denotes a special case of spacetime corresponding to the flat spacetime. Any spacetime that is not flat corresponds to a metric $g_{\mu\nu} \neq \eta_{\mu\nu}$ and hence the metric tensor $g_{\mu\nu}$ encodes the information about the curvature of spacetime.

To derive gravitational-waves, consider a small perturbation of the the flat spacetime. In this case, the metric can be written in the form

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.5)$$

where $|h_{\mu\nu}| \ll 1$ and $h_{\mu\nu}$ represents the small metric perturbation away from Minkowski space. Such configuration is called weak-field limit. Because of this, the second-order terms of $h_{\mu\nu}$ can be discarded and the Einstein Tensor $G_{\mu\nu}$ and the stress-energy tensor $T_{\mu\nu}$ can be expressed in linear terms of $h_{\mu\nu}$. In this case, the Einstein field equations (1.1) is said to be linearized. Now, under the choice of a special coordinates marked out by the

worldlines of a free-falling mass¹, the linearized Einstein field equations in vacuum can be written as

$$\left(\nabla^2 - \frac{1}{c^2} \frac{d}{dt}\right) h_{\mu\nu} = 0, \quad (1.6)$$

which is a wave equation. Now, recall that $h_{\mu\nu}$ is a small perturbation of the flat spacetime. The fact that it satisfies the wave equation means that the small perturbation can propagate through space at the speed of light and these perturbations are known as gravitational waves.

For simplicity, assume that the gravitational wave propagates along the z direction, the transverse traceless gauge requires the perturbation to take the form

$$h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (1.7)$$

This gives two polarizations, $+$ and \times , for the gravitational wave, one along

$$\hat{h}_+ = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (1.8)$$

and the other along

$$\hat{h}_\times = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (1.9)$$

To see what the physical effect of the polarization is, consider a gravitational wave polarized in the \hat{h}_+ direction so the $h_\times = 0$ and $h_+ = h e^{i(\omega t - kz)}$, where ω is the angular frequency of the gravitational wave and k is the wavenumber. The spacetime interval can then be written as

$$ds^2 = -c^2 dt^2 + (1 + h_+) dx^2 + (1 - h_+) dy^2 + dz^2. \quad (1.10)$$

Suppose there is a free-falling particle located at $x = L \cos \theta$, $y = L \sin \theta$, and $z = 0$, where θ is the angle between the position vector of the particle and the x -axis. The

¹This corresponds to fixing the so-called gauge freedom using the transverse traceless gauge (TT gauge).

proper distance between the particle the origin can then be evaluated as

$$\begin{aligned}
 s &= \int ds = \int \sqrt{(1+h_+)dx^2 + (1-h_+)dy^2} \\
 &\approx \int \left(1 + \frac{1}{2}h_+ \cos(2\theta)\right) dr \\
 &= L \left(1 + \frac{1}{2}h \cos(2\theta)e^{i\omega t}\right).
 \end{aligned} \tag{1.11}$$

Now, instead of one particle, suppose there is a ring of free-falling particles populating the full 2π angle of θ . An observer in the proper frame measuring the distance between the particle and the origin along the x -axis, i.e. $\theta = 0, \pi$, with respect to the origin would be oscillating, with an expansion followed by a contraction. However, the proper distance of a particle along the y -direction, i.e. $\theta = \pi/2, -\pi/2$, would be a oscillation with a contraction followed by an expansion, which is the opposite of that in the x -direction. A similar analysis could be done for the \hat{h}_\times polarization and would result in a proper distance that reads $L[1 + (1/2)h \sin(2\theta)e^{i\omega t}]$. In this case, the expansion-contraction principle axis along the $\theta = \pi/4, \pi + \pi/4$ axis while the contraction-expansion principle axis is along the $\theta = -\pi/2, \pi/2 + \pi/4$ axis, i.e. along the cross directions hence the name h_\times . The time evolution of the ring of test masses over a period of $2\pi/\omega$ is shown in Fig. 1.1 for the two polarization modes.




Figure 1.1: Ring of test masses with gravitational waves going into the paper at time stamps $t = 0, \frac{\pi}{2\omega}, \frac{\pi}{\omega}, \frac{3\pi}{2\omega}, \frac{2\pi}{\omega}$. Top: + polarization. Bottom: \times polarization.

If the observer measures the radius of the ring along the direction of the expansion, and compare the measurements at time $t = 0$ and $t = \pi/2\omega$, the distance would change by a fraction $\Delta l = lh/2$, where l is the original distance. Rewriting the relationship gives

$$h = \frac{2\Delta l}{l}, \tag{1.12}$$

which is similar to the definition of strain, which is $\Delta l/l$. From here, it can be seen that the gravitational wave is a transverse wave of strain. This makes instruments that can measure a fractional change in distance, such as a Michelson interferometer, a good candidate for direct detection of gravitational waves.

1.2 Gravitational wave sources

As mentioned previously, gravitational waves are generated by sources whose mass quadrupole moments vary in time. The metric perturbation from a source is given by

$$h_{\mu\nu} = \frac{2G}{Rc^4} \ddot{I}_{\mu\nu}, \quad (1.13)$$

where G is the gravitational constant, R is the distance between the observer and the source, and $I_{\mu\nu}$ is the reduced quadrupole moment defined by

$$I_{\mu\nu} = \int dV \left(x_\mu x_\nu - \frac{1}{3} \delta_{\mu\nu} r^2 \right) \rho(\mathbf{r}), \quad (1.14)$$

where V is the volume, $\delta_{\mu\nu}$ is the Kronecker delta and $\rho(\mathbf{r})$ is the mass distribution. The reason why the dominating component is the quadrupole term, as oppose to the dipole term in electromagnetism, is due to the conservation of momentum and angular momentum [5]. Suppose there are two point masses with mass M and orbiting along the z -axis with an orbital frequency of f_{orb} and a radius of r_0 . The quadrupole moments can then be straightforwardly calculated as

$$I_{xx} = 2Mr_0^2 \left(\cos^2(2\pi f_{\text{orb}}t) - \frac{1}{3} \right), \quad (1.15)$$

$$I_{yy} = 2Mr_0^2 \left(\sin^2(2\pi f_{\text{orb}}t) - \frac{1}{3} \right), \quad (1.16)$$

and

$$I_{xy} = 2Mr_0^2 \cos(2\pi f_{\text{orb}}t) \sin(2\pi f_{\text{orb}}t). \quad (1.17)$$

Plugging these into Eqn. (1.13) yields

$$h_{xx} = -h_{yy} = \frac{32\pi^2 G}{Rc^4} Mr_0^2 f_{\text{orb}}^2 \cos(4\pi f_{\text{orb}}t), \quad (1.18)$$

and

$$h_{xy} = h_{yx} = \frac{32\pi^2 G}{Rc^4} Mr_0^2 f_{\text{orb}}^2 \sin(4\pi f_{\text{orb}}t). \quad (1.19)$$

As can be seen, the amplitude of the gravitational wave scales as

$$h = \frac{32\pi^2 G}{Rc^4} M r_0^2 f_{\text{orb}}^2. \quad (1.20)$$

Suppose there there are there is a binary system of two neutron stars with masses at the Chandrasekhar limit of about $M = 3 \times 10^{30}$ kg. They are orbiting in such that they are almost touching. This gives a orbiting radius of around $r_0 = 20$ km and a frequency of $f_{\text{orb}} = 400$ Hz. They are located at the Virgo Cluster, which is the nearest galaxy clusters to the Earth, at around $R = 4.5 \times 10^{23}$ m, corresponding to 15 Mpc (Megaparsec). This gives an gravitational wave amplitude of

$$h \approx 10^{-21}, \quad (1.21)$$

which is extremely small.

As can be seen, gravitational waves are very tiny. While gravitational waves can easily be produced by massive objects that has a changing quadrupole moment, only systems that are massive, compact, or violent enough can produce a gravitational field strong enough that can be observed with current technology. This means that terrestrial sources are unlike to be detected. Fortunately, there are several candidates of astronomical gravitational wave sources. Example includes the aforementioned binary neutron star mergers, binary black hole mergers, core-collapse supernovae, and stochastic background. They are only mentioned here and details are given in Refs. [5, 6].

1.3 A brief history of gravitational wave detection

In 1974, the first indirect detection of gravitational waves was made via the discovery of a binary pulsar system (PSR B1913+16) [8], which has a decaying orbital rate matching that predicted by general relativity as a loss of energy due to the emission of gravitational wave. The first direct detection of of gravitational waves was made in 2015 by the two detectors of the Laser Interferometer Gravitational-Wave Observatory (LIGO), which are located in the US [9]. The two detectors simultaneously observed a transient gravitational-wave signal (GW150914), with an increasing frequency from 35 Hz to 250 Hz with a peak gravitational-wave strain of 10^{-21} . The signal has lasted for about 0.2 seconds and it matches the waveform predicted by general relativity for the inspiral and merger of a pair of black holes and the ringdown of the remnant black hole.

The first detection of gravitational waves opens up a unique window on the universe. Following the first detection, two more detection of gravitational waves were made during the first observation run (O1). The Virgo detector, which is located in the Italy, joined the second observing run, O2, and has detected a gravitational wave generated by binary

neutron star mergers (GW170817). The particular event has an electromagnetic counterpart in the form of gamma ray burst and it was detected by the Fermi Gamma-Ray Burst Monitor (GRB 170817A) 1.7 seconds after the coalescence of the neutron stars [10]. The close arrival times of the two waves can be used to place strong constraints on the fundamental physics of gravity, including the speed of gravitational wave. In fact, the difference between the speed of gravity and the speed of light was constrained to be between -3×10^{-15} and $+7 \times 10^{-16}$ times the speed of light, improving the previous estimate by around 14 orders of magnitude [11]. The LIGO and Virgo detectors continued to operate for a third observation run (O3) in 2019 after substantial upgrades between the runs. Due to the improved sensitivities of the detector networks, the total number of detected gravitational wave candidates has been increase from 11 (after the end of O2) to 90 (after the end of O3) [12].

KAGRA is a new underground cryogenic gravitational-wave detector located in Japan [13, 14]. Different from LIGO and Virgo, KAGRA implements two new technologies that are important for the third-generation gravitational-wave detectors such as the Einstein Telescope [15], that is, being built underground and operating at cryogenic temperature. These new technologies enable the detector to run at higher sensitivities. The KAGRA detector has been constructed and still being developed to reach the designed sensitivity [14]. In 2020, KAGRA and GEO 600, a gravitational-wave detector in Germany [16], conducted an two-week observing run (O3GK), during the suspension of the O3 [17]. The binary neutron star range of the two detectors were 0.6 megaparsec and 1 parsec for KAGRA and GEO 600, respectively. But, no observations were made. Upgrades of the KAGRA detector is still ongoing and it is planned to join the forth observing run (O4), which is planned to be commenced in March 2023.

Chapter 2

Gravitational-Wave Detector

The gravitational-wave detector is instrument designed to achieve an outstanding sensitivity such that it is capable of detecting a fractional change in length in the order of 10^{-21} . Current gravitational-wave detectors are interferometers that has a complicated optical configuration. But, the basic principles of a gravitational-wave detector can be illustrated by just a few fundamental components, such as the Michelson interferometer and the Fabry-Perot cavity. The Michelson interferometer and Fabry-Perot cavity are explained in Sec. 2.1 and 2.2, respectively. In Sec. 2.3, several noise sources of the gravitational-wave detector are also discussed.

2.1 Michelson interferometer

A gravitational wave causes space to contract and expand in the two transverse directions. To detect gravitational waves, this calls for a measurement of differential length changes in two orthogonal directions. One particular instrument fulfilling this purpose is the Michelson interferometer. The Michelson interferometer is originally used to detect the relative motion between the Earth and the luminiferous aeither [18]. The simplest version of the interferometer shown in Fig. 2.1 is only composed of 5 components, a laser, a beamsplitter, two mirrors, and a photodiode.

In a Michelson interferometer, the laser is divided into two beams traveling along orthogonal paths, say along the x and y directions. The two beams travel some distances l_x and l_y before reaching the mirrors, which are denoted test mass X and test mass Y. They are reflected by the test masses and they travel the same distances before the recombination at the beamsplitter. The recombined beam is an interference of the two beams and it is absorbed by the photodiode, which measures the light power of the recombined beam. Suppose, right before beam travels through the beamsplitter, the electric field of the laser reads

$$E(x, t) = E_0 e^{i(\omega t - kx)}, \quad (2.1)$$




Figure 2.1: A simple Michelson interferometer.

where E_0 is the amplitude of the electric field, ω is the angular frequency, t is time, and k is the wave number. The returned beam at the photodiode traveled along the x direction reads

$$E_x(t) = \frac{i}{2} E_0 e^{i(\omega t - 2kl_x)}, \quad (2.2)$$

picking up an amplitude transmission coefficient of $i/\sqrt{2}$ and an amplitude reflection coefficient of $1/\sqrt{2}$ from the beamsplitter. Similarly, the beam traveled along the y direction reads

$$E_y(t) = \frac{i}{2} E_0 e^{i(\omega t - 2kl_y)}. \quad (2.3)$$

The electric field at the output, i.e. the photodiode, reads

$$\begin{aligned} E_{\text{out}}(t) &= E_x(t) + E_y(t) \\ &= \frac{i}{2} E_0 e^{i(\omega t - 2kl_x)} + \frac{i}{2} E_0 e^{i(\omega t - 2kl_y)} \\ &= \frac{i}{2} E_0 e^{i\omega t} (e^{-i2kl_x} + e^{-i2kl_y}) \\ &= \frac{i}{2} E_0 e^{i(\omega t - kl_x - kl_y)} (e^{ik(l_x - l_y)} + e^{-ik(l_x - l_y)}) \\ &= iE_0 e^{i(\omega t - kl_x - kl_y)} \cos(k(l_x - l_y)) \end{aligned} \quad (2.4)$$

The power detected by the photodiode is proportional to $|E_{\text{out}}(t)|^2$ and the power at the

input is proportional to $|E(x, t)|^2$ with the same proportionality. This gives

$$\begin{aligned} \frac{P_{\text{out}}}{P_{\text{in}}} &= \frac{E_0^2 \cos^2(k(l_x - l_y))}{E_0^2} \\ P_{\text{out}} &= \frac{P_{\text{in}}}{2} [1 + \cos(2k(l_x - l_y))] , \end{aligned} \quad (2.5)$$

where P_{out} is the power readout of the photodiode and P_{in} is the laser power at the input.

If the Michelson interferometer is used as a gravitational-wave detector, each mirror needs to be a free-falling mass and instead of having it fixed on a rigid structure as in the original Michelson-Morley experiment. This is by decoupling the mirrors from the ground motion and the details are given in Chapter 3. With the mirrors suspended, they can be thought as two of the test masses in a ring of test masses shown in Fig. 1.1. As gravitational wave propagates through the ring, assuming that the interferometer arms are aligned in with the polarization of the gravitational wave, one arm gets compressed and the other gets stretched. This induces a differential change in arm length. Suppose the gravitational wave caused a length change of $\Delta l/2$ in the x direction and a change of $-\Delta l/2$ in the y direction. The Michelson interferometer readout (2.5) becomes

$$P'_{\text{out}}(\Delta l) = \frac{P_{\text{in}}}{2} [1 + \cos(2k(l_x - l_y + \Delta l))] . \quad (2.6)$$

The differential change in length Δl is a small quantity. Therefore, the power readout can be written as a first-order approximation

$$P'_{\text{out}}(\Delta l) = P_{\text{out}}(0) + \left. \frac{dP_{\text{out}}(\Delta l')}{d\Delta l'} \right|_{\Delta l'=0} \Delta l , \quad (2.7)$$

which gives

$$P'_{\text{out}}(\Delta l) = \frac{P_{\text{in}}}{2} [1 + \cos(2k(l_x - l_y)) - 2k\Delta l \sin(2k(l_x - l_y))] . \quad (2.8)$$

As can be seen, the Michelson interferometer has effectively become transducer of gravitational wave signal Δl with a calibration factor of $-kP_{\text{in}} \sin(2k(l_x - l_y))$.

Here, the difference in the interferometer arm lengths $l_x - l_y$ is a design factor. Intuitively, it would seem reasonable to put the initial lengths such that $\sin(2k(l_x - l_y)) = 1$, giving

$$\begin{aligned} l_x - l_y &= \frac{1}{2k} \left(\frac{\pi}{2} \pm n\pi \right) \\ &= \frac{\lambda}{4} \left(\frac{1}{2} \pm n \right) , \end{aligned} \quad (2.9)$$

where the wave number is $k = 2\pi/\lambda$, λ is the wavelength of the laser, and $n = 0, 1, 2, \dots$

This would put the detector at an operating point where it is most sensitive to the gravitational wave and this particular operating point is called the mid-fringe. However, this is not true when readout noise comes into consideration.

2.1.1 Photon shot noise

One fundamental readout noise of the interferometer is the photon shot noise. The photon shot noise is due to the quantum nature of light. Measuring the laser power at the output of the Michelson interferometer is effectively the same as counting the number of photons arriving at the photodiode. The number of photon arrival is not a constant but is a random variable that follows a probability distribution called the Poisson distribution. The Poisson distribution can be estimated as a Gaussian distribution with a standard deviation of

$$\sigma_{N_{\text{photon}}} = \sqrt{N_{\text{photon}}}, \quad (2.10)$$

where N_{photon} is the mean number of photon arriving at the photodiode. Now, the light power is related to the mean number of photon by the Planck-Einstein relation

$$P_{\text{out}} = N_{\text{photon}} \hbar \omega, \quad (2.11)$$

where \hbar is the reduced Planck constant. The fluctuation in laser power at the mid-fringe is then

$$\begin{aligned} \sigma_{P_{\text{mid}}} &= \sigma_{N_{\text{photon}}} \hbar \omega \\ &= \sqrt{N_{\text{photon}}} \hbar \omega \\ &= \sqrt{\frac{P_{\text{in}}}{2\hbar\omega}} \hbar \omega \\ &= \sqrt{\frac{P_{\text{in}} \hbar \omega}{2}}, \end{aligned} \quad (2.12)$$

where the mean number of photon is approximated from the power readout with the absence of the gravitational-wave signal, i.e. $N_{\text{photon}} = P_{\text{in}}/(2\hbar\omega)$. The fluctuation in power cannot be distinguished from the change in power caused by the change in arm length Δl , which is caused by the gravitational wave. The fluctuation in power can be converted to an equivalent change in arm length via the calibration factor, which is kP_{in} at the mid-fringe. Dividing the change in arm length by the common arm length l of the interferometer (assuming $l_x - l_y$ is much smaller than l) gives a strain fluctuation at

mid-fringe of

$$\begin{aligned}
 \sigma_{h_{\text{mid}}} &= \frac{1}{l} \frac{\sigma_{P_{\text{mid}}}}{kP_{\text{in}}} \\
 &= \frac{1}{l} \sqrt{\frac{P_{\text{in}} \hbar \omega}{2(2\pi/\lambda)^2 P_{\text{in}}^2}} \\
 &= \frac{1}{l} \sqrt{\frac{\hbar c \lambda}{4\pi P_{\text{in}}}}.
 \end{aligned} \tag{2.13}$$

Now, instead of the mid-fringe, consider the operating point near the dark fringe. Here, the dark fringe is defined the operating point where the Michelson readout is null. This is achieved by setting the difference in arm length by

$$\begin{aligned}
 l_x - l_y &= \frac{1}{2k} (\pi \pm 2n\pi) \\
 &= \frac{\lambda}{4} (1 \pm 2n),
 \end{aligned} \tag{2.14}$$

where $n = 0, 1, 2, \dots$. Suppose the operating point is set to some distance d from the dark fringe, where $d \ll \lambda$. The power readout is then

$$P_{\text{dark}}(\Delta l) = P_{\text{in}} [1 - \cos(2kd) + 2k\Delta l \sin(2kd)]. \tag{2.15}$$

And, expanding up to a second order of d gives

$$P_{\text{dark}}(\Delta l) = P_{\text{in}} [2k^2 d^2 + 4k^2 d \Delta l]. \tag{2.16}$$

This gives a calibration factor of $4k^2 d P_{\text{in}}$. Without the presence of the gravitational wave, the power detected by the photodiode is $2k^2 d^2 P_{\text{in}}$, which converts to a mean number of photon of $N_{\text{photon}} = 2k^2 d^2 P_{\text{in}} / (\hbar \omega)$. This gives the power fluctuation at the dark fringe

$$\begin{aligned}
 \sigma_{P_{\text{dark}}} &= \sqrt{N_{\text{photon}} \hbar \omega} \\
 &= \sqrt{\frac{2k^2 d^2 P_{\text{in}}}{\hbar \omega} \hbar \omega} \\
 &= kd \sqrt{2P_{\text{in}} \hbar \omega}.
 \end{aligned} \tag{2.17}$$

The fluctuation of interpreted strain when operating at the dark fringe is then

$$\begin{aligned}
 \sigma_{h_{\text{dark}}} &= \frac{1}{l} \frac{\sigma_{P_{\text{dark}}}}{4k^2 d P_{\text{in}}} \\
 &= \frac{1}{l} \frac{kd \sqrt{2P_{\text{in}} \hbar \omega}}{4k^2 d P_{\text{in}}} \\
 &= \frac{1}{2l} \sqrt{\frac{\hbar c \lambda}{4\pi P_{\text{in}}}}.
 \end{aligned} \tag{2.18}$$

Comparing Eqn. (2.13) and (2.18), it is obvious that the strain fluctuation, and hence strain noise, at the dark fringe is smaller by a factor of 2. For this reason, the operating point of a gravitational wave detector is chosen to be at the dark fringe so it is insensitive to important noise sources [5]. In this sense, the gravitational-wave detector is said to be a null instrument. However, the difference in arm length $l_x - l_y$ cannot not be zero for the same reason that the gravitational wave detector needs to be sensitive to gravitational waves. The difference is known as the Schnupp asymmetry and more details are given in Ref. [19].

The amplitude spectral density $h_{\text{shot}}(f)$ of the shot noise, with a standard deviation of $\sigma_{h_{\text{shot}}}$ can be obtained by the following relationship

$$\int_0^{\frac{1}{2}} h_{\text{shot}}(f_n)^2 df_n = \sigma_{h_{\text{shot}}}^2, \quad (2.19)$$

where f_n is the normalized frequency, which is the frequency divided by the sampling frequency. Note that the shot noise $h_{\text{shot}}(f)$ is a white noise, meaning that it has no frequency dependency. The frequency dependency (f) is kept here to indicate that it is an amplitude spectral density. Using the shot noise at dark fringe, i.e. $\sigma_{h_{\text{shot}}} = \sigma_{h_{\text{dark}}}$, this gives

$$\begin{aligned} h_{\text{shot}}(f) &= \sqrt{2}\sigma_{h_{\text{dark}}} \\ &= \frac{\sqrt{2}}{2l} \sqrt{\frac{\hbar c \lambda}{4\pi P_{\text{in}}}}, \end{aligned} \quad (2.20)$$

i.e. picking up a factor of $\sqrt{2}$ compared to Eqn. (2.18).

2.1.2 Radiation pressure noise

The shot noise derived in Eqn. (2.13) and (2.18) is inversely proportional to the square root of the power of the injected laser, i.e. $\propto \sqrt{1/P_{\text{in}}}$. In order to reduce strain noise, one would be tempted to increase the laser power in exchange for a reduction of photon shot noise. However, even if it is possible to increase the laser power indefinitely, there exists another fundamental quantum noise that would suggest otherwise. The noise is called the radiation pressure and like the photon shot noise, it is a quantum noise due to the quantum nature of light.

In the Michelson interferometer, the laser is reflected by the mirror, which is assumed to be free falling. Assuming that the light is reflected completely by the test mass, the light exerts a force P/c on the mass, where P is the power of the light. Now, recall that the power fluctuation is due to the shot noise is $\sqrt{N_{\text{photon}}}\hbar\omega$ in Eqn. (2.12), and with an average power $P_{\text{in}}/2$ in one arm, the power fluctuation of the laser in one arm happens to

be exactly that at the mid-fringe, i.e.

$$\sigma_{P_{\text{rad}}} = \sqrt{\frac{P_{\text{in}}\hbar\omega}{2}}. \quad (2.21)$$

It follows that the fluctuation in force exerting on the mirror is

$$\begin{aligned} \sigma_{F_{\text{rad}}} &= \frac{\sigma_{P_{\text{rad}}}}{c} \\ &= \sqrt{\frac{P_{\text{in}}\hbar\omega}{2c^2}} \\ &= \sqrt{\frac{\pi P_{\text{in}}\hbar}{c\lambda}} \end{aligned} \quad (2.22)$$

The corresponding amplitude spectral density of the force fluctuation $F(f)$ due to radiation pressure is obtained by the same procedure in Eqn. (2.20), which gives

$$F(f) = \sqrt{\frac{2\pi P_{\text{in}}\hbar}{c\lambda}}. \quad (2.23)$$

The amplitude spectral density of the mirror displacement $X(f)$ caused by the radiation pressure can be obtained from the Newton's second law in the frequency domain,

$$\begin{aligned} m(2\pi f)^2 X(f) &= F(f) \\ X(f) &= \frac{1}{m(2\pi f)^2} \sqrt{\frac{2\pi P_{\text{in}}\hbar}{c\lambda}} \\ &= \frac{1}{mf^2} \sqrt{\frac{P_{\text{in}}\hbar}{8\pi^3 c\lambda}}, \end{aligned} \quad (2.24)$$

where m is the mass of the mirror. Finally, the strain noise due to the radiation pressure $h_{\text{rad}}(f)$ can be obtained by dividing the mirror displacement by the length of the interferometer arm and multiplying by a factor of 2, taking into account the two mirrors, gives

$$\begin{aligned} h_{\text{rad}}(f) &= \frac{2}{l} X(f) \\ &= \frac{1}{mf^2 l} \sqrt{\frac{P_{\text{in}}\hbar}{2\pi^3 c\lambda}}. \end{aligned} \quad (2.25)$$

As can be seen, the radiation pressure noise is proportional $\sqrt{P_{\text{in}}}$. This suggests that the shot noise and the radiation pressure noise cannot be minimized simultaneously and there exists an optimal P_{in} in which the total quantum noise is minimum. Nevertheless, the radiation pressure noise alone can be minimized by increasing the mass m of the mirrors.

The total quantum noise $h_{\text{quantum}}(f)$ is given by a quadrature sum of the shot noise $h_{\text{shot}}(f)$ and the radiation pressure noise $h_{\text{rad}}(f)$, assuming that they are not correlated.

This gives the quantum noise

$$h_{\text{quantum}}(f) = \sqrt{h_{\text{shot}}(f)^2 + h_{\text{rad}}(f)^2}. \quad (2.26)$$

The quantum noise has a minimum when $h_{\text{shot}} = h_{\text{rad}}$. Using Eqn. (2.20) and (2.25), this gives

$$\begin{aligned} \frac{\sqrt{2}}{2l} \sqrt{\frac{\hbar c \lambda}{4\pi P_{\text{in}}}} &= \frac{1}{mf^2 l} \sqrt{\frac{P_{\text{in}} \hbar}{2\pi^3 c \lambda}} \\ P_{\text{in}} &= \frac{1}{2} \pi c \lambda m f^2. \end{aligned} \quad (2.27)$$

Suppose the optimal laser power is chosen to observing gravitational waves at $f = 100$ Hz. In terms of order of magnitude, the speed of light is $c \sim 10^8$, the wavelength, assuming an infrared laser, is $\lambda \sim 10^{-6}$, and a mass $m \sim 10^0$, this gives the optimal laser power of $P_{\text{in}} \sim 10^6$ W. This is still considerably higher than lasers that are used in gravitational-wave detectors today, which is 200 W for the advanced LIGO detector [20]. The quantum noise for such configuration is shown in Fig. 2.2 as a reference.

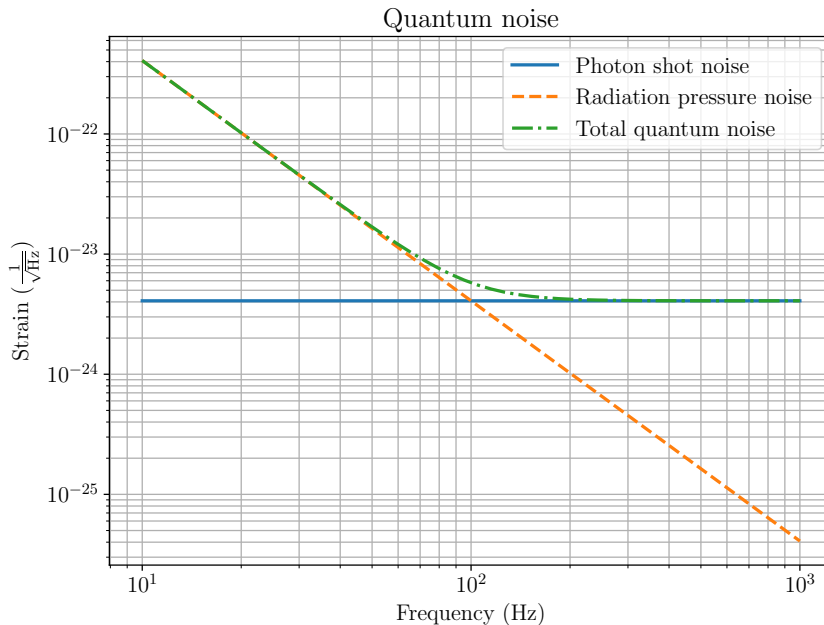


Figure 2.2: Quantum noise of a Michelson interferometer with $\lambda = 1064$ nm, $m = 1$ kg, and $l = 4$ km. Blue solid: Photon shot noise. Orange dashed: Radiation pressure noise. Green dash-dot: Total quantum noise.

2.1.3 Optimal arm length

The remaining design freedom of the Michelson interferometer is the arm length l . Navily speaking, the interferometer arm should be as long as long as possible. This is because

the detected signal is the differential change in arm length Δl , which can be converted to the gravitational wave strain h by

$$2h = \frac{\Delta l}{l}. \quad (2.28)$$

Effectively, the arm length l is an amplification converting the gravitational wave strain to the differential change via $\Delta l = 2lh$. However, this is not true since the response of the Michelson interferometer to the gravitational wave is as simply derived.

Suppose the gravitational wave has a $+$ polarization and the interferometer arms are aligned along the x and y direction. Consider a photon traveling along the x direction, the photon travels along an worldline in spacetime with an invariant interval of

$$\begin{aligned} d\tau^2 &= -g_{\mu\nu}dx_\mu dx_\nu \\ &= -(\eta_{\mu\nu} + h_{\mu\nu})dx_\mu dx_\nu, \end{aligned} \quad (2.29)$$

where $g_{\mu\nu}$ is the metric tensor, $\eta_{\mu\nu}$ is the Minkowski metric with a $(-, +, +, +)$ signature, and $h_{\mu\nu}$ is the metric perturbation due to the gravitational wave. For a photon traveling along the x direction, the invariant interval is zero and this gives

$$dt^2 = \frac{1}{c^2} (1 + h_{11}(t)) dx^2 \quad (2.30)$$

The time that the photon takes to complete a round trip t_{round} along the interferometer arm is then the sum of the forward and the return time given by

$$\int_0^{t_{\text{round}}} dt = \frac{1}{c} \int_0^l \sqrt{1 + h_{11}(t)} dx - \frac{1}{c} \int_l^0 \sqrt{1 + h_{11}(t)} dx. \quad (2.31)$$

With a binomial expansion of $\sqrt{1 + h_{11}(t)} = 1 + h_{11}(t)/2 + O(h_{11}(t)^2)$ and keeping the first order terms of $h_{11}(t)$, this gives the return time

$$t_{\text{round}} = \frac{2l}{c} + \frac{1}{c} \int_0^l h_{11}(t) dx. \quad (2.32)$$

As for a photon traveling along the y direction, the round trip time has a similar express but with the element $h_{22}(t) = -h_{11}(t)$. This gives a difference in return time of

$$\Delta t_{\text{round}} = \frac{2}{c} \int_0^l h_{11}(t) dx, \quad (2.33)$$

which can be approximated by

$$\Delta t_{\text{round}} \approx 2 \int_0^{l/c} h_{11}(t) dt. \quad (2.34)$$

This can be interpreted as a differential change in arm length

$$\Delta l = c\Delta t_{\text{round}} = 2c \int_0^{l/c} h_{11}(t) dt. \quad (2.35)$$

Assuming that the gravitational wave is sinusoidal, i.e. $h_{11}(t) = h \cos(\omega_{\text{gw}}t)$, where ω_{gw} is the angular frequency of the gravitational wave, this gives

$$\begin{aligned} \Delta l &= 2c \int_0^{l/c} h \cos(\omega_{\text{gw}}t) dt \\ &= \frac{2hc}{\omega_{\text{gw}}} \sin\left(\omega_{\text{gw}} \frac{l}{c}\right) \\ &= \frac{2hc}{\omega_{\text{gw}}} \sin\left(2\pi \frac{l}{\lambda_{\text{gw}}}\right), \end{aligned} \quad (2.36)$$

where λ_{gw} is the wavelength of the gravitational wave. Here, it is clear that the maximum differential change in arm length can be found by setting the arm length

$$l = \frac{\lambda_{\text{gw}}}{4}. \quad (2.37)$$

For a gravitational wave at 100 Hz, the optimal arm length is roughly 750 km, which is roughly 1/9 the radius of the earth. Clearly, building a Michelson interferometer that long does not seem practical and alternative strategies need to be sought.

2.2 Optical cavity

2.2.1 Folded Interferometer

As discussed in Sec. 2.1.3, the optimal arm length for a simple Michelson interferometer is impractically long. Instead, consider a not so simple Michelson interferometer where the optical path is folded. This is achieved by placing an mirror at the each input of the interferometer arm next to the beamsplitter. The input mirrors have openings for the beam to enter and exit the interferometer arm. The input mirrors and the end mirrors face the each other such that the laser beam in each arm travels N round trips before recombination. This effectively increased the arm length and downsized the interferometer by N times, making the length scale more reachable. Suppose it is possible to afford an arm length of 4 km, the corresponding round trips needed is roughly $N = 188$. Still, a folded Michelson interferometer can be difficult to achieve due to the diffraction of light [5]. To achieve a folding scheme, one needs to build meter-scale mirrors, which can be difficult to engineer to tight specifications such as reflectivity and surface smoothness [5].

2.2.2 Fabry-Perot cavity

Instead of using a big mirror with an opening, the folded optical paths can superpose on top of each other. This can be achieved by using a partially reflective mirror at the input so the light can leak into the interferometer arm. Such configuration is called a




Figure 2.3: Fabry-Perot cavity.

Fabry-Perot cavity and it is shown in Fig. 2.3. The Fabry-Perot cavity is formed by two partially reflective mirrors with amplitude reflectivity r_1, r_2 and amplitude transmissivity t_1, t_2 . The two mirrors are facing each other and are placed with a separation distance of l . As can be seen, the input laser has an electric field of E_{in} and due to the partial reflection of the mirrors, this gives 3 additional light fields E_{refl} , E_{cav} , and E_{trans} , which corresponds to the electric fields of the reflected light, the light inside the cavity, and the transmitted light, respectively.

The reflected light is a superposition of, the light that gets reflected before entering the cavity, $-r_1 E_{\text{in}}$, the light that transmitted into the cavity but is reflected by the second mirror and transmitted through the first mirror, $t_1 r_2 t_1 E_{\text{in}} e^{-i2kl}$, and so on. The electric field of the reflected light can then be expressed as

$$\begin{aligned} E_{\text{refl}} &= -r_1 E_{\text{in}} + r_2 t_1^2 E_{\text{in}} e^{-i2kl} + r_1 r_2^2 t_1^2 E_{\text{in}} e^{-i4kl} + r_1^2 r_2^3 t_1^2 E_{\text{in}} e^{-i6kl} \dots \\ &= -r_1 E_{\text{in}} + r_2 t_1^2 E_{\text{in}} e^{-i2kl} \sum_{n=0}^{\infty} (r_1 r_2 e^{-i2kl})^n. \end{aligned} \quad (2.38)$$

The last term is just an infinite sum of a diverging geometry series so it can finally be written as

$$\begin{aligned} E_{\text{refl}} &= -r_1 E_{\text{in}} + r_2 t_1^2 E_{\text{in}} e^{-i2kl} \frac{1}{1 - r_1 r_2 e^{-i2kl}} \\ &= \left(-r_1 + \frac{r_2 t_1^2 e^{-i2kl}}{1 - r_1 r_2 e^{-i2kl}} \right) E_{\text{in}}. \end{aligned} \quad (2.39)$$

Here, the amplitude reflectivity of the Fabry-Perot cavity can be written as

$$r_{\text{FP}} = \left(-r_1 + \frac{r_2 t_1^2 e^{-i2kl}}{1 - r_1 r_2 e^{-i2kl}} \right) \quad (2.40)$$

such that $E_{\text{refl}} = r_{\text{FP}} E_{\text{in}}$. Similarly, the electric field of the transmitted electric field can

be written as

$$\begin{aligned}
 E_{\text{trans}} &= t_1 t_2 E_{\text{in}} e^{-ikl} + r_1 r_2 t_1 t_2 E_{\text{in}} e^{-i3kl} + t_1 r_1^2 r_2^2 E_{\text{in}} e^{-i5kl} + \dots \\
 &= t_1 t_2 E_{\text{in}} e^{-ikl} \sum_{n=0}^{\infty} (r_1 r_2 e^{-i2kl})^n \\
 &= \frac{t_1 t_2 e^{-ikl}}{1 - r_1 r_2 e^{-i2kl}} E_{\text{in}}.
 \end{aligned} \tag{2.41}$$

This gives the amplitude transmissivity of the Fabry-Perot cavity

$$t_{\text{FP}} = \frac{t_1 t_2 e^{-ikl}}{1 - r_1 r_2 e^{-i2kl}}. \tag{2.42}$$

Finally, the electric field of the light that is trapped in the cavity (at mirror 1 traveling towards mirror 2) reads

$$\begin{aligned}
 E_{\text{cav}} &= t_1 E_{\text{in}} + t_1 r_1 r_2^2 E_{\text{in}} e^{-i2kl} + \dots \\
 &= \frac{t_1}{1 - r_1 r_2 e^{-i2kl}} E_{\text{in}}.
 \end{aligned} \tag{2.43}$$

Since the reflectivity r_1 and r_2 are less than 1, it is easy to see that the electric field inside the cavity attains a maximum at $2kl = 2n\pi$, where $n = 1, 2, \dots$. This gives the resonance condition

$$2l = n\lambda, \tag{2.44}$$

where λ is the wavelength of the laser. In other words, the Fabry-Perot cavity is said to be in resonance when the round trip distance is an integer multiple of the wavelength. In this case, the light can be seen as being trapped between the mirrors, similar to that in a folded interferometer scheme.

To further examine the equivalence between the folded interferometer and the Fabry-Perot cavity, it would be interesting to compare the phase changes of the light due to a shift Δl between the mirrors. For the Fabry-Perot cavity, consider a special case with $t_2 = 0$ and $r_2 = 1$, i.e. the end mirror is fully reflecting the light. Under the resonance condition, the reflected electric field reads

$$E_{\text{refl}} = \left(-r_1 + \frac{t_1^2 e^{-i2k\Delta l}}{1 - r_1 e^{-i2k\Delta l}} \right) E_{\text{in}}. \tag{2.45}$$

As can be seen, the reflected electric field composes of an term corresponding to the light field that has never entered the cavity, $-r_1 E_{\text{in}}$. The other term can be seen as the light that enters the cavity and escaped, it reads

$$E_{\text{esc}} = \frac{t_1^2 e^{-i2k\Delta l}}{1 - r_1 e^{-i2k\Delta l}} E_{\text{in}}. \tag{2.46}$$

Suppose the light escaping from the cavity can be written in the form $Ae^{-i\phi}$, the fractional change of phase $\Delta\Phi$ due to a change in cavity length Δl is approximated, to the first order, as

$$\left(\frac{\Delta\phi}{\Delta l}\right)_{\text{FP}} = \frac{2k}{1-r_1}. \quad (2.47)$$

In comparison, for a folded cavity, the phase of the beam picks on a phase of $2k\Delta l$ per round trip. This gives the a fractional phase change of

$$\left(\frac{\Delta\phi}{\Delta l}\right)_{\text{Folded}} = 2kN, \quad (2.48)$$

for the folded cavity. By comparing Eqn. (2.47) and (2.48), it is easy to see that the equivalent number of round trips in a Fabry-Perot cavity can be written as

$$N_{\text{FP},\phi} = \frac{1}{1-r_1}. \quad (2.49)$$

As a verification, for $r_1 = 0$, i.e. the first mirror does not exist, the number of equivalent round trips is 1, which reduces to the Michelson interferometer case. It is worth mentioning that $N_{\text{FP},\phi}$ denotes the equivalent round trips by a phase consideration. There exists another equivalent number of round trips that is derived from the storage time of the cavity, i.e. the time that takes the light to escape as the input beam is shut. For details, refer to the folded interferometer arm chapter in Ref. [5].

With the folded scheme made possible with Fabry-Perot cavities, the photon shot noise and radiation pressure noise can be re-evaluated. A folded interferometer is equivalent to a simple Michelson interferometer with an arm length of Nl . The photon shot noise can simply be written by replacing l by Nl in Eqn. (2.20), giving a shot noise that is N times lower, i.e.

$$h_{\text{shot}}(f) = \frac{\sqrt{2}}{2Nl} \sqrt{\frac{\hbar c \lambda}{4\pi P_{\text{in}}}}, \quad (2.50)$$

assuming an operation at the dark fringe. As for the radiation pressure noise, it is enhanced in two ways. Firstly, the radiation pressure introduces a fluctuation test mass displacement, which induces a fluctuation in arm length that is $2N$ times that of an unfolded arm. Secondly, since the light is encounter each mirror N times, the force is multiplied by a factor of N . With l replaced by Nl in Eqn. (2.25), this gives the radiation pressure noise for a Fabry-Perot interferometer of

$$h_{\text{rad}}(f) = \frac{N}{m f^2 l} \sqrt{\frac{2\hbar P_{\text{in}}}{\pi^3 c \lambda}}. \quad (2.51)$$

The corresponding optimal input laser power is

$$P_{\text{in}} = \frac{1}{4} \frac{\pi c m \lambda f^2}{N^2}. \quad (2.52)$$

With the same parameters used in Sec. 2.1, assuming a round trip number of $N \sim 100$, this gives an optimal power in the 10^4 W scale and for $f = 10$, the optimal power is in the 10^2 W scale, which is attainable. As can be seen, with the Fabry-Perot cavity, it solves not just the optimal arm length problem but also the optimal power problem.

2.2.3 Power and signal recycling cavities

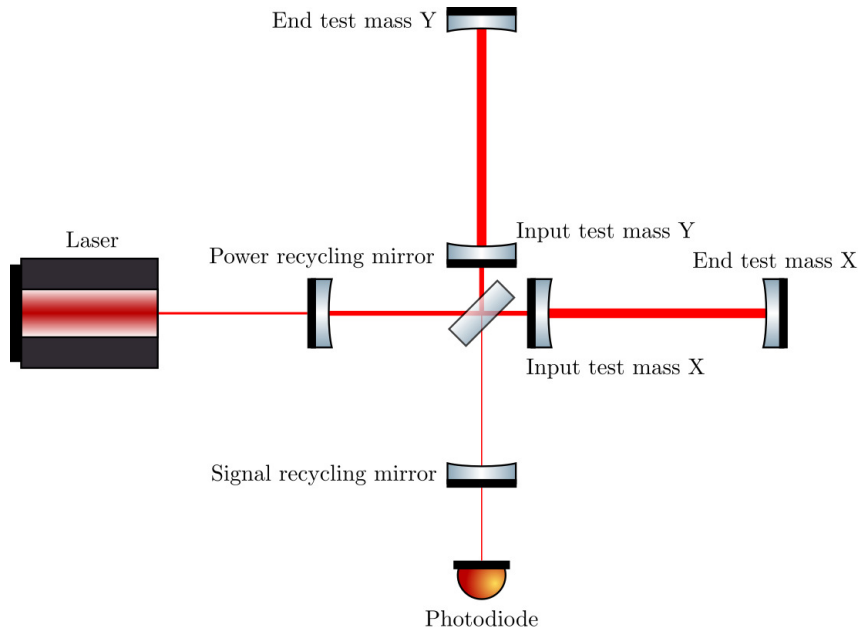


Figure 2.4: Dual-recycling Fabry-Perot Michelson Interferometer. (Not to scale)

Recall that the gravitational-wave detector is a null instrument, meaning that the light power detected by the photodiode is almost zero at all times. By conservation of energy, this means that the almost all light that goes into the interferometer eventually exits the interferometer via the input, as if the interferometer is reflecting light back to the input. While the reflected light at the input would have been eventually dissipated, this is a waste of energy. To make better utilization of the laser, the exiting light can be reflected back into the interferometer, which effectively increase the input power of the interferometer. This is achieved by placing a partially reflective mirror, which is known as the power recycling mirror, between the input laser and the beamsplitter, as shown in Fig. 2.4. The power recycling mirror forms a optical cavity with the interferometer, enhancing the input power of the interferometer by the number of round trips N . With the power recycling scheme, the shot noise of the detector is reduced by a factor of \sqrt{N} .

Similarly, a partially reflective mirror can be placed between the photodiode and the beamsplitter. The mirror in this configuration is called the signal recycling mirror as shown in Fig. 2.4. Instead of the laser frequency f_l , the signal recycling cavity can be set up to resonant at a different frequency $f_l \pm f_{\text{sig}}$, where f_{sig} is the signal frequency [5]. Since the gravitational wave signal is encoded on the sideband through phase modulation of the laser frequency at frequency f_{sig} . By reflecting the signal sideband to the interferometer allows it to be enhanced.

2.3 Noises in gravitational-wave detectors

Detecting gravitational waves means a search of extremely tiny signals. A detector looking for such small signals is inevitably contaminated by noise from various sources. While these noises degrade the quality of the gravitational-wave signals, suppressing these noises and hence achieving better sensitivity is one of the everlasting goals in experimental gravitational-wave physics. There are several noise sources in a gravitational-wave detector, including quantum noise, Newtonian noise, thermal noise, laser noise, and seismic noise, to list a few. Here, a few important noise sources are short-listed for a brief introduction.

2.3.1 Quantum noise

Quantum noise refers to the detection noise that related to the quantum nature of the laser used in the detector. The two fundamental components of the quantum noise, photon shot noise and radiation pressure noise, have already been briefly discussed in Sec. 2.1.1, 2.1.2, and 2.2.2. And, this section only serves as a brief reiteration. The photon shot noise refers to the noise resulting from counting the number of photons, which is a random variable that follows the Poisson distribution. The shot noise has is a white noise and it scales with the input laser power as $\sqrt{1/P_{\text{in}}}$. The radiation pressure noise is resulting from the fluctuating number of photons exerting force on the test masses of the detector. This leads to a arm length fluctuation and the radiation pressure noise scales as $\sqrt{P_{\text{in}}}$. Comparing the laser power scaling between the photon shot noise and radiation pressure noise, it is clear that the two are not simultaneously minimizable by a increasing or decreasing of the laser power. Instead, there exists a optimal trade-off between the shot noise and the radiation pressure noise via an optimal choice of laser power. With the optimal laser power, the resulting quantum noise is called the standard quantum limit, which is a fundamental limit of a gravitational-wave detector.

2.3.2 Thermal noise

The mirrors of the interferometer and the structures that are holding them in place are thermally connected to the surrounding environment and are allowed to exchange energy in the form of heat, which translates to random motion. The random flow of energy can cause the mirror surface to fluctuate and eventually causing the optical path to fluctuate with it. Internal friction from materials can also be translate to thermal noise and this can be estimated from the fluctuation-dissipation theorem [5]. The thermal noise of such systems scales with $\sqrt{k_B T}$, where k_B is the Boltzmann constant and T is the temperature and has a frequency scaling of $f^{5/2}$. This calls for the use of materials with low mechanical loss, or in other words, high quality factor, Q . For this reason, fused silica [21] was chosen for LIGO and Virgo as the material for the test mass substrate and for suspension wires due to its high quality factor ($\sim 10^7$) at room temperature. As for KAGRA, sapphire was chosen instead due to its high thermal conductivity and high Q value at cryogenic temperatures [22].

2.3.3 Seismic noise

Seismic noise is due to the vibration of the terrestrial environment. As opposed to the quantum noise and thermal noise there are no fundamental constants, such as the Planck constant \hbar and Boltzmann constant k_B , that sets the scale of the seismic noise. The seismic noise is due to natural phenomenon such as ocean waves and atmospheric activities. Human activities also contributes to the seismic noise in a local scale. The seismic noise is not a static noise and the amplitude can vary by orders of magnitude from time to time. Also, the seismic noise at different locations may also vary due to different levels of activity and geological reasons. Therefore, the seismic noise can be hard to derive from first principles. Instead, the seismic noise is characterized from worldwide observations of seismic noise. The resulted models are the well-known Peterson low noise model (NLNM) and high noise model (NHNM) [23]. They are obtained by combining the seismic data measured from all around the world and modeling the upper and lower bounds.

Fig. 2.5 shows the Peterson noise models and the KAGRA mean seismic noise. There are several features in the seismic noise spectrum that are worth mentioning. First of all, at 10 Hz, the KAGRA mean seismic noise reads $\sim 10^{-11}$ m/ $\sqrt{\text{Hz}}$. For an interferometer baseline in the order of 10^3 m, the seismic noise correspond to a strain noise in the order of $\sim 10^{-14}$ / $\sqrt{\text{Hz}}$. This calls for a required suppression of at least 10^7 times if gravitational waves in the order of 10^{-21} / $\sqrt{\text{Hz}}$ were to be observed.

At frequencies around 0.2 Hz, there is a peak in the seismic noise spectrum that is due to seismic activity known as the secondary microseism. The primary microseism is located at around 0.06 Hz and is also observable in both noise models and the KAGRA seismic noise in Fig. 2.5. The root mean square value of the KAGRA ground motion,




Figure 2.5: Seismic noise. Blue solid: Peterson low noise model (NLNM). Orange dashed: Peterson high noise model (NHNM). Green dash-dot: KAGRA mean seismic noise along the Y-arm of the interferometer.

integrated from high frequency down to 0.1 Hz, is in the order of 10^{-7} m, and in some cases, in the order of 10^{-6} m [24]. Recall that the interferometer needs to be operated at the dark fringe, meaning that the differential displacement of the interferometer arms need to stay well within a fraction of the wavelength, which is in the order of 10^{-6} m for an infrared laser. This means that the suppression of seismic noise is needed not only at the observation band of the gravitational-wave detector, but also extends to lower frequency, or else the interferometer would not be able to operate at all. Therefore, the suppression of seismic noise is very important. The following chapters shall continue on the discussion of this interesting topic.

Chapter 3

Vibration Isolation

Seismic noise can affect a gravitational-wave detector in two ways. First of all, the seismic noise is an important source of displacement noise that can cause the test masses of the gravitational-wave detector to vibrate at the observation band (10 Hz – few kHz). With an interferometer baseline in the order of 10^3 m, the displacement levels of the test masses must be at least below 10^{-18} m/ $\sqrt{\text{Hz}}$ for detecting gravitational wave strains in the order of 10^{-21} . At 10 Hz, the seismic noise at KAGRA is around 10^{-11} m/ $\sqrt{\text{Hz}}$. This would require the seismic noise transmitted to the test masses to be attenuated by at least 10^7 times above 10 Hz. In gravitational-wave detectors, this is achieved by suspending the test masses as multiple pendulums, which passively isolate the test masses from the seismic noise.

At lower frequencies, seismic isolation is as important. The seismic noise spectrum is typically highest at frequency between 0.1 Hz - 0.5 Hz. This peak is known as the microseismic peak and is caused by the ocean waves. The magnitude of this peak depends on the weather. At times, it could cause the ground motion to attain a displacement level higher than 1 μm , which is comparable to the wavelength of the main laser used at KAGRA (1064 nm). Moreover, the resonances of the multiple-pendulum suspensions amplify the seismic noise at lower frequencies, causing the test masses to move excessively. This can result in a misalignment of the optics and eventually causing the interferometer to dysfunction. Therefore, the displacement levels of the test masses also need to be suppressed at lower frequencies. This kind of vibration isolation is achieved by control systems using active elements such as sensors and actuators. The basic principles of passive isolation and active isolation are discussed in Sec. 3.1 and 3.2, respectively.

3.1 Passive isolation

3.1.1 Oscillators: spring-mass and pendulums

The fundamental principle of passive vibration isolation is based on the dynamics of the harmonic oscillator, such as a spring-mass system and a pendulum. Consider a spring-




Figure 3.1: Spring mass system.

mass system as shown in Fig. 3.1. In the figure, a mass m is attached to a suspension point via a spring k . The suspension point is free to move and it has a displacement $x(t)$. The displacement of the mass $x_m(t)$, assuming an equilibrium position at $x_m(t) - x(t) = 0$, is governed by the differential equation

$$m \frac{d^2 x_m(t)}{dt^2} = -k [x_m(t) - x(t)] . \quad (3.1)$$

Now, suppose the suspension point is oscillating at an angular frequency ω , it can be described by

$$x(t) = A e^{i\omega t} , \quad (3.2)$$

where A is the amplitude of the oscillation, i is the imaginary number, ω is the angular frequency, and t is time. Let us construct a trial solution

$$x_m(t) = A_m e^{i(\omega t + \phi)} , \quad (3.3)$$

where A_m is the amplitude of the oscillation of the mass and ϕ is the relative phase between $x(t)$ and $x_m(t)$. Substituting the trial solution (3.3) into the governing equation

(3.1) gives

$$\begin{aligned}
 -m\omega^2 A_m e^{i(\omega t + \phi)} &= -k [A_m e^{i(\omega t + \phi)} - A e^{i\omega t}] \\
 -m\omega^2 A_m e^{i\phi} &= -k [A_m e^{i\phi} - A] \\
 A_m e^{i\phi} &= \frac{\frac{k}{m}}{\frac{k}{m} - \omega^2} A
 \end{aligned} \tag{3.4}$$

It is easy to see that when $\omega^2 < k/m$, the phase is $\phi = 0$, and when $\omega^2 > k/m$, the phase is π (or an even integer multiple of π). However, at $\omega^2 = k/m$, the phase is undefined and the amplitude A_m explodes¹. For this reason, $\omega_n = \sqrt{k/m}$ is called the resonance frequency of the spring-mass system, where the motion of the suspension point is being amplified. At $\omega \ll \omega_n$, Eqn. (3.4) reads $A_m \approx A$, so the mass follows almost exactly the motion of the suspension point, i.e. the motion is not attenuated. However, at $\omega \gg \omega_n$, the amplitude reads $A_m \approx (\omega_n/\omega)^2 A$, i.e. the suspension point displacement is attenuated by a factor of $(\omega_n/\omega)^2$. For example, suppose the oscillation frequency is 10 times greater than then resonance frequency, the attenuation factor is approximately 100 times. As can be seen, by suspending the mass with a spring, the displacement level of the mass can be attenuated at frequencies above the resonance and this is basically how passive isolation works.

To put things into perspective, consider a case where test masses in a gravitational-wave detector are suspended by the spring-mass mechanism. To achieve the aforementioned seismic noise attenuation of 10^7 at 10 Hz, the resonance frequency of the spring-mass system needs to be as low as ~ 0.00316 Hz (a period of ~ 5 min). Suppose the spring-mass mechanism is achieved by suspending the test mass like a pendulum, the resonance frequency can be computed as $\omega_n = \sqrt{g/l}$, where g is the acceleration due to gravity and l is the length of the pendulum. Here, small angles are assumed, which would be valid given the nature of small displacement in gravitational-wave detectors. The required resonance frequency then converts to a pendulum length of roughly 24 841 m, which is just 30 times taller than the tallest building on Earth. Clearly, it is not practical to build a pendulum that tall and alternative designs must be sought. In reality, passive isolation in gravitational-wave detectors rely on two other mechanisms, anti-springs and stacked spring-mass systems. The former adds negative stiffness to a spring-mass system, which makes low resonance frequency achievable in a relatively compact system. And, the latter allows a higher-order attenuation, which makes high seismic noise attenuation achievable with multiple spring-mass systems with resonance frequencies that are not unreasonably low.


¹Of course, the amplitude of the mass cannot be amplified indefinitely in reality. This is due to energy loss via viscous damping.

3.1.2 Anti-springs

In KAGRA, there are two anti-spring mechanisms, geometric anti-springs (GAS) [25] and inverted pendulum. GAS filters are mainly used for vibration isolation along the vertical direction². As for the inverted pendulum, it is used as the first seismic isolation stage to isolate seismic noise along the horizontal direction. The inverted pendulum can be tuned to a resonance frequency lower than 0.1 Hz and hence KAGRA vibration isolation systems with an inverted pendulum is also capable of attenuating the secondary microseismic peak.

Geometric anti-spring

(a) Geometric anti-spring



(b) Geometric anti-spring (equivalent model)




Figure 3.2: (a): Conceptual 2D geometric anti-spring. (b): An equivalent model.

GAS filters, providing vibration isolation in the vertical direction, are essentially horizontally compressed spring blades. Fig. 3.2 subplot (a) shows a 2D representation of the GAS filter. The spring blades are metal cantilevers that are originally flat. The two blades are compressed horizontally towards the mass. This causes the blades bend upward, giving the signature fishing rod shape of a GAS filter. The compressed spring blades has two effects on the suspended mass. Firstly, they act as a vertical spring hanging the mass, like a fishing rod. This alone provides some vertical isolation via the spring-mass mechanism. Secondly, they provide a horizontal compression on the mass at the equilibrium position. And, when the mass deviates from the equilibrium position, the two horizontal forces has a component in the same direction in which the mass deviates, which is opposite to the restoring force provided by the first effect. This effectively reduces the overall restoring force in the vertical direction, achieving the anti-spring effect.

The GAS filter can be modeled by a three-spring system as shown in subplot (b) in Fig. 3.2. As can be seen, the suspended mass is attached to three springs, two horizontally and one vertically. The horizontal springs have a spring constant of k_x and are compressed

²The reason for seismic isolation along the vertical direction is due to the fact that the Earth is round while the interferometer arm is straight. With a interferometer baseline of 3 km and an Earth radius of ~ 6000 km, the vertical vibration (in the radial direction) translates to a coupling of $3/6000 = 1/300$ in the direction along the interferometer arm.

by a length of x_c . The vertical spring have a spring constant of k_y . Again, the suspension point and the suspended mass has a vertical displacement of $y(t)$ and $y_m(t)$, respectively. Suppose the horizontal springs have a nominal length of x_0 , as the vertical spring is being stretched by an amount of $y_m(t) - y(t)$, the length of the horizontal spring becomes

$$x_{\text{spring}}(t) = \sqrt{[y_m(t) - y(t)]^2 + [x_0 - x_c]^2} \quad (3.5)$$

The vertical force that each horizontal spring that acts on the suspended mass is

$$F_h(t) = k_x [x_0 - x_{\text{spring}}(t)] \sin \theta, \quad (3.6)$$

where

$$\theta = \sin^{-1} \left(\frac{y_m(t) - y(t)}{x_{\text{spring}}(t)} \right). \quad (3.7)$$

This gives the vertical force due to a horizontal spring

$$F_h(t) = k_x \left[\frac{x_0}{\sqrt{[y_m(t) - y(t)]^2 + [x_0 - x_c]^2}} - 1 \right] [y_m(t) - y(t)]. \quad (3.8)$$

Now, consider a binomial expansion of the first term:

$$\begin{aligned} \frac{1}{\sqrt{[y_m(t) - y(t)]^2 + [x_0 - x_c]^2}} &= \frac{1}{(x_0 - x_c) \sqrt{1 + \left(\frac{y_m(t) - y(t)}{x_0 - x_c} \right)^2}} \\ &= \frac{1}{x_0 - x_c} \left[1 + O \left(\left[\frac{y_m(t) - y(t)}{x_0 - x_c} \right]^2 \right) \right]. \end{aligned} \quad (3.9)$$

Linearizing the vertical force to the first order of $y_m(t) - y(t)$ gives

$$F_h(t) = k_x \left(\frac{x_c}{x_0 - x_c} \right) [y_m(t) - y(t)]. \quad (3.10)$$

With the vertical force due to the compressed horizontal springs derived, the equation of motion of the suspended mass now reads

$$\begin{aligned} m \frac{d^2 y_m(t)}{dt^2} &= -k_y [y_m(t) - y(t)] + 2F_h(t) \\ &= - \left(k_y - 2k_x \frac{x_c}{x_0 - x_c} \right) [y_m(t) - y(t)], \end{aligned} \quad (3.11)$$

where $k_y - 2k_x x_c / (x_0 - x_c)$ is the effective spring constant along the vertical direction. While $x_c < x_0$, it can be seen that the compression of the horizontal spring is indeed contributing a negative stiffness and hence achieving the anti-spring effect.

Inverted pendulum




Figure 3.3: An inverted pendulum.

The inverted pendulum is used to isolate a platform from horizontal vibration. To illustrate its anti-spring effect, consider an inverted pendulum as shown in Fig. 3.3. The mass m is suspended from the suspension point via a massless inverted pendulum leg with length l . The leg is connected to the suspension point via a flexible joint with a bending stiffness of k_θ . The equation of motion of the suspended mass reads

$$m \frac{d^2 x_m(t)}{dt^2} = -\frac{k_\theta}{l} \theta + mg \sin \theta, \quad (3.12)$$

where θ is the angle between the inverted pendulum and the vertical. With the small angle approximation, the angle θ can be written as

$$\theta = \frac{x_m(t) - x(t)}{l}. \quad (3.13)$$

Substituting θ into the equation of motion gives

$$m \frac{d^2 x_m(t)}{dt^2} = -\frac{k_\theta}{l^2} [x_m(t) - x(t)] + mg \sin \left(\frac{x_m(t) - x(t)}{l} \right). \quad (3.14)$$

Again, linearization to the first order of $x_m(t) - x(t)$ gives

$$m \frac{d^2 x_m(t)}{dt^2} = -\left(\frac{k_\theta}{l^2} - \frac{mg}{l} \right) [x_m(t) - x(t)], \quad (3.15)$$

where $k_\theta/l^2 - mg/l$ is the effective stiffness. To enhance the anti-spring effect, one can simply increase the mass m and this is exactly how the inverted pendulums at KAGRA were tuned to have resonance frequency below 0.1 Hz.

3.1.3 Cascaded spring-mass systems

In Sec. 3.1.1, it is shown that a simple spring-mass system can be used to achieve an displacement attenuation of $(\omega_n/\omega)^2$, where ω_n is the angular resonance frequency of the spring-mass system and ω is the angular frequency of the vibration of the suspension point. Here, the frequency ω is arbitrary, and in fact, the displacement of the suspension point can contain a superposition of vibrations at various frequencies and the attenuation relation is still valid for all these independent vibrations. In the case where the suspension point is a stochastic process, e.g. the seismic noise, its frequency content above the resonance frequency of the spring-mass system attenuated by a frequency-dependent attenuation of $(\omega_n/\omega)^2$ and this is known as low-pass filtering. And, the exponent of the attenuation is 2, therefore the simple spring-mass system is referred as a second-order low-pass filter. With a second-order filtering, it was shown that the resonance frequency has to be unreasonably low in order to observe gravitational wave. While it was shown that the anti-spring effect can be manipulated to reduce the resonance frequency of a spring mass system, resonance frequencies below 10 mHz is still unrealistic in practice due to non-linear effects leading to bistability or even instability. For example, the lowest resonance frequency of the inverted pendulums at KAGRA is around ~ 50 mHz, which is still one order of magnitude higher than that of the required. To achieve even higher level seismic attenuation at higher frequencies, one might consider to increase the order of filtering instead and this can be achieved by cascading spring-mass systems.




Figure 3.4: Double spring-mass system.

To see how stacking spring-mass systems can increase the order of low-pass filtering, consider a double spring-mass system as shown in Fig. 3.4. Again, here the displacement of the suspension point $x(t)$ is assumed to be

$$x(t) = Ae^{i\omega t} \quad (3.16)$$

and the displacements of the first and second masses are

$$x_1(t) = A_1 e^{i(\omega t + \phi_1)}, \quad (3.17)$$

and

$$x_2(t) = A_2 e^{i(\omega t + \phi_2)}, \quad (3.18)$$

respectively, where A , A_1 , and A_2 are the oscillation amplitudes and ϕ_1 and ϕ_2 are the relative phases. The governing equations for the two masses are

$$m_1 \frac{d^2 x_1(t)}{dt^2} = -k_1 [x_1(t) - x(t)] - k_2 [x_1(t) - x_2(t)] \quad (3.19)$$

and

$$m_2 \frac{d^2 x_2(t)}{dt^2} = -k_2 [x_2(t) - x_1(t)]. \quad (3.20)$$

Following a similar analysis in Sec. 3.1.1 yields

$$A_1 e^{i\phi_1} = \frac{\frac{k_1}{m_1} \left(\frac{k_2}{m_2} - \omega^2 \right)}{\omega^4 - \left(\frac{k_1+k_2}{m_1} + \frac{k_2}{m_2} \right) \omega^2 + \frac{k_1 k_2}{m_1 m_2}} A \quad (3.21)$$

and

$$\begin{aligned} A_2 e^{i\phi_2} &= \frac{\frac{k_2}{m_2}}{\frac{k_2}{m_2} - \omega^2} A_1 e^{i\phi_1} \\ &= \frac{\frac{k_1 k_2}{m_1 m_2}}{\omega^4 - \left(\frac{k_1+k_2}{m_1} + \frac{k_2}{m_2} \right) \omega^2 + \frac{k_1 k_2}{m_1 m_2}} A \\ &= \frac{\frac{k_1 k_2}{m_1 m_2}}{(\omega^2 - \omega_1^2)(\omega^2 - \omega_2^2)} A, \end{aligned} \quad (3.22)$$

where ω_1 and ω_2 are two angular resonance frequencies such that

$$\omega_1^2, \omega_2^2 = \frac{\left(\frac{k_1+k_2}{m_1} + \frac{k_2}{m_2} \right) \pm \sqrt{\left(\frac{k_1+k_2}{m_1} + \frac{k_2}{m_2} \right)^2 - 4 \frac{k_1 k_2}{m_1 m_2}}}{2}, \quad (3.23)$$

which can be easily shown as positive values. As can be seen from Eqn. (3.22), as $\omega \gg \omega_1$ and $\omega \gg \omega_2$, the amplitude A_2 can be approximated as

$$A_2 \approx \frac{1}{\omega^4} \frac{k_1 k_2}{m_1 m_2} A, \quad (3.24)$$

which is clearly a forth-order attenuation of the displacement of the suspension point. However, the attenuation at the first stage, i.e. mass 1, remains to be second order.

Increasing the number of spring-mass stages can effectively increase the order of low-pass filtering, achieving higher-order attenuation at high frequencies. However, the result for systems with more than 2 stages can be more difficult to compute algebraically. To

layout the steps, one needs to express the government equations in a matrix form

$$-\omega^2 \mathbf{M} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} + \mathbf{K} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} = \begin{bmatrix} k_1 x \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (3.25)$$

where \mathbf{M} is the mass matrix, \mathbf{K} is the stiffness matrix, and here $\ddot{x}_j = \omega^2 x_j$, $j = 1, 2, \dots, N$ is assumed. And then, compute $(-\omega^2 \mathbf{M} + \mathbf{K})^{-1}$ to obtain the attenuation at each stage. While \mathbf{M} is a diagonal matrix composes of the values of the masses, it is easy to show that the denominator of $(-\omega^2 \mathbf{M} + \mathbf{K})^{-1}$, i.e. $\det(-\omega^2 \mathbf{M} + \mathbf{K})$, is a $2N^{\text{th}}$ -order polynomial of ω . It follows that that the attenuation at the last stage of a N -stage spring-mass system is proportional to $1/\omega^{2N}$. Fig. 3.5 shows the ratio between the displacement amplitude of




Figure 3.5: Suppression ratio between of spring-mass systems. Blue solid: Simple spring-mass system. Orange dashed: Double spring-mass system. Green dash-dot: Triple spring-mass systems. Red dotted: Quadruple spring-mass systems.

the suspended mass and the suspension point of simple spring-mass system and higher-order spring-mass systems (up to $N = 4$). Here, the spring constants and masses are set to unity for illustration purpose. As can be seen, higher-order systems attain higher suppression ratio at frequencies above the resonances. In fact, a suppression ratio of 10^{-8} can be achieved by a quadruple system (red dotted) with resonances around 1 Hz. And, at last, with a combination of the anti-spring effect, i.e. lowering the spring constants, and increasing the number of stages of spring-mass systems (or pendulums), the required attenuation for observing gravitational waves is easily achievable with a relatively practical

form factor of systems. These systems are referred to suspensions or vibration isolation systems and some real versions of suspensions are shown in Chapter 4.

3.2 Active isolation

This idea of active isolation is to, instead of relying of mechanical components such as springs or pendulums, use an actuator to apply a force on the suspended object to reject the perturbed motion. The required force can be derived from the sensed displacement of the suspension point or the suspension body. The former is known as feedforward control while the latter is known as feedback control. The two active isolation schemes can be used simultaneously to achieve disturbance rejection at different frequency bandwidths. Notably, the two schemes are utilized in LIGO for active seismic isolation [26] while Virgo [27, 28] and KAGRA [29, 28, 30, 31] only utilize feedback control. In this section, the feedforward and feedback control schemes are discussed in Sec. 3.2.2 and 3.2.3, respectively. In Sec. 3.2.1, some prerequisite concepts such as transfer functions and block diagrams are introduced as an interlude.

3.2.1 Interlude: Transfer functions

Before going into the details of active isolation, it is worth introducing the concept of transfer functions, which would simplify the discussion of active isolation. The transfer function represents a linear time-invariant (LTI) system with one input and one output. Mathematically, it is the ratio between the output and the input variables in the Laplace domain, i.e. the amplification/attenuation. To exemplify, consider a single spring-mass system discussed in Sec. 3.1.1. Here, the displacement of the suspension point $x(t)$ is an input and the displacement of the mass $x_m(t)$ is the output. Taking the Laplace transform of the equation of motion (3.1) gives

$$\begin{aligned} \mathcal{L}\left\{m\frac{d^2x_m(t)}{dt^2} + kx_m(t)\right\} &= \mathcal{L}\{kx(t)\} \\ (ms^2 + k)X_m(s) &= kX(s) \\ \frac{X_m(s)}{X(s)} &= \frac{\omega_n^2}{s^2 + \omega_n^2}, \end{aligned} \tag{3.26}$$

where $\mathcal{L}\{\cdot\}$ denotes the Laplace transform, s is the Laplace variable and is a complex number, $X(s)$ and $X_m(s)$ are the Laplace transform of $x(s)$ and $x_m(t)$, respectively, and $\omega_n \equiv \sqrt{k/m}$ is the angular resonance frequency of the spring-mass system. Here, as can be seen, by evaluating the Laplace variable along the positive imaginary axis, i.e. $s = j\omega$, gives the frequency-dependent attenuation relation $\omega_n^2/(\omega_n^2 - \omega^2)$ as shown in Eqn. (3.4). This is called the frequency response of a transfer function. The frequency

response is a complex-valued frequency series and therefore cannot be plotted simply in one plot. Instead, the frequency response of a transfer function is shown as a magnitude plot and a phase plot, where the magnitude represents the amplification/attenuation (and sometimes referred as “gain”) and the phase represents the relative phase between the input and output signals. The combination of the two plots is called a Bode plot and is often used to characterize frequency-dependent relationships between signals.

The introduction of the transfer function allows one to represent dynamic processes with block diagrams. In a block diagram, each arrow represents a process variable, such as the displacements, and each block represents a system corresponding to a transfer function. Fig. 3.6 shows a block diagram representing the aforementioned spring-mass




Figure 3.6: The block diagram representation of a spring-mass system.

system. The diagram represents an equation

$$X_m(s) = P_x(s)X(s), \quad (3.27)$$

where

$$P_x(s) \equiv \frac{\omega_n^2}{s^2 + \omega_n^2} \quad (3.28)$$

is the transfer function between the suspension point displacement and the mass displacement and is sometimes called the “plant”. As can be seen, the block diagrams conveniently represents the dynamics of the spring-mass system. From here, more complicated systems, such as the feedforward control and feedback control systems, can be represented by interconnecting the process variables via junctions.

3.2.2 Feedforward control

Suppose there is an actuator that can be used to apply a force to the mass of the spring-mass system shown in Fig. 3.1. The force then can be used to cancel, or at least attenuate, the motion of the mass due to the motion of the suspension point. With an external force, the equation of motion of the suspended mass reads

$$m \frac{d^2 x_m(t)}{dt^2} + k x_m(t) = k x(t) + f(t), \quad (3.29)$$

where $f(t)$ is the external force applied on the suspended mass. It is easy to see that when the force is set to $f(t) = -kx(t)$, then the suspended mass is completely decoupled from

the motion of the suspension point. And, this can be achieved by having a measurement of $x(t)$ and then converting that sensed signal to an actuation signal via a multiplicative factor $-k$. Such control method is known as feedforward control³ and the corresponding spring-mass system is shown in Fig. 3.7. As shown in subplot (a), a sensor is placed at the




Figure 3.7: (a): Feedforward control of a spring-mass system. (b): Equivalent block diagram of the feedforward control scheme.

suspension point measuring the displacement $x(t)$. The signal is sent to a feedforward controller, which is represented by a transfer function $K_{ff}(s)$. The controller then outputs a force $f(t)$ that acts on the suspended mass.

Fig. 3.7 subplot (b) shows an equivalent block diagram representing the feedforward control of a spring-mass system. Without loss of generality, the input-output relationship can be easily derived as

$$X_m(s) = [P_x(s) + K_{ff}(s)P(s)] X(s), \quad (3.30)$$

where $P(s)$ is the transfer function from the actuation force $F(s)$ to the displacement of the suspended mass $X_m(s)$, and is given by

$$P(s) = \frac{1}{s^2 + \frac{k}{m}}. \quad (3.31)$$

Here, the subscript is omitted due to its frequent use throughout the thesis. In Eqn. 3.30, the feedforward controller $K_{ff}(s)$ is something that needs to be designed. Again, it is obvious when the feedforward controller is set to $K_{ff}(s) = -P_x(s)/P(s) = -k$, the attenuation factor then becomes 0. However, this only serves as an ideal case for illustrating the idea of feedforward control. In reality, the measurement of $x(t)$ is not perfect. The measurement contains noise and could ruin the passive isolation performance when it is not properly filtered. Therefore, in practice, the feedforward controller needs to be shaped to pass the signal at frequencies where signal-to-noise ratio is high and rolled-off otherwise.

³Note that the signal path points forward and hence the name “feedforward”.

It is worth to mention that the feedforward scheme is not as important for Virgo and KAGRA suspensions due to the low resonance frequencies of the suspensions. The Virgo and KAGRA suspensions have resonances lower than the secondary microseism so the microseismic peak is suppressed adequately. However, the resonances at lower frequencies become the dominating motion of the suspensions and they require also suppression. With noiseless seismometers, the resonances can in principle suppressed by feedforward control. But in reality, seismic noise at such low frequencies can only be measured by seismometers with very low signal-to-noise ratio (as shown in much later chapters). This means that the feedforward scheme would rather inject noise, than providing cancellation force at low frequencies. These motions are suppressed by feedback control instead. LIGO seismic isolation systems do not have inverted pendulums so secondary microseism is not attenuated passively. Therefore, it is reasonable for LIGO to use feedforward control to attenuate the seismic noise at lower frequencies.

3.2.3 Feedback control

For feedforward control, the sensing element is placed on the suspension point. When the actuation force is derived from a sensing element placed on the suspended object itself, it becomes a feedback control scheme shown in Fig. 3.8. In subplot (a), the displacement

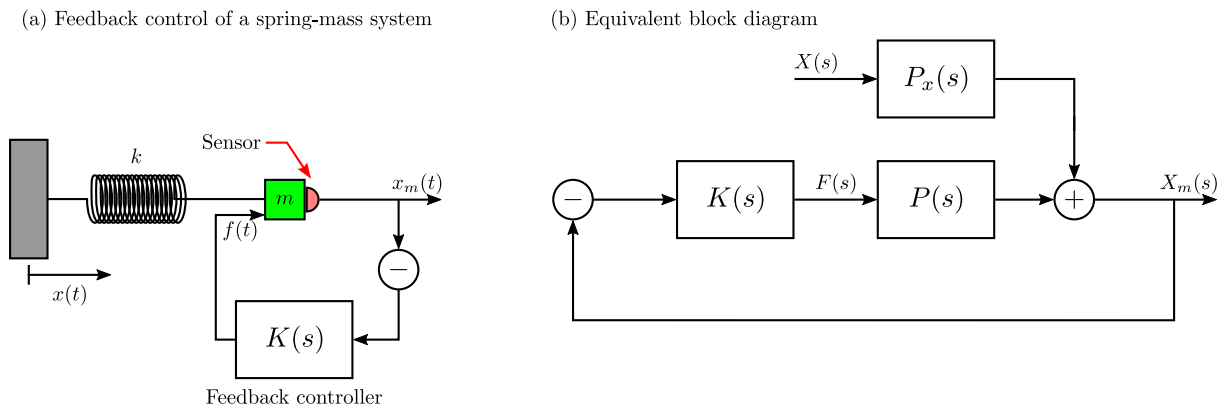


Figure 3.8: (a): Feedback control of a spring-mass system. (b): Equivalent block diagram of the feedback control scheme.

of the suspended mass $x_m(t)$ is measured by a sensor. The signal is negated and is sent to a feedback controller $K(s)$, which then converts the measured signal to a feedback force. The negation of the signal indicates a negative feedback system and this exists for conventional reasons related to stability of the feedback. Again, the subscript for the feedback controller is omitted because of its frequent use in later sections in the thesis.

Fig. 3.8 subplot (b) shows the block diagram representing the feedback control of the

spring-mass system⁴. The displacement of the suspended mass can be derived as

$$\begin{aligned} X_m(s) &= P_x(s)X(s) - [K(s)P(s)] X_m(s) \\ &= \frac{1}{1 + K(s)P(s)} P_x(s)X(s). \end{aligned} \quad (3.32)$$

Recall that $P_x(s)X(s)$ is the natural response of the spring-mass system, the term $1/[1 + K(s)P(s)]$ denotes an additional amplification/attenuation over the natural response. The passive isolation $P_x(s)$ only attenuates seismic noise above the resonance frequency but amplifies the noise at around the resonance frequency. With a feedback control scheme, the resonances can be suppressed by shaping the controller $K(s)$ such that $|K(j\omega)P(j\omega)|$ is higher than 1 around the resonance frequencies, where j denotes the imaginary number. In fact, the feedback control scheme can provide additional suppression at any frequency as long as $|K(j\omega)P(j\omega)|$ is higher than 1. The frequency band at which $|K(j\omega)P(j\omega)|$ is greater than 1 is referred to the control bandwidth. And, the system is said to be controlled within these frequencies. Also, it is easy to see from Eqn. (3.32) that, as $K(s)P(s) \rightarrow \infty$, the displacement of the mass is completely decoupled from the motion of the suspension point. However, this is, again, only an ideal scenario. In reality, the measurements contain noise and the major challenge of controller design is to shape the controller such that it suppresses the motion while not injecting too much control noise. In fact, motion of the suspension point can be re-injected when using feedback control with relative sensors, even if it is passively isolated. This ruins the passive isolation performance and so feedback control systems must be carefully designed in order to achieve the required isolation. Optimizing the feedback control system for active isolation is the central theme of this thesis and the concept will be revisited in Part II and III.

It is worth mentioning that a major difference between feedforward control and feedback control is stability. The feedforward control is stable as long as all components, such as the controllers and the plants, in the feedforward control scheme are internally stable. However, this is not the case for a feedback control system and the controllers need to be carefully selected to maintain the stability of the system. Let us define the closed-loop transfer function $G(s)$ as

$$G(s) \equiv \frac{1}{1 + K(s)P(s)}. \quad (3.33)$$

The closed-loop transfer function can be represented as a fraction of polynomials, i.e. $G(s) = N(s)/D(s)$, where $N(s)$ and $D(s)$ are polynomials of s . The roots of the polynomial $N(s)$ are called “zeros” while the roots of the polynomial $D(s)$ are called “poles”. Mathematically speaking, the criterion for a stability closed-loop system is that the closed-loop transfer function $G(s)$ must not contain poles with positive real parts (unstable poles). To show why this is the case, suppose $G(s)$ contains a pole at λ , such that it can

⁴Note that the signal path points backward and hence the name “feedback”.

be expressed as a partial fraction

$$G(s) = \frac{a}{s - \lambda} + \dots, \quad (3.34)$$

where a is an arbitrary constant. Then, the inverse Laplace transform gives

$$\mathcal{L}^{-1}\{G(s)\} = g(t) = ae^{\lambda t} + \dots. \quad (3.35)$$

As can be seen, if λ contains a positive real part, then $g(t)$ diverges as t increases, which indicates instability. Therefore, the controller must be designed such that the closed-loop transfer function contains no unstable poles.

In practice, there are two popular tools, Bode plot and Nyquist plot, that can be used to identify the stability of a feedback system without knowing the poles of the closed-loop transfer function. As mentioned previously, the Bode plot is a plot of the frequency response of a transfer function. While the frequency response of a transfer function is generally a complex-valued series, the Bode plot is composed of a magnitude plot and a phase plot. To evaluate the stability of a closed-loop transfer function, the Bode plot of the open-loop transfer function $K(s)P(s)$ is plotted. The unity gain frequency is defined as the frequency where the magnitude response $|K(j\omega)P(j\omega)|$ has a value of 1. For a stable closed-loop system, the phase response at the unity gain frequency must be above -180° . Also, there may be more than one unity gain frequencies and the above criterion must be satisfied at all unity gain frequencies. At the unity gain frequency, the value between the -180° and the phase response is called the phase margin, which is one of the stability margins that is used to characterize the stability of a closed-loop system. Another stability margin, gain margin, is referred to the relative magnitude between unity and the magnitude response as phase response falls below -180° . The gain margin corresponds to the acceptable amplification of the open-loop transfer function before the closed-loop system becomes unstable. In principle, the stability margins can be arbitrary as long as the system is stable. But, in practice, typically, the phase margin is set above 30° - 60° and the gain margin is set to at least $1/3$ (or -10 dB) to account for uncertainties in the control system.

The Bode plot method works for a majority of cases. It is very easy to use and is by far the most popular approach used in KAGRA. However, the Bode plot method could fail in some cases and the Nyquist plot should be used for a proper analysis. The idea of using a Nyquist plot is to identify unstable zeros in the characteristic equation $1 + K(s)P(s)$, recall that the zeros of $1 + K(s)P(s)$ are the poles of the closed-loop transfer function. The Nyquist plot can be obtained by evaluating $1 + K(s)P(s)$ along a contour starting from $-j\infty$ to $j\infty$ and coming back to $-j\infty$ via a semicircular arc on the right-half of the complex plane. The contour is mapped to another complex plane via the function $1 + K(s)P(s)$ resulting in the Nyquist plot. The Cauchy's argument principle states that

the Nyquist plot encircles the origin in the clockwise direction equal to the number of unstable zeros in $1 + K(s)P(s)$ minus the number of unstable poles in $1 + K(s)P(s)$. Alternatively, the open-loop transfer function $K(s)P(s)$ can be used as the map. In this case, the number of clockwise encirclement around the -1 point is counted. For a stable plant, the closed-loop system is said to be stable the number of clockwise encirclement is zero. From here, the closed-loop stability can be determined.

Part II

Active vibration isolation systems in KAGRA

Chapter 4

KAGRA Vibration Isolation Systems

The KAGRA vibration isolation systems are tower suspensions that are capable of providing the necessary seismic isolation for the main optics of the gravitational-wave detector. The multiple-pendulum composition of the suspensions provides the necessary high-frequency passive seismic isolation for the main optics. Not all main optics need to be as still as the test mass mirrors and they have different sensitivity requirements. Hence, there exist multiple types of vibration isolation systems in KAGRA, each correspond to a different extent of passive isolation. The different types of suspensions are introduced in Sec. 4.1 Within a suspension, there are different types of suspension stages, such as the pre-isolator and the marionette, serving different purposes in vibration isolation. A brief introduction of these stages is given in Sec. 4.2. Nevertheless, each stage of the suspensions are equipped with sensors and actuators, which allows the low-frequency active isolation and active damping of the resonance motions of the suspension stages. The various types of sensors and actuators used in KAGRA vibration isolation systems are introduced in Sec. 4.3.

4.1 Types of vibration isolation systems

Fig. 4.1 shows the four types of vibration isolation systems in KAGRA, including the Type-A [30, 31], Type-B [32, 29], Type-Bp [33], and Type-C suspensions. As shown in the figure, the Type-A suspension is a ~ 13.5 m tall suspension and it is the tallest and most complicated vibration isolation system among others. Listing stages from the ground (top) to the optics (bottom), the Type-A suspension is composed of a pre-isolator, five vertical GAS filters (including the one at the pre-isolator), a platform, a marionette, an intermediate mass, and a test mass. Effectively, it is a 9-stage multiple pendulum providing a horizontal seismic noise attenuation proportional to $\sim f^{-18}$ at high frequency. The Type-A suspension provides the greatest amount of passive isolation and there are four such vibration isolation systems at KAGRA, each suspending the test masses of the gravitational-wave detector, as shown in Fig. 4.2. These test masses are the input and




Figure 4.1: Types of vibration isolation systems in KAGRA. Retrieved from Refs. [32, 33]. Left to right: Type-A suspensions, Type-B suspensions,, Type-Bp suspensions, and type-C suspensions. (Not to scale)




Figure 4.2: Optical layout of the KAGRA interferometer with suspension types labeled. Retrieved from Refs. [32, 33]. Red solid circle: Type-A suspensions. Yellow dashed circle: Type-B suspensions. Green dash-dot circle: Type-Bp suspensions. Blue dotted circle: Type-C suspensions.

end test masses along the X-arm and the Y-arm (ITMX, ITMY, ETMX, and ETMY) (the directions of X and Y relative to the Earth are arbitrary but they are orthogonal) which defines the 3-kilometer Fabry-Perot cavities of the interferometer.

The Type-B suspension can be seen as a truncated version of the Type-A suspension. Its composition is very similar to that of the Type-A suspension. It has a pre-isolator,

three vertical GAS filters, an intermediate mass, and a test mass stage¹. It has 5 pendulums stages providing a $\sim f^{-10}$ horizontal seismic noise attenuation at high frequency. Compared to the Type-A suspensions, there are a few differences besides having fewer stages. First of all, the marionette stage is missing in the Type-B suspension and therefore the steering of the optics relies on the active components at the intermediate mass and the optics stage. The Type-A payload, consisting of the marionette, the intermediate mass, and the payload, hangs from the platform via four wires per stage. Their recoil masses constitutes a multiple pendulum separated from the main pendulum chain and are also hanging from the platform. In contrast, the Type-B payload, consisting of the intermediate mass and the optics, hangs from the bottom filter (the last vertical filter). And, the recoil masses do not compose a separate chain like that in the Type-A suspension. Instead, they each hangs from the upper stage, i.e. the intermediate recoil mass hangs from the bottom filter and the optics recoil mass hangs from the intermediate mass. There are four Type-B suspensions in KAGRA and they suspend the signal recycling mirrors (SRM, SR2, and SR3) and the beamsplitter (BS), as shown in Fig. 4.2.

The Type-Bp suspensions are supposed to be the compactified version of the Type-B suspension and hence the suffix “p”, which stands for “pico”. But, the differences between the Type-Bp and Type-B suspensions are significant. Firstly, the Type-Bp suspension do not have the pre-isolator, unlike the Type-A and Type-B suspensions. Instead, the first vertical filter, which is a standard filter, is attached on a motorized horizontal platform that can be used to adjust the coarse alignment of the whole chain. The Type-Bp suspension is effectively a triple-pendulum, with the bottom filter, the intermediate mass, and the optics being the three pendulums. The Type-Bp suspensions are used for the power recycling mirrors including the PRM, PR2, and the PR3.

Nevertheless, the Type-C suspensions table-top suspensions that are significantly smaller than the other suspensions. For this reason, multiple Type-C optics can share the same vacuum chamber, such as the input and output optics of the input mode cleaner as shown in Fig. 4.2. The Type-C suspensions are used for the auxiliary optics such as the input and output mode cleaners. They contain actuators and sensors at the optics stage only for the purpose of steering and active damping of the optics. There are no active seismic isolation element in the Type-C suspensions. Therefore they are only briefly mentioned here and further discussions are oriented around the Type-A, Type-B, and the Type-Bp suspensions.

¹The optics of the Type-B suspension are not the test masses of the gravitational-wave detector but the optics stage is also referred to the test mass stage by convention.

4.2 Suspension stages

The suspensions are composed of different suspension stages that serve different purposes in vibration isolation. For example, the Type-B suspension contains a pre-isolator stage, 2 GAS filter stages (3 if counting the top filter in the pre-isolator), an intermediate mass stage, and a test mass stage, as shown in Fig. 4.3. The different type of stages are the building blocks of a suspension and, in this section, the function and the composition of the suspension stages are explained.




Figure 4.3: Type-B suspension.

4.2.1 Pre-isolator

The pre-isolator is the first seismic isolation stage of the Type-A and Type-B suspension and it interfaces the ground and the rest of the suspension and eventually the optics. The pre-isolator is often referred to the inverted pendulum (IP), which is the most iconic

structure in the pre-isolator. The inverted pendulum is a miniaturized version of the inverted pendulum of the Virgo suspension [34]. The three inverted pendulum legs connects the pre-isolator base, which is the attached to the ground via an elastic flexure, and the pre-isolator table. This suspends the pre-isolator table like a stable inverted pendulum and, with small displacements, the inverted pendulum legs effectively constrain the motion of the pre-isolator table on a horizontal plane. As shown in Fig. 4.4, arc weights are placed along the edge of the pre-isolator, which adds weights to the inverted pendulum and reducing its resonance frequency because of the anti-spring effect. In practice, thin pieces of arc weights are added in a symmetric way until the resonance frequency becomes low enough to attenuate the secondary microseism between 0.1 Hz and 0.5 Hz, but not too low such that the inverted pendulum becomes bistable. Typically, the resonance frequency is around 0.06 Hz and 0.1 Hz. This provides seismic attenuation to the secondary microseism of 4 to ~ 11 times at 0.2 Hz.




Figure 4.4: The Type-B pre-isolator.

There are three columns extending from the base, supporting a ring platform where three LVDT assemblies around the pre-isolator table can be found. The LVDT stands for linear variable differential transformer, which is a type of relative displacement sensor to be described in Sec. 4.3.1. The assembly is composed of an LVDT measuring the a relative displacement between the pre-isolator table and the ground, and a coil-magnet actuator that can be manipulated to apply force approximately to the same direction where the LVDT is sensing. The three LVDT directions spans the full control space of the pre-isolator on the horizontal place so the pre-isolator can be fully controlled in two horizontal directions and one rotational direction. On the same ring platform, three linear stepper motors are attached. The sliding blocks of the stepper motors are connected to

the pre-isolator table via soft copper blade spring so they do not alter the dynamics of the inverted pendulum. But, the motorized sliding blocks can be used to adjust the static alignment of the pre-isolator table since it has a significantly larger dynamic range compared to the LVDT. However, the stepper motors are too noisy to be used for feedback control so they are turned off after the initial position adjustment.

On the pre-isolator table, besides the arc weights, there are the three geophone pods and a top vertical filter. The geophone pods are attached along the circumference of the pre-isolator like the LVDT assemblies. Inside the pod, there is an L-4C geophone from Mark products and an pre-amplifier circuit board designed by NIKHEF [29]. The geophone converts the velocity, relative to an inertial frame, of the pre-isolator table to voltage, which can be used in a feedback system for active isolation. As for the top vertical filter, it is composed of three (or more) compressed spring blades forming one vertical anti-spring that hangs the rest of the suspension chain. As part of the top filter, there is a so-called “fishing rod”, named after its resemblance of a fishing rod, that is part of the linear stepper motor assembly that serves the same purpose as those for the horizontal direction. The top filter anti-spring is considerable larger than the standard filters that forms the pendulum chains and the resonance frequency was tuned in situ to around 0.5 Hz. The blade tips of the anti-spring blades are connected to the keystone, which holds the suspension wire connecting to the stages below. There is an LVDT assembly around the keystone that provide sensing and actuation of the keystone in the vertical degree of freedom.

4.2.2 Standard filters

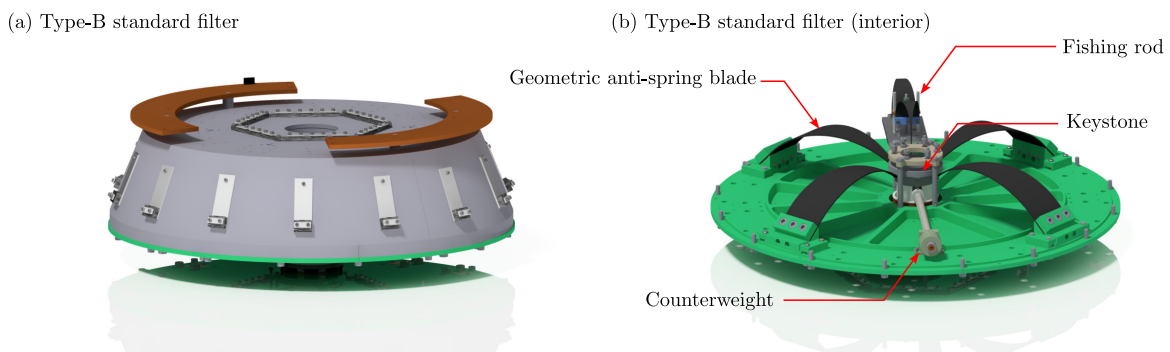


Figure 4.5: (a): The Type-B standard filter. (b): The interior of a standard filter

The standard filters are pendulums of the suspension. On one hand, the weight of the filter assembly act as the mass of the pendulum. As it is hanged from an upper stage via a single suspension wire, it literally acts as a pendulum providing passive isolation in the horizontal direction. On the other hand, as shown in subplot (a) in Fig. 4.5, enclosed in the standard filter is a GAS filter that suspends the next stage of the suspension, much

like the top filter, providing passive isolation in the vertical direction. Everything related to the GAS filter is similar to that of the top filter, i.e. it has an LVDT assembly and a fishing rod, but in a smaller scale. While the standard filter can, in principle, move freely subjected to the constrain of a pendulum, there are no sensors and actuators that can be used to control these degrees of freedom. But, there is one exception, that is, the bottom filter.




Figure 4.6: (a): Type-Bp bottom filter. (b): Type-B bottom filter.

The bottom filter is the last pendulum above the payload stages, such as the marionette and the intermediate mass. It is very similar to the standard filter besides it can be controlled more than just the GAS degree of freedom. The way it can be manipulated varies for Type-A/Type-Bp and the Type-B systems. For Type-A/Type-Bp suspensions, there exist structures extended from the bottom filter enclosure, each holding the primary coil of an LVDT, as shown in subplot (a) in Fig. 4.6. The secondary coils are attached to a structure that extends from the base structure of the suspension (effectively, the ground) so there are LVDTs measuring the relative displacement between the bottom filter and the ground. There are six of such LVDT assemblies so the bottom filter can be controlled in all six degrees of freedom.

As for the Type-B suspensions, the bottom filter cannot be controlled like those in the Type-A and Type-Bp suspensions. Instead, for the Type-B suspensions (and Type-Bp), there are four moving masses on top of the bottom filter that can be moved using piezo linear actuators (picomotors), as shown in subplot (b) in Fig. 4.6. Moving the masses changes the center of gravity of the bottom filter and hence determines the tilt of the bottom filter. The tilt eventually translates to the tilt of the intermediate recoil mass. Therefore, it can be used to adjust the tilt of the intermediate recoil mass relative to the intermediate mass. Like the stepper motors, the picomotors are too noisy to be used for feedback control purposes and hence they are only used for initial adjustments.

4.2.3 Marionette and intermediate stages

The marionette and the intermediate masses stages above the suspended optics. In the case of the Type-A suspension, the marionette suspends from a transitional stage called




Figure 4.7: (a): The intermediate mass and recoil mass. (b): Half section view.

the platform, which is suspended from the bottom GAS filter. The intermediate mass is then suspended from the marionette. In the case of Type-B and Type-Bp suspensions, there is no marionette stage. Instead, there is only the intermediate mass stage. The intermediate mass stage in the Type-A adds an extra f^{-2} noise attenuation at high frequency.

The marionette in the Type-A suspension and the intermediate mass stage in the Type-B and Type-Bp suspensions serve a similar purpose, that is, to provide large range steering to the optics that is suspended below. Therefore, here, only the intermediate mass stage is described. As shown in Fig. 4.7, the intermediate mass is surrounded by an intermediate recoil mass. There are six OSEMs attached around the intermediate mass, three horizontal and three vertical (In the case of Type-A suspensions, they are photosensors instead of OSEMs.). The OSEM stands for Optical sensing and electromagnetic actuator, which is to be introduced in Sec. 4.3.3. They are similar to the LVDTs that are used for the GAS filter and pre-isolator and they provide both sensing and actuation in a compact form factor [35]. The OSEM body (dark grey as shown in Fig. 4.7) is attached on the recoil mass and it contains a measurement unit and a voice coil. The measurement unit measures the displacement of the OSEM flag, which is shown in beige color in the figure, in the direction of the flag axis. And, the voice coil is also aligned to the flag axis and it can be manipulated to apply force on a magnet attached to the flag. With the six horizontal and vertical OSEMs, the intermediate mass can be controlled in all six degrees of freedom.

4.2.4 Test mass

Like the intermediate mass stage, the test mass stage is composed of a recoil mass and the test mass, i.e. the suspended optics, which appears to be transparent in Fig. 4.8. The optics and the recoil mass are each suspended from the intermediate mass via two loops of suspension wires clamped at the intermediate mass (In the case of the Type-A suspension, the recoil mass is suspended from the intermediate recoil mass.). As shown in Fig. 4.8, there are four OSEMs at the back side of the optics. Unlike the OSEMs at




Figure 4.8: (a): Test mass stage (front). (b): Test mass stage (back).

the intermediate mass stage, the OSEMs at the test mass stage do not provide sensing of the suspended optics. This is due to the frequent breakage of the OSEM flags as the motion of the optics became excessive. Instead, the OSEM only provide actuation at four corners of the optics. The sensing of the optics displacement is done by an optical lever at the exterior of the vacuum chamber, which is described in Sec. 4.3.5 and 6.2.1. Both the arrangement of the OSEMs and the optical lever only allows the optics to be sensed and actuated in the longitudinal direction (parallel to the optical axis of the suspended optics) and two rotational degrees of freedom, pitch and yaw.

4.3 Sensing and actuation technologies

In this section, the sensors and actuators in the KAGRA vibration isolation systems are discussed.

4.3.1 Linear variable differential transformer (LVDT)




Figure 4.9: Top filter LVDT assembly at the pre-isolator stage.

The linear variable differential transformer (LVDT) is a relative displacement sensor. It is used in majorly in the tower part of the suspension including the pre-isolator and the GAS filters. As an example, Fig. 4.9 shows the cross-section view of the top GAS filter in the pre-isolator. As can be seen, the sensing element of the LVDT is composed of two separate pieces, the primary coil and the secondary coil. For KAGRA vibration isolation systems, a coil-magnet actuator is also installed at close proximity of the LVDT and the two forms an LVDT assembly. In Fig. 4.9, the coil-magnet actuator of the top filter is located below the LVDT coils and in this case, the coil is attached to the moving keystone while the magnetic yoke is attached to the pre-isolator table (underneath).

The quantity that the LVDT is measuring is the relative displacement between the primary coil and the secondary coil in the direction parallel the coil axis². When the primary coil and the secondary coils are mounted at different objects, e.g. the ground and the pre-isolator table, moving relative to each other, the LVDT then sense the relative displacement between the two objects. While the primary coil and the secondary coils are aligned coaxially, they do not need to be in contact. This means the sensing surfaces can move freely and the sensing do no alter the dynamics of the sensed objects. For this reason, the LVDT becomes a perfect candidate for active vibration isolation systems that need to behave like oscillators, i.e. pendulums and spring-mass systems.




Figure 4.10: Conceptual readout circuit of the LVDT.

The working principle of the LVDT is as follows. As shown in Fig. 4.10, the primary coil is being injected with an excitation signal $V_e(t)$. The excitation is sinusoidal such that

$$V_e(t) = V_0 \sin(\omega t), \quad (4.1)$$

where V_0 is the amplitude, ω is the frequency, and t is time. The induced voltages in the secondary coils are $V_1 \sin(\omega t)$ and $V_2 \sin(\omega t)$, where V_1 and V_2 are the amplitudes of the voltage induced in the first and second secondary coils, respectively. Their difference is being amplified to give a differential signal

$$V_d(t) = A(V_2 - V_1) \sin(\omega t), \quad (4.2)$$

²Typical LVDT contains a moving ferromagnetic core and the LVDT measures the displacement of the core. The coils are immobile. However, KAGRA LVDTs do not contain the ferromagnetic core possibly because of the presences of the coil-magnet actuator in proximity.

where A is the amplification of the amplifier. The differential signal and the excitation signal are mixed, i.e. multiplied, to give a signal

$$V_d(t)V_e(t) = AV_0(V_2 - V_1) \sin^2(\omega t). \quad (4.3)$$

In practice, the relative phase between the differential signal and the excitation signal is non-zero and there exists a phase shifting circuitry that needs to be tuned to maximize the readout. For simplicity, let us assume that the relative phase is zero. Then, by the double angle identity of cosine, the product of the two signals can be rewritten as

$$\begin{aligned} V_d(t)V_e(t) &= AV_0(V_2 - V_1) \left[\frac{1 - \cos(2\omega t)}{2} \right] \\ &= \frac{1}{2}AV_0(V_2 - V_1) - \frac{1}{2}AV_0(V_2 - V_1) \cos(2\omega t), \end{aligned} \quad (4.4)$$

which is a superposition of a static term and an oscillatory term with frequency 2ω . The oscillatory term is then removed by low-passing the product of the two signals with a real low-pass filter circuitry. And, the final readout is the remained static term that is proportional to the voltage difference in the two coils $V_2 - V_1$. Now, the differential voltage $V_2 - V_1$ of the secondary coils depends on the relative position of the primary coil in the axial direction, which determines the number of primary coil turns directly facing each secondary coil. Therefore, the low-passed differential signal is linearly dependent on relative axial displacement between the coils. And, when the primary coil and the secondary coils are mounted on two objects moving relative to each other, the readout is a signal proportional to the relative displacement. To obtain the proportionality constant, the LVDTs were calibrated with calipers in practice and this is possible because of the large linear range of the LVDTs.

Note that here the induced voltages in the secondary coils V_1 and V_2 are assumed to be constants, but in reality, they are not. In practice, the excitation signal is chosen such that the frequency is much higher than that of the suspension displacements that are sensed. To give perspective, the excitation signal used at KAGRA is a 10 kHz signal with 5 V peak-to-peak. Whereas, the motion of the suspension degrees of freedom can be assumed to be dominated by resonant motions at below few 10 Hz. Therefore, the difference is large enough for the above analysis to remain valid.

4.3.2 Geophone and inertial sensors

The geophone is a type of sensor that converts velocity of the sensed object to a voltage. And, the velocity readout is relative to an inertial frame of reference. For this reason, the geophone is a type of inertial sensors, which can be used to achieve active seismic isolation when used in an active control scheme, such as feedback control. The pre-isolators of the

KAGRA suspensions are all installed with the L-4C geophone and they are meant to be used for active seismic isolation of the main optics. The geophone is not the only type of inertial sensors that is used at KAGRA. Seismometer is also a type of inertial sensors that is used in to monitor the ground motion at KAGRA. There are other types of inertial sensors, such as the monolithic folded pendulum [28, 36], and the development of inertial sensors to achieve higher sensitivity is still an active area in active isolation research for gravitational-wave detectors.

The inertial sensors can be mainly divided into two types, one with passive sensing and the others with active elements. But, they are all based on the same principle, that is, sensing the relative motion of a suspended proof mass and inferring the inertial motion of the suspension point. In the case of the passive ones, the relative motion of the suspended proof mass is measured in the form of relative velocity between suspension point and the proof mass. Examples of passive inertial sensors are geophones and seismometers. In the geophones, the relative velocity is measured by a coil and a magnet, which are fixed either on the proof mass and the enclosure. The inertial motion of the suspension point can then be inferred from the relative velocity using the inverse dynamics of the suspended proof mass. The detailed mechanism of the geophone is given in Sec. 6.1.2, where the inverse dynamics of the geophone is derived. As for the active inertial sensors, the relative motion between the proof mass and the suspension point is being actively nullified by a feedback control system. The feedback force, being proportional to the acceleration of the suspension point, can then be used to infer the inertial acceleration of the suspension point. The advantage of an active scheme is that it allows the incorporation of a readout mechanism with lower range but with higher sensitivity, such as an interferometer [36].

4.3.3 Optical sensing and electromagnetic actuator (OSEM)

The OSEM stands for optical sensing and electromagnetic actuator and is compact device that contains both non-contact sensing and actuation in one small form factor. The original version of the OSEM is called BOSEM (B for Birmingham) and was developed for the LIGO detector [37]. Several iterations have been developed as an improvement to the original BOSEM and the version presented here is the one used in KAGRA [35]. As shown in Fig. 4.11, the OSEM is used at the intermediate mass stage of the vibration isolation systems. For Type-A systems, another sensor called photocoupler position sensor, photo-reflective sensors, or simply photosensors are for the marionette and the intermediate mass instead and this is to be discussed in Sec. 4.3.4. The OSEM was originally planned to be used also for the sensing and actuation at the test mass stage. However, the design of the OSEM requires the sensor flag to be physically inserted into a small extrusion hole, which makes it vulnerable to large relative motion where the sensor flag is dislocated due to mechanical interference [35]. This is very undesirable for gravitational-wave detectors




Figure 4.11: Section view of the OSEM.

where the optics are sealed in complicated vacuum systems and resetting the sensor flag could potential take time off the commissioning and observation period. Therefore, the sensing part of the OSEM at the test mass stage is eliminated [32, 35].

The sensing part of the OSEM is essentially a shadow sensor, which is a photodiode detecting the presences of a shadow (or the absence of light). As shown in Fig. 4.11, the OSEM body is divided to two sides corresponding to the transmitter and the receiver. The transmitter consist of an light-emitting diode (LED) and a small collimating lens (not labeled). The light from the LED travels through a small slit and reach the receiver side, which contains a photodiode. As the sensor flag is inserted between the transmitter and the receiver, it partially blocks the light. This is reflected in the reduced voltage readout of the photodiode. And, since the amount of reduced readout a function of the amount of clipped light, which depends on the position of the sensor flag, displacement sensing can be achieved. The calibration relationship between photodiode readout and the displacement of the flag is sigmoidal, which roughly approximated by an error function assuming a Gaussian light profile, as discussed later in Sec. 6.1.1. The linear region of the calibration function is around 1 mm [35] and this defines the linear range of the OSEM sensor. As for the actuation part, the sensor flag is, in fact, magnetically attached to a larger magnet as shown in Fig. 4.11. A coil is wrapped around coil bobbin of the OSEM body and actuation signals can be injected into the coil so a force can be applied to the magnet where the sensor flag is attached.

4.3.4 Photocoupler position sensor (photosensor)

The photosensor is a novel sensor specially developed for KAGRA and is used to replace the OSEMs for the sensing of the marionette and intermediate mass stages of the Type-A system. The reason why the photosensor is used instead of the OSEM is due to the cryogenic nature of the payload of the Type-A suspension. The drastic change in

temperature and the slight difference in thermal expansion of the suspension wire causes the suspension to drift during cool down [38]. And, this calls for an increase of the 1 mm linear sensing range of the OSEM. To resolve the issue, the photosensor is developed and it has a linear range of 10 mm.




Figure 4.12: Conceptual diagram of a photosensor.

Fig. 4.12 shows a conceptual diagram of a photosensor. In the figure, the photosensor measures the relative displacement between a reflective surface and the LED. In the case of KAGRA suspensions, the stainless steel surfaces of the marionette and the intermediate mass are polished to increase the reflectivity for this purpose. Unlike the OSEM, the photosensor has one transmitter but two receivers, i.e. one LED and two photodiodes. As the reflective surface moves towards the LED and the photodiodes, the reflected beam shifts (much like the length sensing optical lever in Sec. 6.2.1) and this changes the amount of light that enters the photodiodes. Also, when the reflective surface tilts (along the axis out of the screen), the amount of light that enters each photodiode also changes. But, the change in the two readouts are differential in this case, one increase and the other decreases. Therefore, the common mode of the readouts can be used as a displacement readout. However, for a two-sensor configuration, the readout cannot be decoupled from a tilt along photodiode-LED-photodiode line (vertical in Fig. 4.12). The ideal configuration is to have two such units of photodiode-LED-photodiode placed such that the photodiodes form four corners of a square. In this way, the common mode of the four photodiode readouts give a displacement readout that is decoupled from all tilt of the reflective surface.

4.3.5 Optical lever

The optical lever is given its name because part of the test mass displacement sensing device is composed of an optical lever. Technically, an optical lever is an optical device that senses the tilt of a reflective surface by measuring the shift a reflected beam spot.

The distance from the reflected surface to the reflected beam spot forms a lever arm that amplifies the tilt motion of the reflected surface and hence is given the name. The tilt sensing of the test mass stage of the KAGRA suspensions rely on the optical lever mechanism and is therefore the test mass sensor is referred as the optical lever. However, the same optical lever beam can be used to measure a shift in the reflective surface besides a tilt. This is achieved by measuring the reflected beam spot at some position behind a convex lens. The details of the test mass sensing is given in Sec. 6.2.1, where the relationship between the optical lever readout and the displacements of the suspended optics is derived. This section serves as a brief introduction.




Figure 4.13: Cross-section view of a Type-B suspension.

The optical lever consists of a transmitter and a receiver. They are optical components that are located at the exterior of the vacuum chamber containing the vibration isolation systems. Fig. 4.13 shows the cross-section view of a Type-B suspension. As can be seen, the optical lever beam enters the vacuum chamber via a lower viewport. In this case, the light source is a superluminescent diode mounted on a optical breadboard next to the lower viewport (not shown in the figure). The optical lever beam then is reflected off the suspended optics and exits the vacuum chamber via the upper viewport. The beam is then steered and made horizontal along an optical breadboard, by a steering

mirror, containing the readout optics. It is worth noting that this configuration, where the incidence plane is parallel to a vertical plane, is exclusive to the Type-B systems. For the Type-A and Type-Bp suspensions, optical lever beam travels along a horizontal plane instead.




Figure 4.14: Conceptual readout optics of the optical lever. Diagram created using component template provided by Ref. [39].

A conceptual diagram of the optical lever readout optics is shown in Fig. 4.14. As shown in the figure, the beam reflected from the suspended optics is divided into two paths via a beamsplitter. One of the beam hits the tilt-sensing quadrant photodiode (QPD), forming the traditional optical lever. The other beam travels through a convex lens and then eventually reaches the length-sensing QPD. The QPDs measure the shift of the beam spots caused by a tilt or a shift of the suspended optics. With a proper placement of the length-sensing QPD behind the lens, the length-sensing QPD becomes only sensitive to the longitudinal shift of the suspended optics but not a tilt in pitch or yaw.

There are four photodiode readouts in a QPD, each measuring the light power of beam that enters the corresponding quadrant. Suppose the voltage readout of the quadrants A, B, C, and D (following the convention in Fig. 4.14) are V_A , V_B , V_C , and V_D , respectively, a horizontal shift of the beam spot can be written as a proportionality

$$x \propto \frac{V_B + V_D - V_A - V_C}{V_A + V_B + V_C + V_D}, \quad (4.5)$$

where x is the beam spot displacement in the horizontal direction. Similarly, the vertical beam spot displacement can be written as

$$y \propto \frac{V_A + V_B - V_C - V_D}{V_A + V_B + V_C + V_D}, \quad (4.6)$$

where y is the vertical beam spot displacement. From here, the beam spot displacements can be calibrated in physical units by moving the QPDs in horizontal and vertical directions relative to a static beam using a micrometer stage. With the beam spot dis-

placements obtained, the longitudinal, pitch, and yaw displacements can be derived using ray transfer matrix analysis [40] and the detailed analysis of the optical lever is given in Sec. 6.2.1, where the geometric sensing matrix of the optical lever is derived.

Chapter 5

Control Topology

In this chapter, the control architecture of the vibration isolation system used in KAGRA is described in detail. This chapter gives background information about the active isolation systems at KAGRA. The purpose is to establish the stage in which the research is built upon.

Active control is done via converting some sensing readouts into some actuation signals, achieving feedback or feedforward control. In general, this transformation can be arbitrary and even non-linear. However, control systems without a defined architecture can be hard to realize and design without involving advanced control techniques. A control architecture defines how signals are wired, how signals flow and transform, and how sensed signals are eventually converted into actuation signals.

KAGRA inherited the LIGO real-time code generator (RCG) [41], which is developed by the LIGO Control and Data Systems group (CDS). RCG uses MathWorks Simulink [42] as a drawing tool to design the control architecture as a Simulink model. The Simulink model is then compiled into software to be run on front end control computers to realize real-time control. Simulink provides many components (parts) that can be used as building blocks to construct a control architecture. But, not all Simulink parts are supported by RCG. Instead, RCG has a library called CDS parts that contain essential parts required to assemble and realize a control architecture. There are many parts related to real-time implementation, but not related to control system design, such as parts that interface actual hardware like an analog-to-digital converter (ADC). The critical building blocks related to control architecture design are matrix and filter. There are also summing junction and gain blocks. But, they are simply special cases of matrices and filters and they are not required for explaining the control architecture used at KAGRA.

Fig. 5.1 shows a simplified example control architecture for one stage of the suspension that has three sensors, three control basis, and three actuators. Note that the number of sensors, control basis, and actuators are generally not the same. This is a general layout for suspension stages with multiple degrees of freedom (DoFs) and it shows how sensing signals are passed through various components to become actuator signals that are used to




Figure 5.1: Example control architecture with three sensors, three control degrees of freedom, and three actuators.

command the actuators. In the figure, each arrow represents a process variable. Filters are linear-time-invariant (LTI) filters that can be represented by transfer functions, which can be quantified by frequency-dependent amplitude and phase responses. And, matrices are linear transformations in the usual mathematical sense. Filters and elements in matrices can be accessed and modified in real-time whereas the inter-connections between them are defined within the architecture. Most of the filters and matrices require some initial configuration, design, or tuning. The intended purpose of these components are described in following paragraphs.

There are three levels in the control architecture, as enclosed by red, green, and blue brackets in Fig. 5.1. They correspond to three different signal bases (coordinate systems):

1. Sensor basis,
2. Control basis, and
3. Actuator basis.

Generally speaking, while they span the same physical space, these coordinate systems are not necessarily aligned. The sensor basis is defined by the directions in which the sensors physically measure. The control basis is a user defined coordinate system that is used for control. Unless otherwise specified, the control basis is chosen to be the three translation directions and the rotational directions along the axes of these directions. Detail description of the control basis is given in Sec. 5.1. Lastly, like sensor basis, actuator basis is defined by the directions where the actuators physically apply force.

Sensor signals from the left of Fig. 5.1 are raw sensor readouts sourced from the analog-to-digital converters. The raw readouts have arbitrary scale and must be converted into proper physical units such as micrometer μm . The sensors are typically calibrated to obtain a calibration factor that converts the raw sensor readouts to proper displacement units. For inertial sensors, such as geophones and accelerometers, they need frequency-dependent calibration filters, which correspond to the inverse frequency responses of the sensors. This is done via input filters as shown area enclosed by the red bracket in Fig. 5.1. The aforementioned calibration factors are also usually implemented in the input filters.

Details of sensor calibration are given in Sec. 6.1. Besides calibration, the input filters are typically be used for other signal processing purposes such as whitening/dewhitening.

The input filters give raw displacement readouts that correspond to displacements that the sensors are directly measuring. They are often not measuring in the directions that align with the control basis and the raw displacements contain a linear combination of the desired control displacements. Therefore, the input signals need to be transformed to displacement signals that aligns with the control coordinate system. This is achieved via a sensing matrix as shown in Fig. 5.1. Then sensing matrix can be initially derived from the physical locations of the sensors using geometry. This type of sensing matrix is referred to the geometric sensing matrix. The derivation is often straightforward as the geometric locations of the sensors can be obtained from the computer aided design (CAD) file of the hardware. Refs. [43, 44] provide detail derivation of the geometric sensing matrices for the beamsplitter (BS) suspension and the signal-recycling (SR) mirror suspensions, except for the sensors (optical levers) at the optics (test mass) stage. The derivation of the optical lever sensing matrices is given in Sec. 6.2.1 as an example. Although the geometric sensing matrix provides an initial sensing matrix that theoretically aligns the sensors to the control basis, cross-coupling ratios between different displacement readouts at the $10^{-2} - 10^{-1}$ level are often observable, especially at the test mass stage. Therefore, they can be further tuned to finely align the sensor basis to the desired control basis. This fine alignment process is called “diagonalization” and is detailed in Sec. 6.2.1. Again, the sensing matrix is not necessarily a square matrix.

The output of the sensing matrix gives the calibrated displacements in the desired control basis. The signal here is often called the “error signal”, which is passed directly into the control filter. Often, there is a summing junction in between the sensing matrix and the control filters. The summing junction takes a reference “setpoint”, which is the desired static position of the corresponding DoF, and outputs the difference between the setpoint and the displacement. This enables position control of the suspension. But, for simplicity, it is omitted here since it is simply an offset to the measured displacement.

The control filters in Fig. 5.1 are feedback control filters. They are frequency-dependent filters and they ultimately determine the control bandwidth and performance of the control system. The filters convert the displacement signals into actuation signals in the control coordinate system. The design of these filters are rather involved. In this thesis, two methods are proposed for designing these filters:

1. PID-based controller with critical design criteria (Sec. 6.4), and
2. \mathcal{H}_∞ optimization (Sec. 8.2.3).

Both methods work with minimal design freedom and the result control filters are optimal in two different manners.

The control filters output actuation commands in the control basis and they need to be transformed into the actuation basis. This is achieved via an actuation matrix as shown in Fig. 5.1. The idea behind an actuation matrix is similar to that of a sensing matrix but works in an opposite way. Another difference between a sensing matrix and an actuation matrix is that the actuation matrix preserve no information about calibration in any sense. This is because the actuators are not calibrated and this is generally not necessary. Therefore, the actuation matrix does not only convert the actuation signals from the control coordinate system to the actuation coordinate system, but there is also a freedom to scale the signals by an arbitrary amount. Again, the actuation matrix can be straightforwardly derived from the positions of the actuators and matrices for the BS and SR suspensions are given in Refs. [43, 44]. An example is also given in Sec. 6.2.2. Like sensing matrices, the geometric actuation matrix is often not sufficient and require fine adjustment (diagonalization). Sec. 6.2.2 gives detailed explanation on this topic.

At last, the output filters may be used as whitening filters depending on the hardware configuration. The output filters are sometimes used as “balancing gains”, which equalizes the strength of each individual actuators. The effect of balancing is degenerating with actuation matrix diagonalization, which is an attempt to align the actuators to the control basis. The output filters send signals to the digital-to-analog converter (DAC) in the Simulink model and the front end computer passes these signals to the actual DAC, which eventually converts the digital signals to voltage to be applied on the actuators. This completes a signal flow on the software side. The actuators actuate the suspension and the motion is sensed by the sensors and this completes a feedback loop.

It is worth mentioning that Fig. 5.1 is a simplified stage-wise picture of the control architecture of the real control system. In reality, there are more components, such as the aforementioned setpoints, and they will be introduced appropriately in relevant sections.

Fig. 5.1 not only shows how control signal flows within the real-time model, but also defined some tasks that are required to partially complete the control system. The tasks are mentioned in previous paragraphs and here is a brief reiteration:

1. Calibration,
2. Aligning the sensors to the control basis (sensor diagonalization),
3. Designing feedback control filters to fulfill control objectives, and
4. Aligning the actuators to the control basis (actuation diaognalization).

These tasks are referred to part of the suspension commissioning tasks. Sections in Chap. 6 provides well tested methods for tackling these tasks and the results are shown.

5.1 Control basis

This section describes how the control coordinate system is defined.

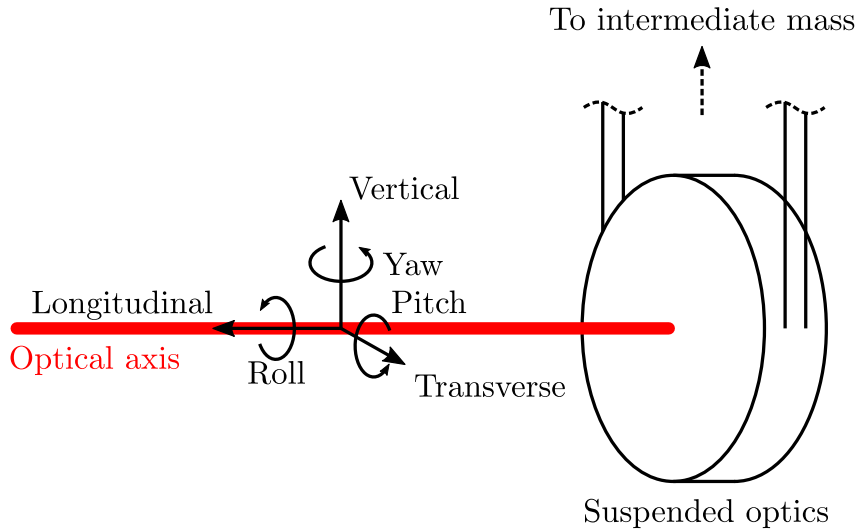


Figure 5.2: Local coordinate system of a suspended optics.

Fig. 5.2 shows a simplified diagram of a suspended optics with all six directions labeled. The red line denotes the optical axis of the main interferometer. The optical axis defines first translational degree of freedom of the suspended optics. This direction is referred as the longitudinal direction. The vertical direction is defined as the direction opposite the gravity, i.e. upward. Then, the transverse direction is defined such that the vertical direction is a cross-product between longitudinal and transverse direction. As for the rotational degrees of freedom, roll, pitch, and yaw, they are simply defined as the rotation along the longitudinal, transverse, and vertical axis, respectively.

Being a multiple pendulum, each stage of the suspension has their own locally defined degrees of freedom that aligns with the coordinate system of the optics. In principle, almost all stages are free to move in all six degrees of freedom but not all degrees of freedom are important and not all are controllable. To exemplify, Fig. 5.3 shows a simplified diagram of a Type-B suspension. Firstly, the inverted pendulum (IP) suspends a table (shown in red) from the ground. The inverted pendulum legs constrain the table to move on an horizontal plane so the table can only move in longitudinal, transverse, and yaw direction. All three degrees of freedom are measured relative to the ground. For geometric anti-springs (GAS), they only have vertical degree of freedom and it is defined as the relative motion between the keystone (shown in green, blue, and yellow), which hangs the rest of the suspensions, and the suspension point, i.e. the body of the GAS filter enclosure (the trapezoids in Fig. 5.3). For type-A and type-Bp suspensions, there are sensors and actuators around the entirety of the bottom filter (BF) so all six freedoms can be controlled. But, for type-B suspensions, the motion of the bottom filter is not sensed so only the keystone (shown in yellow) is controlled. For the marionette (MN)

and intermediate mass (IM, shown in purple), they can be controlled in all six degrees of freedom and all displacements are sensed relative to the recoil masses. Lastly, the test mass (TM, shown in cyan), i.e. the suspended optics, can move in all six degrees of freedom but only longitudinal, pitch, and yaw are controllable, with displacement measured relative to the ground. This gives a total of 15 controllable degrees of freedom for a Type-B suspension (3 inverted pendulum, 3 GAS filters, 6 intermediate mass, and 3 test mass degrees of freedom). For Type-A suspensions, there are 29 controllable degrees of freedom (3 inverted pendulum, 5 GAS filters, 6 bottom filter, 6 marionette, 6 intermediate mass, and 3 test mass degrees of freedom). And for Type-Bp suspensions, there are 17 controllable degrees of freedom (2 GAS filters, 6 bottom filter, 6 intermediate mass, and 3 test mass degrees of freedom). Note that the controllable degrees of freedom are not necessarily independent. For example, moving the inverted pendulum table in the longitudinal direction would move the whole chain below it, including the test mass, in the same direction. Also, in Fig. 5.3, the vertical degree of freedom of the intermediate mass happens to be exactly the the bottom filter GAS DoF.

The desired control basis doesn't not need to be exactly the controllable degrees of freedom. The use of a stage-wise Cartesian basis and Euler angles is chosen for simplicity. It aligns with the interferometer arms and this simplifies tasks such as alignment and sensor correction (Sec. 7.2). In principle, the controllable degrees of freedom can be transformed into other unusual basis for control. One example would be the eigenvectors of the suspension. In such basis, all degrees of freedom will be decoupled with each other and each degree of freedom behave like a simple spring-mass, which may simplify the controller design [30]. However, it is obvious the decomposition is not possible for all modes of the suspension as many degrees of freedom are not sensed. Using the eigenbasis decouples the dynamics of the system but couples other things such as sensor noises and disturbances as it couples many controllable degrees of freedom. This can complicate optimal controller designs where the disturbances and noises are considered. Also, controllers in the eigenbasis can be mapped into other basis. Therefore, using the eigenbasis has no significant advantages over the Cartesian and Euler bases as long as the controllers are optimized to achieve certain objectives. Hence, unless otherwise specified, the stage-wise Cartesian basis and Euler angles (e.g. IP longitudinal, IM transverse, TM pitch, etc.) will be used for further discussions in this thesis But, it is worth to mention that the methods to be discussed are not limited to a certain control basis.

5.2 Control problem formulation

Assuming that the sensing matrices and actuation matrices properly aligns the sensors and actuators to the desired control basis, each control degree of freedom becomes an individual system with one sensor and one actuator. Each individual system has one

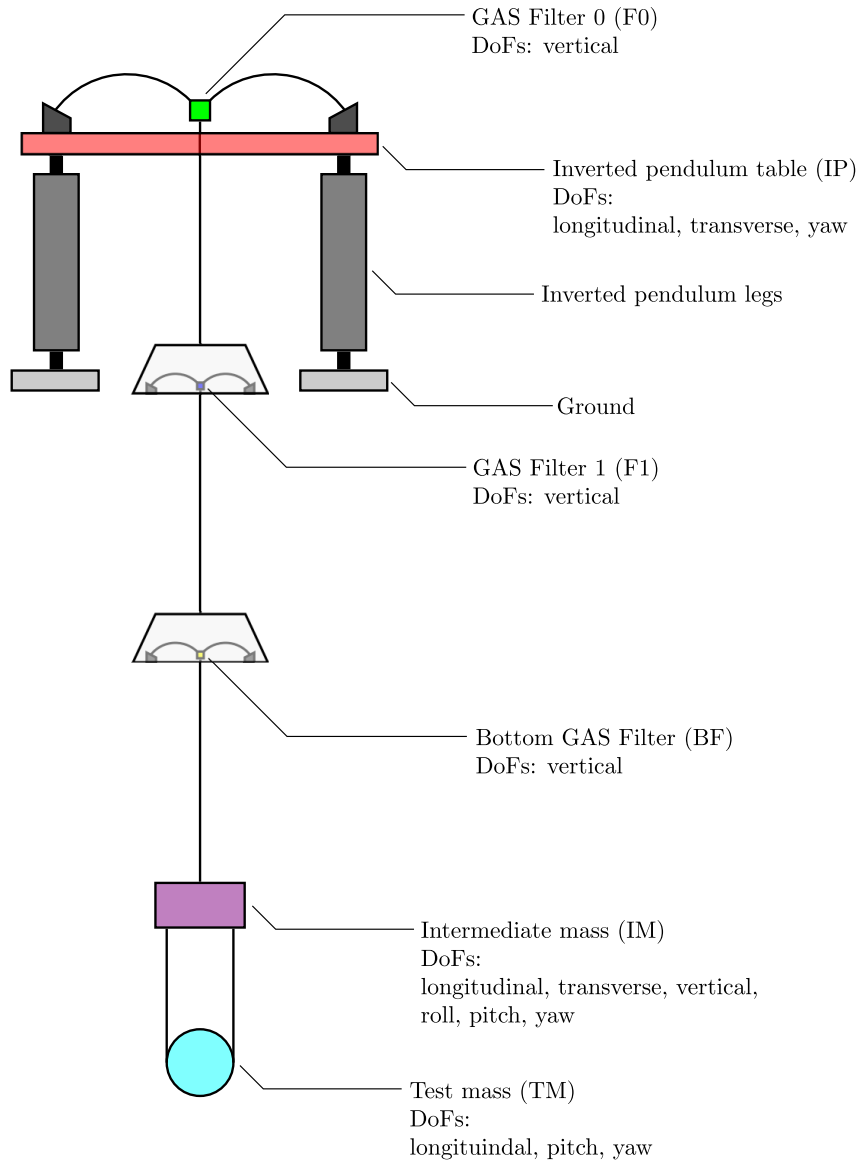


Figure 5.3: A simplified diagram of a type-B suspension (Not to scale). All sensed objects are colored and labeled: Inverted pendulum table (red), GAS filter 0 keystone (green), GAS filter 1 keystone (blue), bottom GAS filter keystone (yellow), intermediate mass (purple), and test mass (cyan).

controller which can be designed independently. To illustrate this, consider a pendulum system as shown in Fig. 5.4. In Fig. 5.4, a pendulum is hanged from a platform and the control degree of freedom is the horizontal position of the pendulum, assuming small angles. The platform is not fixed in space and can be perturbed by random disturbances. As a result, the disturbance propagates and the pendulum starts to move. This is a general description a free-swing condition of any stages in a suspension. The pendulum analogize the stage while the platform analogize the stage above it. Some examples would be the inverted pendulum table (the pendulum) and the ground (the platform), or the test mass (the pendulum) and the intermediate mass (the platform).

The horizontal displacement of the pendulum is measured by a sensor as shown in

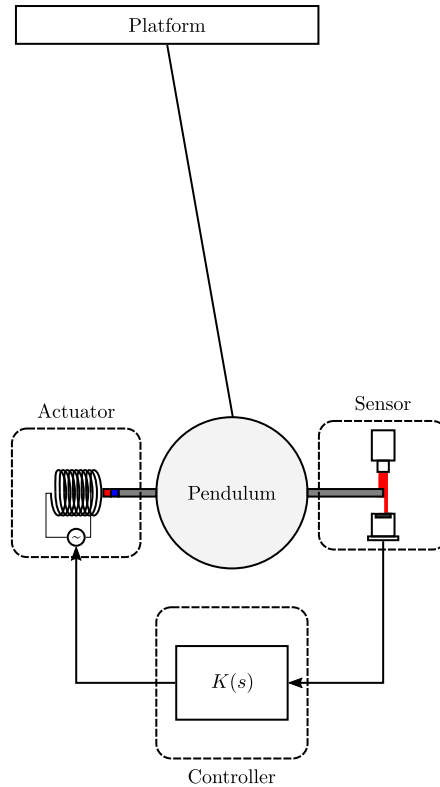


Figure 5.4: An actively controlled pendulum.

Fig. 5.4. The sensor signal is taken by a controller, which outputs an actuation signal. The actuation signal drives an actuator, which applies force on the pendulum. Under this configuration, the idea is to design a controller $K(s)$ such that it commands the actuator to apply a force in the opposite direction of the disturbed motion and hopefully attenuates and cancels it. This holds the pendulum at a desired position, effectively rejecting the disturbances coming from the motion of the platform. And, this is conceptually what needs to be achieved with active isolation¹.

To formally describe the problem, consider a control block diagram shown in Fig. 5.5. Fig. 5.5 shows a general control system with three exogenous inputs $R(s)$, $D(s)$, $N(s)$ and one output $X(s)$. They represent the reference (setpoint), disturbance, noise, and displacement, respectively. Here, s is the complex variables and all quantities are described in the Laplace domain. They blocks labeled $K(s)$ and $P(s)$ denotes the control filter and the plant. And, the variable $U(s)$ represents an actuation command.

In the block diagram, the system has a disturbance $D(s)$ that causes the displacement $X(s)$ to be non-zero. The displacement $X(s)$ is sensed by a noisy sensor with sensing noise $N(s)$, which gives a sensor readout $X(s)+N(s)$. The sensor readout is negatively feedback and the reference $R(s)$ is subtracted from it to give an error signal $R(s) - X(s) - N(s)$. The error signal is then filtered by a controller $K(s)$ that outputs an actuation command $U(s) = K(s)[R(s) - X(s) - N(s)]$ that goes into the plant $P(s)$. The output of the plant

¹Some would use the word damping or simply control.




Figure 5.5: Control block diagram of one degree of freedom

$P(s)U(s)$ summed with the disturbance $D(s)$ gives back the displacement $X(s)$. This gives the signature equation of the displacement:

$$X(s) = \frac{K(s)P(s)}{1 + K(s)P(s)}R(s) + \frac{1}{1 + K(s)P(s)}D(s) - \frac{K(s)P(s)}{1 + K(s)P(s)}N(s). \quad (5.1)$$

The variables names state the nature of the variables but it is worth properly defining them and give further elaboration to avoid confusion and ambiguity. The reference $R(s)$ is the setpoint of system. It represents the desired value of the displacement $X(s)$. For local suspension control, this is usually a value that is determined by the coarse alignment of the interferometer. But, it is not necessarily a predetermined value and it can even be interconnected to other systems. For example, the yaw reference of the inverted pendulum can be interconnected to the yaw readout at the test mass stage. With position control engaged, this effectively offloads the actuation signals from the payload to the inverted pendulum stage. Another example would be the yaw reference of the test mass being interconnected to the interferometer alignment-sensing yaw signal.

The disturbance $D(s)$ is defined as the displacement $X(s)$ when the controller is $K(s) = 0$ (or equivalently, $U(s) = 0$), i.e. the free-swing displacement. For example, the disturbance at the inverted pendulum stage $D_{\text{IP}}(s)$ is the pendulum attenuated ground motion, $D_{\text{IP}}(s) = P_{X_g \rightarrow X_{\text{IP}}}(s)X_g(s)$, where $P_{X_g \rightarrow X_{\text{IP}}}(s)$ is the transfer function from ground displacement $X_g(s)$ to the displacement of the inverted pendulum table. On the other hand, the noise $N(s)$ is also defined as the ‘‘sensing noise’’, which is slightly different from the ‘‘sensor noise’’. Sensing noise is defined as any unwanted signal that is present in the sensor readout, i.e. signal that is not $X(s)$, whereas sensor noise refers to the intrinsic sensor noise. Sensing noise and sensor noise can be the same but are generally not the same. Take the sensing of the inverted pendulum table as an example. The longitudinal sensor at the inverted pendulum stages measures the relative displacement the inverted pendulum table and the ground. So the readout is $X_{\text{IP}}(s) - X_g(s) + N_{\text{sensor}}(s)$, where $X_{\text{IP}}(s)$ is the displacement, $X_g(s)$ is the ground displacement, and $N_{\text{sensor}}(s)$ is the

sensor noise. In this case, the sensing noise is $N(s) = -X_g(s) + N_{\text{sensor}}(s) \neq N_{\text{sensor}}(s)$. Alternatively, the noise $N(s)$ is defined as the negative displacement $-X(s)$ as the controller $K(s) \rightarrow \infty$, i.e. when the control gain is high, assuming that $R(s)$ is a static value at DC and $N(s)$ is a stochastic process with zero mean. This can be easily seen from Eqn. (5.1).

The plant $P(s)$ is also known as the actuation path, or the transfer function from the actuation signal $U(s)$ to the displacement $X(s)$. It refers to the ratio between the displacement and the displacement and the actuation signal. This plant is physical and represents the dynamics of the system. The transfer function of the plant can be derived from the equation of motion. But for a multiple pendulum suspension, most of the time, the frequency response of the plant is measured by, again, taking the ratio between the sensor readout and the actuation signal, which are both measurable. In this case, the actuation signal has to be large enough to the actuated displacement is much greater than the disturbance, and is also might greater than the sensing noise. The transfer function of the plant is then modeled from the measured frequency response.

The controller $K(s)$ is sometimes called a regulator, or simply the control gain. It can be represented by a generic transfer function in polynomial form,

$$K(s) = \frac{b_0 + b_1s + b_2s^2 \dots}{a_0 + a_1s + a_2s^2 \dots}, \quad (5.2)$$

where $a_0, a_1, a_2, b_0, b_1, b_2, \dots$ are scalar coefficients. Another common expression of controller would be the zero-pole-gain model (ZPK model),

$$K(s) = \frac{(s - z_1)(s - z_2)(s - z_3) \dots}{(s - p_1)(s - p_2)(s - p_3) \dots}, \quad (5.3)$$

where z_1, z_2, z_3 are the zeros and p_1, p_2, p_3 are the poles of the controller transfer function. The ZPK expression is often preferred for manual shaping of the controller since the corner frequencies are directly determined by the zeros and poles. The feedback controller is a design freedom of the control system as can take almost any arbitrary number of design parameters, i.e. it can have any arbitrary number of zeros and poles so long as the system is stable.

The reference term is assumed to be a constant at DC and it simply acts as a static offset to the displacement whereas the disturbance $D(s)$ and noise $N(s)$ correspond to two frequency-dependent stochastic processes. Both the disturbance and the noise contribute to the residual motion of the system and they need to be attenuated. To formulate the control design problem, consider a simplified version of Eqn. (5.1) with the reference term stripped,

$$X(s) = \frac{1}{1 + K(s)P(s)}D(s) - \frac{K(s)P(s)}{1 + K(s)P(s)}N(s). \quad (5.4)$$

Eqn. (5.4) shows that the displacement $X(s)$ is a linear combination of a disturbance term $1/[1+K(s)P(s)]D(s)$ (filtered disturbance) and a noise term $K(s)P(s)/[1+K(s)P(s)]N(s)$ (filtered noise). Assuming that the disturbance and noise two uncorrelated, the power spectral density (PSD) of is displacement is simply the sum of the PSD of the disturbance term and the PSD of the noise term. Therefore, the amplitude spectral density (ASD) of the displacement is a quadrature sum of the term filtered terms and it reads

$$\hat{X}(f) = \left[\left| \frac{1}{1 + K(s)P(s)} \right|^2 \hat{D}(f)^2 + \left| \frac{K(s)P(s)}{1 + K(s)P(s)} \right|^2 \hat{N}(f)^2 \right]^{\frac{1}{2}}, \quad (5.5)$$

where f is the frequency (Hz), $s = j\omega = j2\pi f$, j is the imaginary number, $\hat{X}(f)$ is the ASD of the displacement, $\hat{D}(f)$ is the ASD of the disturbance, and $\hat{N}(f)$ is the ASD of the noise.

The displacement $\hat{X}(f)$ is the quantity that needs to be minimized (in some sense). From Eqn. (5.4), it is clear that when $K(s)P(s)$ is large, the disturbance term becomes zero as $1/[1 + K(s)P(s)] \rightarrow 0$. This is the idea of active isolation. However, when $K(s)P(s)$ is large, the coupling term of the noise $\hat{N}(f)$ becomes $K(s)P(s)/[1 + K(s)P(s)] \rightarrow 1$, which fully couples the the noise to the displacement. The two coupling terms² are complementary, i.e.

$$\frac{1}{1 + K(s)P(s)} + \frac{K(s)P(s)}{1 + K(s)P(s)} = 1. \quad (5.6)$$

Therefore, this necessarily means that the disturbance term and the noise term cannot be minimized simultaneously. The problem here is clear. Qualitatively, that is, to design a controller $K(s)$ such that $K(s)P(s)$ is large when $\hat{D}(f) > \hat{N}(f)$, and is small when $\hat{N}(f) > \hat{D}(f)$. And, it has to be done in a way that the system is stable, i.e. $1 + K(s)P(s)$ has no zero on the right-half the complex plane In such a way, the displacement $\hat{X}(f)$, under the influence of a feedback control system, will be dominated by $\min \{\hat{D}(f), \hat{N}(f)\}$, which is the minimum amount of displacement level.

The minimum between the disturbance $\hat{D}(f)$ and noise $\hat{N}(f)$ denote the lower bound of the displacement $\hat{X}(f)$. While the actual control performance is realized by the design of the controller $K(s)$, $\hat{D}(f)$ and $\hat{N}(f)$ determines the potential of the control performance. Therefore, to maximize the possibility of minimizing $\hat{X}(f)$, it is important that the disturbance and noise are minimized as well. The source of disturbance in a suspension system ultimately comes from motion of the ground, which is hard to attenuate. But, recall the pendulum and platform analogy in Fig. 5.4, for stages hanging from the pre-isolator, the source of disturbance can be treated as the displacement of the upper

²The term $1/[1 + K(s)P(s)]$ is called the ‘‘sensitivity’’ and $K(s)P(s)/[1 + K(s)P(s)]$ is called the ‘‘complementary sensitivity’’ in control theory. One famous phenomenon related to the sensitivity function is the waterbed effect, which states that suppression at some frequencies is necessarily coupled with amplification at other frequencies. This also poses an interesting challenge in controller design, where the amplification needs to be placed at frequencies that is not important.

stage, which is minimized by solving control problem that is being discussed. This means that the control of an upper stage is important as the residual motion gets propagated to lower stages. This also determines the priority of control problems: Control is more important at upper stages and they need to be solved first.

As for the sensing noise, for most control systems, its source is ultimately the intrinsic noise of the sensor used to measure the displacement. The intrinsic noise is determined by the hardware and the technology being used to achieve sensing. The noise performance of a sensor can only be improved by advancement in sensing technology and this is currently an active research area. Once the hardware is determined, there is no in situ treatment for reducing the intrinsic sensor noise. However, in some cases, the sensing noise is possible to be minimized. In particular, there are two control schemes at the inverted pendulum stage that can be used to reduce sensing noise: 1) Sensor fusion using complementary filters, and 2) sensor correction. Details of the two approaches are given in Sec. 7.1 and Sec. 7.2. Both control methods assumes that the sensors measure the same quantity and they are inter-calibrated. The method of inter-calibration is given in Sec. 6.1.3.

Nevertheless, there are several regression tasks that need to be done before the controller can be designed. As discussed earlier, the control performance is not solely determined by the controller $K(s)$, but rather, the product of the controller and the plant $K(s)P(s)$, i.e. the open loop transfer function (OLTF). Therefore, in order to analyze control performance before deployment of the controller, the plant $P(s)$ needs to be modeled and this is covered in Sec. 6.3.1. If the control performance were to be analyzed or to be used during the controller design, the ASD of the disturbance $\hat{D}(f)$ and noise $\hat{N}(f)$ also need to be modeled. Details of spectrum modeling is given in Sec. 6.3.2. At last, Sec. 6.4 discusses a baseline method to design the controller $K(s)$ for critical damping and alignment. The whole Chap. 6 is dedicated to discuss methods that have been used during the commissioning of the suspensions at KAGRA. Advanced methods using \mathcal{H}_∞ synthesis is discussed in part. III.

Chapter 6

Suspension Commissioning Techniques

This chapter presents an essential set of techniques that are used to prepare the local control systems for the active isolation systems at KAGRA. Throughout the chapter, complementary examples (real or simulated) are given wherever possible.

The methods discussed in this chapter are developed particularly to tackle tasks discussed in Chap. 5. However, they are not the only way to tackle these tasks and they may not be the best way. But, they are well defined, simple, easy to use, and robust. Most importantly, these methods have already been tested extensively and used during the commissioning phase of O4. Therefore, they serve as fallback methods for suspension commissioning and can be considered the baseline for further developments.

This chapter is a long one containing 4 sections, each corresponds to solutions to the discrete problems that needed to be tackled for commissioning an active isolation system. Sec. 6.1 discusses the calibration of sensors including relative displacement sensors, such as the LVDT, and inertial sensors, such as the geophone. Sec. 6.2 discusses the alignment of the sensors and actuators to a user-defined control basis using sensing matrices and actuation matrices. Sec. 6.3 discusses the modeling of frequency response measurements with transfer functions and the modeling of frequency spectrum. At last, Sec. 6.4 discusses the design of the feedback controller for active damping and coarse alignment control of the optics. While each section can be read without the knowledge from an earlier section, it is recommended that the actual tasks are tackled in the order that these sections are presented. This is because the tasks are not entirely independent. For example, the design of a feedback controller requires a model of the frequency response, and the measurement of a frequency response requires proper sensors and actuators alignment. Therefore, it is recommended to proceed the suspension setup according to the order that these tasks are mentioned in this chapter.

6.1 Sensor calibration

In this section, the calibration of displacement sensors and inertial sensors is covered. The goal of sensor calibration is to obtain a calibration factor or filter that converts the raw sensor readout, which is a voltage converted to a digital value via the analog-to-digital converter (ADC), to a value in displacement unit such as micrometer μm . The calibration factor/filter is implemented in the input filters at the very first stage of the signal processing flow as discussed in Sec. 5.

In general, there are two types of calibration functions, frequency-dependent and frequency-independent. Displacement sensors, such as linear variable differential transformers (LVDT) and OSEMs, are typically frequency-independent, meaning that they map the measured displacement to a voltage. However, the mapping is not a linear relationship. The displacement-voltage curve of a sensor usually follows a sigmoidal trend, i.e. an “S” shape. The relationship is only approximately linear around the center of the curve, i.e. the operating point. For this type of calibration, the common treatment is to approximate a linear relationship around the center and the sensor is said to have a linear range. Sec. 6.1.1 gives two methods to achieve this kind of calibration.

On the other hand, inertial sensors, such as geophones and accelerometer, are typically frequency-dependent since the readout is typically the velocity of a suspended proof mass. The proof mass is suspended from a structure attached to the object that the sensor is sensing. So, the inertial sensors measure inertial motion of an object in direction and a filter is needed to compensate for the dynamics of the suspended proof mass. The calibration filter is like any control filter and it has frequency response correspond to the inverse dynamics of the sensor. Sec. 6.1.2 provides a derivation of the geophone calibration filter and gives an example on calibrating such sensors.

It is worth mentioning that the control performance of an active isolation system is not depended on the quality of sensor calibration. Therefore, sensor calibration needs not to be perfect, except for frequency-dependent calibration, which could introduce unwanted frequency response (phase distortion or amplification) to the plant. However, there are several benefits of having calibrated sensors. First of all, calibrated sensors provide a good reference of the control performances, such as the residual motion in terms of ASD or root-mean-square value (RMS). This enables comparison between different active isolated systems. Having sensors calibrated also means that the sensing noise levels are calibrated. This can aid controller design where noise level is considered.

6.1.1 Sigmoidal calibration curve

All displacement sensors, except inertial sensors, in KAGRA active isolation systems have a sigmoidal calibration curve, including LVDTs, OSEMs, photosensors, and optical levers. While their calibration curves follow a general sigmoidal curve, there are subtle differences

between different types of sensors. And hence, there are more than one way to calibrate them. As for which method is the best, there is no consent in KAGRA. However, all methods will be discussed along with their pros and cons.

The calibration problem can be generalized into an optimization problem:

$$\min_{\theta} \left(\frac{1}{N} \sum_{n=1}^N [f(x_n; \theta) - y_n]^2 \right), \quad (6.1)$$

where $f(x_n; \theta)$ is a model of the calibration curve, x_n is the measured independent variable, i.e. displacement, y_n is the dependent variable, i.e. the sensor readout, θ is the parameter vector being optimized, and N is the number of measurements. Eqn. (6.1) represents a simple least square regression problem, which is straightforward to solve. Unless otherwise specified, the problems in this section are solved using the `Scipy` optimization function `scipy.optimize.minimize()` with default options and $\theta = \vec{0}$ as an initial guess. It is worth noting that the cost function needs not to be a 2-norm of the difference between the calibration function $f(x)$ and the measurements y_n . The 2-norm is a popular choice for regression but is susceptible to biased due to outliers. In case of outliers, 1-norm, i.e. absolute value, can be used to replace the 2-norm in Eqn. (6.1) for fitting.

Here, two simple models have been used as $f(x; \theta)$ for calibrating displacements sensors at KAGRA. They are

1. Linear function, and
2. Error function.

The choice of these functions is physics informed. Linear function is the candidate for LVDTs since the electronic readout of an LVDT is proportional to the displacement being measured. The error function models calibration curves of sensors with optical readout, especially those with a Gaussian beam. Hence, the error function works well with OSEMs, photosensors, and optical levers.

Linear function

Linear variable differential transformers (LVDTs) are sensors that map displacement linearly to output voltage, as is indicated by the name. Therefore, a linear function

$$f(x; m, c) = mx + c, \quad (6.2)$$

where m is the slope and c is the y-intercept, is supposedly a perfect candidate to be used in Eqn. (6.1) as a calibration function for LVDTs. However, there are some caveats, which will be pointed out with an example below.

Fig. 6.1 shows a example calibration of one of the LVDTs at the inverted pendulum stage of the SR3 suspension. The measurements was done on site with the following




Figure 6.1: SR3 IP H1 LVDT calibration plot. Blue dots: Measurement data. Orange “×”: Masked data set with readouts between $(-16384, 16384)$. Orange solid line: Calibration curve fit to the half-range readouts. Green “+”: Data set found by Algorithm 1 corresponding to 1% non-linearity. Green dashed line: Calibration curve fit to the 1% non-linearity data set.

procedure. The inverted pendulum was rotated manually in yaw (with hands) at regular space steps. At each step, all three LVDTs were measured at the same time with three calipers and the digital readings were taken as well. For each LVDT, the caliper was used to measure the distance between a metal bracket (moving with the LVDT) and the security structure (fixed on the ground). The constant rotation intervals are defined such that the calipers read 1 millimeter intervals so Fig. 6.1 shows the measurement data with the caliper reading from 4 mm to 20 mm.

The measurement was terminated as soon as non-linearity had clearly been observed. And, it is clear that the raw readout (blue dots) shown in Fig. 6.1 exhibits non-linear relationship with measured displacement around the boundary and the data points look sigmoidal rather than a straight line. This also indicates that the LVDT doesn’t work linearly for the full 16-bit ADC range, i.e. readout from -32768 to 32767 counts. Therefore, the LVDT only works within a smaller range (linear range) that is often hard to define. Not all data points are used to fit the calibration curve. But, as for which data points are used, this is often not well-defined at KAGRA and they are chosen by experience.

Here, two ways are proposed to mitigate the inconsistency problem arising from manually chosen data points,

1. Define a operating range and only use data points within the range, and
2. Specify an acceptable non-linearity and use Algorithm 1 to obtain a calibration

curve.

The first approach is simple. Instead of having data points chosen from experience, simply define a usable range for the sensor and only use data points within the operating range for fitting a calibration curve. Fig. 6.1 shows an example of this approach. Orange crosses are data points with readout in between $(-16384, 16384)$, which is the half-range of the 16-bit ADC. In this case, the range is chosen for demonstration purpose only although the choice of a half-range is usually sufficient for defining the linear range in practice. The green solid line in Fig. 6.1 is a least-square fit of the orange data set and it gives a slope of 3852.5 counts/mm. Assuming that the desired displacement unit is μm , the calibration factor that converts the sensor readout to μm is then $1/3852.5 \text{ mm/counts} \times 1000 \mu\text{m}/1 \text{ mm} = 0.2596 \mu\text{m/counts}$.

As for the second approach, Algorithm 1 shows the pseudocode corresponding to the approach. First of all, an acceptable non-linearity needs to be specified first. The non-linearity in sensors is defined as the readout deviation from the ideal readout as a percentage of the full scale:

$$\text{non-linearity} = \frac{\text{readout} - \text{ideal readout}}{\text{full scale}}, \quad (6.3)$$

where the full scale is defined as the difference between the maximum readout and the minimum readout. In the case of a 16-bit signed ADC readout, the full scale is $32767 + 32768 = 65535$ counts. Then, the data point closest to the midrange of the readout is put into an initial fitting data set along with the two neighboring points. These three points are at the middle of the calibration curve. They are assumed to be within the “most linear” part of the calibration curve and therefore they are chosen for an initial least-square fit. This approach automatically fails if this is not true. Alternatively, users can specify a starting point where the calibration curve is considered to be the center of the linear regime. After obtaining a linear function from the three points, the function is used to evaluate the non-linearity of other data points from the original data set. If the non-linearity of a data point falls within specification, e.g. within 1% non-linearity, then the data point is included into the data set for fitting. This process repeats until no data point is added to the fitting data set. One advantage of this is that the non-linear range is automatically defined by the data points within the selected data set, which is defined by specified non-linearity. As for what value should the non-linearity specification be taken, this remains as an unanswered question and further study is needed to evaluate the non-linearity effect in control.

Using Algorithm 1, the data within 1% non-linearity is found and plotted as green “+” in Fig. 6.1. The green dashed line shows the least-square fit to the data. The corresponding calibration factor is $1/3813.6 \text{ mm/counts} \times 1000 \mu\text{m}/1 \text{ mm} = 0.2622 \mu\text{m/counts}$. This happens to be exactly the same as the calibration factor used in KAGRA [45], although

Algorithm 1 Find and calibrate data points within non-linearity specification

$y_n \leftarrow \{y_1, y_2, y_3, \dots, y_N\}$ ▷ The full readout data set (sorted).
 $x_n \leftarrow \{x_1, x_2, x_3, \dots, x_N\}$ ▷ The full displacement data set (sorted).
 $y_{\text{midrange}} \leftarrow \frac{y_{\text{max}} + y_{\text{min}}}{2}$ ▷ Midrange, i.e. the center point of the readout, e.g. -0.5.
 $y_m \leftarrow \arg \min_{y_i \in y_n} |y_i - y_{\text{midrange}}|$ ▷ The readout data point closest to the midrange.
 $y_{\text{fit}} \leftarrow \{y_{m-1}, y_m, y_{m+1}\}$ ▷ Choose the data around the midrange for initial fitting.
 $x_{\text{fit}} \leftarrow \{x_{m-1}, x_m, x_{m+1}\}$
 $y_{\text{fullrange}} \leftarrow y_{\text{max}} - y_{\text{min}}$ ▷ Full scale of the readout, e.g. $32767 - (-32768) = 65535$.

while True **do**

$f(x) = \text{linear_fit}(x_{\text{fit}}, y_{\text{fit}})$ ▷ Fit the data and obtain a linear function $f(x)$.

for i in $\{1, 2, 3, \dots, N\}$ and $y_i \notin y_{\text{fit}}$ **do** ▷ For data that is not in y_{fit} .

 non-linearity = $\frac{y_i - f(x_i)}{y_{\text{fullrange}}}$

if non-linearity < specified non-linearity **then**

$y_{\text{fit}} \leftarrow y_{\text{fit}} \cup \{y_i\}$ ▷ Add the data into y_{fit} .

$x_{\text{fit}} \leftarrow x_{\text{fit}} \cup \{x_i\}$

end if

end for

if the y_{fit} data set didn't change **then**

 break ▷ Terminate and return the slope and intercept of $f(x)$.

end if

end while

the data points used in KAGRA were manually selected. Nevertheless, the two methods happens to result in around 1% difference in calibration factor, which is practically insignificant.

Error function

The choice of the error function is physics driven and, in KAGRA, it is widely use as an intermediate calibration function for optical levers, which are the sensors that measure local displacements of the suspended optics. Consider a Gaussian beam, the intensity of the beam at a distance z from the beam waist is given by a Gaussian profile

$$I(r, z) = k(z) \exp\left(\frac{-2r^2}{w(z)^2}\right), \quad (6.4)$$

where $\exp(\cdot)$ is the exponential function, z is the axial distance from the beam waist, $k(z)$ and $w(z)$ are some functions depending on z , and r is the radial distance from the center axis of the beam. Suppose the beam is measured at fixed distance z_0 from the waist and $k(z_0) = a$ and $\frac{2}{w(z_0)^2} = b$, where a and b are constants, the intensity of the beam, in terms of Cartesian coordinates, is

$$I(x, y) = a \exp(-b(x^2 + y^2)). \quad (6.5)$$

If the beam (symmetrically, the sensor) is further constrained to move along the x -axis, the intensity can be further simplified to

$$I(x) = c \exp(-bx^2), \quad (6.6)$$

where $c \equiv a \exp(-by^2)$ is a constant.

The sensor of the optical lever in KAGRA is a quadrant photodiode (QPD), which is composed of four photodiodes measuring a beam in four quadrants. For simplicity, instead of four photodiodes, consider a sensor consist of two infinitely large photodiodes, corresponding to one on the negative x plane and the other on the the positive x plane. The slit between the two photodiodes now sits at $x = 0$ and the two photodiodes now read an equal amount of light power. Now, denote the light power on the right plane by P_+ and the light power on the left plane by P_- and supposed the beam is shifted by in the x -direction by an amount of x' (or equivalently, the slit shifted by an amount of $-x'$). The light power on the right photodiode reads

$$P_+(x') = \int_{-\infty}^{\infty} \int_{-x'}^{\infty} I(x, y) dx dy = a \int_{-\infty}^{\infty} \int_{-x'}^{\infty} \exp(-b(x^2 + y^2)) dx dy. \quad (6.7)$$

The y -integral gives

$$\int_{-\infty}^{\infty} \exp(-by^2) dy = \frac{1}{\sqrt{b}} \int_{-\infty}^{\infty} \exp\left(-\left(y\sqrt{b}\right)^2\right) d(y\sqrt{b}) \quad (6.8)$$

Here, the Gaussian integral is simply

$$\int_{-\infty}^{\infty} \exp(-t^2) dt = \sqrt{\pi}. \quad (6.9)$$

The y -integral then becomes

$$\frac{1}{\sqrt{b}} \int_{-\infty}^{\infty} \exp\left(-\left(y\sqrt{b}\right)^2\right) d(y\sqrt{b}) = \sqrt{\frac{\pi}{b}}. \quad (6.10)$$

Substituting this into Eqn. (6.7) gives,

$$\begin{aligned} P_+(x') &= \frac{a\sqrt{\pi}}{\sqrt{b}} \int_{-x'}^{\infty} \exp(-bx^2) dx \\ &= a\sqrt{\frac{\pi}{b}} \int_{-x'}^0 \exp(-bx^2) dx + a\sqrt{\frac{\pi}{b}} \int_0^{\infty} \exp(-bx^2) dx \\ &= a\sqrt{\frac{\pi}{b}} \int_{-x'}^0 \exp(-bx^2) dx + \frac{a\sqrt{\pi}}{2b} \int_{-\infty}^{\infty} \exp\left(-\left(x\sqrt{b}\right)^2\right) d(x\sqrt{b}) \\ &= -\frac{a\sqrt{\pi}}{b} \int_0^{-x'\sqrt{b}} \exp\left(-\left(x\sqrt{b}\right)^2\right) d(x\sqrt{b}) + \frac{a\pi}{2b}. \end{aligned} \quad (6.11)$$

Now, the error function $\text{erf}(t)$ is defined as

$$\text{erf}(z) \equiv \frac{2}{\sqrt{\pi}} \int_0^z \exp(-t^2) dt. \quad (6.12)$$

The light power measured by the right photodiode can immediately be seen as

$$\begin{aligned} P_x(x') &= -\frac{a\pi}{2b} \text{erf}(-x'\sqrt{b}) + \frac{a\pi}{2b} \\ &= \frac{a\pi}{2b} \text{erf}(x'\sqrt{b}) + \frac{a\pi}{2b}, \end{aligned} \quad (6.13)$$

where the property that $\text{erf}(z)$ is an odd function is used in the last line.

Now, Eqn. (6.13) shows the light power measured by a photodiode in a shadow sensor, i.e. sensors measuring the motion of a flag blocking an incidence beam. This shows that the calibration function for shadow sensors, like an OSEM, can be perfectly modeled by an error function. For optical levers, the sensor readout is a normalized differential readout between the right and the left photodiode, $(P_+ - P_-)/(P_+ + P_-)$. Here, the light power on the left sensor is

$$\begin{aligned} P_-(x') &= a\sqrt{\frac{\pi}{b}} \int_{-\infty}^{-x'} \exp(-bx^2) dx \\ &= a\sqrt{\frac{\pi}{b}} \int_{-\infty}^0 \exp(-bx^2) dx - a\sqrt{\frac{\pi}{b}} \int_{-x'}^0 \exp(-bx^2) dx \\ &= \frac{a\pi}{2b} - \frac{a\pi}{2b} \text{erf}(x'\sqrt{b}). \end{aligned} \quad (6.14)$$

The differential of the two light power gives

$$P_+(x') - P_-(x') = \frac{a\pi}{b} \text{erf}(x'\sqrt{b}). \quad (6.15)$$

The output of the QPD (and other optical sensors) is a number proportional to the differential light power (light power). Redefining the constants, a and b , this makes

$$f_{\text{erf}}(x; a, b, c, d) = a \text{erf}(b(x - c)) + d \quad (6.16)$$

a perfect candidate as an calibration function for QPD, OSEMs, and photosensors, where c and d are additional constants introduced for generality, and x here is the displacement of the beam spot.

Having Eqn. (6.16) alone is not sufficient since the digital system only allows an implementation of a calibration factor, not a function. Using the Taylor expansion, a small distance x' around $x = c$ is

$$f_{\text{erf}}(x'; a, b, c, d) = d + x' \left[\frac{d}{dx} a \text{erf}(b(x - c)) \right]_{x=c} + O(x^2), \quad (6.17)$$

where $O(\cdot)$ is the big O notation. The derivative of the error function is simply

$$\frac{d}{dx} \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \exp(-x^2). \quad (6.18)$$

The Taylor expansion then becomes

$$\begin{aligned} f_{\operatorname{erf}}(x'; a, b, c, d) &= d + x' \frac{2ab}{\sqrt{\pi}} \left[\exp(-(b(x-c))^2) \right]_{x=c} + O(x^2) \\ &= d + \frac{2ab}{\sqrt{\pi}} x' + O(x^2). \end{aligned} \quad (6.19)$$

This gives a calibration factor of $\sqrt{\pi}/(2ab)$, which converts the raw sensor readout to the measured displacement around $x = c$.

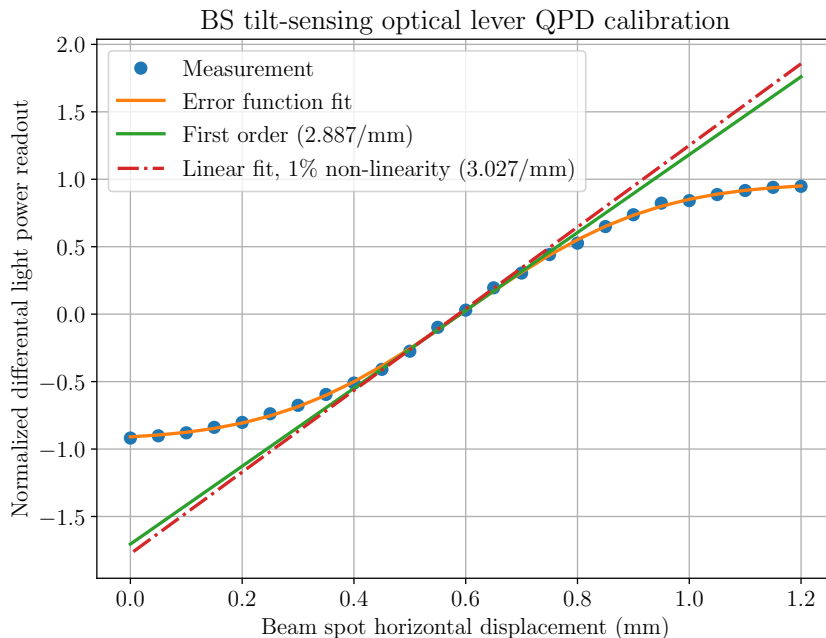


Figure 6.2: Calibration of the BS test mass optical lever tilt-sensing QPD. Blue dots: Measurement. Orange solid line: Error function fit. Green solid line: First order approximation of the error function fit. Red dash-dot line: Linear function fit with 1% non-linearity.

As an example, Fig. 6.2 shows the calibration of the beamsplitter (BS) tilt-sensing optical lever. The corresponding sensor is a QPD sensing a superluminescent diode light beam reflecting off the BS, so the beam spot displacement reflects the angular displacement of the optics. Data points shown in Fig. 6.2 were measured by sliding the QPD horizontally with a micrometer stage in 0.05 mm steps. This simulates incremental displacements of the optical lever beam.

As can be seen, the error function fits very well to the measured data in Fig. 6.2. The fitted error function was obtained by solving the optimization problem described by Eqn. (6.1) using Eqn. (6.16) as the calibration function. The minimization problem should

work with $\vec{0}$ as a starting point. A reasonable guess would be $\{a, b, c, d\} = \{(\max(y_n) - \min(y_n))/2, (y_N - y_1)/(x_N - x_1), x_m, y_m\}$, where y_n is the readout data, x_n is the beam spot displacement data, $n = 1, 2, 3, \dots, N$, N is the number of data, and (x_m, y_m) is the data point at midrange. In practice, it is also found that normalizing y_n and x_n with $\max(|y_n|)$ and $\max(|x_n|)$, respectively, helps with numerical stability.

The same data set is also fitted using the aforementioned non-linearity method as described in Algorithm. 1. As shown in Fig. 6.2, the first order approximation of the fitted error function gives a calibration factor of $1/2.887 = 0.3464$ mm, whereas a linear fit, assuming 1% non-linearity, gives a calibration factor of $1/3.027 = 0.3304$ mm. In comparison, the two methods resulted in an 4.6% difference in calibration factor, which is, again, negligible in practice.

6.1.2 Geophone calibration

Unlike relative displacement sensors such as LVDTs and OSEMs, the readouts from inertial sensors is an inertial readout, like the name suggested. Typically, the readout is either in the unit of velocity or acceleration so the signal needs to be integrated to obtain a displacement signal for control. Moreover, inertial sensors typically don't have a flat frequency response. This means the signals need to be compensated by a filter corresponding to the inverse of the sensor response. To illustrate this, consider a geophone model shown in Fig. 6.3. In the figure, the geophone is attached to a platform, which has an inertial displacement x . The geophone has a proof mass suspended by a spring attached to the enclosure of the geophone and the proof mass has a displacement x_m . A magnet is attached to the proof mass and it induces electrical current to a coil fixed on the enclosure. The readout is the induced voltage v generated by the moving magnet attached to the proof mass. The goal is to obtain a calibration filter, i.e. a transfer function, that converts the readout to the platform displacement x .

The equation of motion of the proof mass gives

$$m \frac{d^2 x_m}{dt^2} = -b \left(\frac{dx_m}{dt} - \frac{dx}{dt} \right) - k (x_m - x), \quad (6.20)$$

where m is the mass of the proof mass, b is the damping coefficient, and k is the spring constant. Rearranging gives

$$\frac{d^2 x_m}{dt^2} + \frac{b}{m} \frac{dx_m}{dt} + \frac{k}{m} x_m = \frac{b}{m} \frac{dx}{dt} + kx, \quad (6.21)$$

and taking the Laplace transform of the equation yields

$$\left(s^2 + \frac{b}{m}s + \frac{k}{m} \right) X_m(s) = \left(\frac{b}{m}s + \frac{k}{m} \right) X(s), \quad (6.22)$$




Figure 6.3: A internal diagram illustrating the sensing principle of a geophone measuring the motion of a platform.

where s is the Laplace variable, $X_m(s) = \mathcal{L}\{x_m\}$ is the Laplace transform of the proof mass displacement, and $X(s) = \mathcal{L}\{x\}$ is the Laplace transform of the platform displacement. The voltage readout v is proportional to the differential velocity

$$v \propto \left(\frac{dx_m}{dt} - \frac{dx}{dt} \right), \quad (6.23)$$

and the Laplace transform gives

$$V(s) = -G [X_m(s) - X(s)] s, \quad (6.24)$$

where $V(s)$ is the Laplace transform of the voltage readout v and G is the sensitivity¹ of the geophone which has a unit of V/(m/sec). Note that the minus sign is conventional. Now, rearranging Eqn. (6.22) gives

$$\left(s^2 + \frac{b}{m}s + \frac{k}{m} \right) [X_m(s) - X(s)] s = -s^3 X(s), \quad (6.25)$$

and substituting $V(s)$ gives

$$\begin{aligned} \frac{V(s)}{sX(s)} &= \frac{Gs^2}{s^2 + \frac{b}{m}s + \frac{k}{m}} \\ &= \frac{Gs^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \end{aligned} \quad (6.26)$$

where $\omega_n \equiv \sqrt{k/m}$ is the natural angular frequency of the geophone and $\zeta \equiv b/(2m\omega_n)$ is the damping factor. From here, it can be seen that the frequency response of the sensor

¹ G is also called the generator constant of a geophone

is $Gs^2/(s^2 + 2\zeta\omega_n s + \omega_n^2)$, so the calibration filter $H(s)$ is simply the inverse:

$$H(s) \equiv \frac{sX(s)}{V(s)} = \frac{s^2 + 2\zeta\omega_n s + \omega_n^2}{Gs^2}, \quad (6.27)$$

which is a filter that converts the voltage readout to the velocity of the platform. To get the platform displacement, simply integrate Eqn. (6.27) in the Laplace domain, i.e. multiplying Eqn. (6.27) by an integrator $1/s$.

The geophones used at KAGRA are the commercial model L-4C from the company Sercel [46]. The typical parameters of the L-4C geophone is shown in Table 6.1. The corresponding calibration filter is shown in Fig. 6.4.

Sensitivity G $\left(\frac{\text{V}}{\text{m/sec}}\right)$	Damping factor ζ	Natural frequency ω_n (rad/sec)
276.8	0.28	2π (1 Hz)

Table 6.1: Typical calibration parameters of the L-4C geophone [47].

The typical parameters in Table 6.1 serve as an good starting point for the calibration filter but they required some fine tuning for each geophone. However, according to the specifications, the damping factor ζ in Table 6.1 is an open-circuit parameter. This means that the damping factor is measured when the leads of the readout coil are opened, which is not the actual operating condition. The actual damping factor depends on the resistance value of the shunt resistor parallel to the coil. Moreover, the damping factor and resonance frequency of the L-4C geophone can differ from specification and even change over time [48]. Therefore, it is recommended that the geophones are calibrated on site right before installing them to the pre-isolators.

The calibration of a geophone can be done by fixing the it on the ground so it measures the seismic noise. The geophone signal is then compared to that of a seismometer, which is assumed to be more sensitive. Solving the optimization problem

$$\min_{G, \zeta, \omega_n} \left(\frac{1}{N} \sum_{i=1}^N \log \left(\left| H(j\omega_i; G, \zeta, \omega_n) - \frac{\dot{X}_g(j\omega_i)}{V(j\omega_i)} \right|^2 \right) W(\omega_i) \right), \quad (6.28)$$

where N is the number of data point in the frequency domain, ω_i is the frequency array, $V(j\omega_i)$ is the raw geophone readout, $\dot{X}_g(j\omega_i)$ is the ground velocity measured by the seismometer, and $W(\omega_i)$ is a frequency-dependent weighting function. The rationale behind this cost function is that the calibrated readout $H(j\omega)V(j\omega)$ needs to be the same as the ground velocity $\dot{X}_g(j\omega)$ at all frequencies. The reason why $V(j\omega)$ and $\dot{X}_g(j\omega)$ are put into a fraction is because it is more convenient to measure a transfer function in KAGRA diagnostic system (inherited from LIGO), so $\dot{X}_g(j\omega_i)/V(j\omega_i)$ is really one complex frequency series. The logarithm is also used to scale the data values, accounting for a large dynamic range. At last, the weighting function $W(\omega_i)$ serves to weigh or filter




Figure 6.4: Bode plot of the L-4C geophone calibration filter. (a) Magnitude response of the calibration filter. (b) Phase response of the calibration filter.

out part of the data where noise is dominated. It can be simply set as the coherence function between the geophone readout and the seismometer readout. Another choice is to set the weighting function to 0 or 1 depending on whether the coherence function exceeds a certain threshold, e.g. 0.5. It is worth noting that the optimization problem denoted by Eqn. (6.28) corresponds to a transfer function fitting problem, which will appear in later sections.

Example

In this example, the goal is to find a calibration filter following the method above. The parameters for this geophone is $G = 304.48 \text{ V}/(\text{m}/\text{sec})$, $\zeta = 0.308$, and $\omega_n = 2.2\pi \text{ rad}/\text{sec}$,

which are all 10% above the typical values.




Figure 6.5: Amplitude spectral density plots of a simulated seismic noise and geophone readout. Blue: Ground velocity ($\mu\text{m/s}/\sqrt{\text{Hz}}$). Orange: Geophone readout noise ($\mu\text{V}/\sqrt{\text{Hz}}$). Green: Geophone raw readout ($\mu\text{V}/\sqrt{\text{Hz}}$).

Fig. 6.5 shows a simulated result of a seismic noise (Blue) and the geophone readout (Green). The geophone readout is also covered in a white readout noise (Orange) with $0.1 \mu\text{m}/\sqrt{\text{Hz}}$. As can be seen, the geophone readout is dominated by readout noise at low frequency (below 0.07 Hz). Because of this, the data points at those frequencies cannot be used and should be masked out.

Fig. 6.6 shows the coherence between the ground velocity and the geophone readout (Blue). As expected, the coherence between the seismic noise and geophone readout is low at below 0.07 Hz so it can be used to define a weighting function in Eqn. (6.28). The weighting function here is chosen to be 1 when the coherence is greater than 0.5, and 0 when the coherence is lower than 0.5. This weighting function is shown as orange line in Fig. 6.6 and it reasonably masks out the useful data that are not noise dominated.

Fig. 6.7 shows the measured inverse geophone frequency response (blue), the ideal calibration filter (orange), and the calibration filter fitted to the measured response. The measured inverse geophone frequency response is obtained via the ratio between the Laplace transform of the ground velocity and the geophone readout, i.e. $\dot{X}_g(s)/V(s)$ in Eqn. (6.28). The fitted calibration filter was obtained by solving the optimization problem Eqn. (6.28). The initial values for local optimization were obtained from Table 6.1. And, the optimized values are $G = 305.96 \text{ V}/(\text{m}/\text{sec})$, $\zeta = 0.3088$, $\omega_n = 2.203\pi \text{ rad}/\text{sec}$, which are all less than 1% away from the true values ($G = 304.48 \text{ V}/(\text{m}/\text{sec})$, $\zeta = 0.308$, and $\omega_n = 2.2\pi \text{ rad}/\text{sec}$). The corresponding calibration filter is shown as green dashed




Figure 6.6: Coherence plot and weighting function. Blue: Coherence between seismic noise and geophone readout. Orange: Weighting function.




Figure 6.7: Geophone calibration filter. Blue: Measured inverse geophone frequency response. Orange: Ideal calibration filter. Green dashed: Calibration filter fitted to the measured response.

line in Fig. 6.7. As shown in the figure, the overlap between the fitted calibration filter and the ideal calibration filter almost perfect.

6.1.3 Inter-calibration of sensors

The inter-calibration of sensors is required when two sensors signals need to be compared or when they need to be combined using special control strategies such as sensor fusion and sensor correction, which are later discussed in Sec. 7.1 and 7.2. In principle, the geophone calibration method presented in Sec. 6.1.2 is already a method for inter-calibration between the geophone and seismometer and is perhaps a more complex as there are three variables. Generally, the inter-calibration of sensors only has one variable, that is, the relative scale between the sensor readouts. Suppose there are two sensors that read a common signal and the amplitude spectral densities of the two readouts are given by $Y_1(f)$ and $Y_2(f)$ and the second readout $Y_2(f)$ is the one that needs to be inter-calibrated. Since the two sensors read the same signal, their relative phase between the two signals should be zero, which makes the amplitude spectral densities comparable. In this case, the readouts can be related as $Y_1(f) = aY_2(f)$, where a is the inter-calibration factor that needs to be obtained. The inter-calibration factor can simply be obtained by solving the optimization problem

$$\min_a \left(\frac{1}{N} \sum_{i=1}^N \log \left[a \frac{Y_2(f_i)}{Y_1(f_i)} \right] W(f) \right), \quad (6.29)$$

where N is the number of data points, f_i is the frequency data, and $W(f)$ is a weighting function to filter out or weight low-coherence data. The cost function is simply derived from the assumption that $\log aY_2(f) - \log Y_1(f) = 0$, where the logarithm is taken to into account the large dynamic range of data values.

6.2 Control basis alignment

As mentioned in Chap. 5, the sensors and actuators in each stage of the suspension are not aligned in the desired directions. They need to be aligned via control matrices, in particular, sensing matrices and actuation matrices, which appears in each stage of the suspensions. The sensing matrix at each stage takes the calibrated sensor signals as inputs and it outputs displacement signals in the desired control basis. Likewise, the actuation matrix at each stage takes the control actuation signals in the control basis and it outputs actuation signals to individual actuators. The control basis at each stage is defined by the Cartesian and Euler basis as discussed in Sec. 5.1.

The sensing matrix can be decomposed into two types of matrices, geometric sensing matrix and diagonalization sensing matrix. Likewise, the actuation matrix can be decomposed into geometric actuation matrix and diagonalization actuation matrix. The geometric matrices are defined by the placement of the hardware so it can be derived using simple geometry. On the other hand, the diagonalization matrices are matrices that finely decouples any residual couplings between different degrees of freedom. The

diagonalization matrices can only be derived from measurable cross-couplings.

For the Type-B suspensions, the geometric sensing matrices for the pre-isolator and intermediate mass are given in Ref. [29], and the actuation matrices for the pre-isolator, intermediate mass, and test mass are also given. Detailed derivations for these matrices are also given in Refs. [43, 44] in the form of Mathematica notebooks. Type-A and Type-Bp suspensions share similar mechanical structures with the Type-B suspensions. Therefore, the aforementioned matrices for Type-A and Type-Bp suspensions can be derived straightforwardly using the same approaches. The derivation of the geometric sensing matrix at the test mass stage is, however, more involved and is not provided anywhere. This is to be discussed in Sec. 6.2.1, which is a re-elaboration of Ref. [49]. Sec. 6.2.1 also discusses diagonalization of sensing matrix and a real example is also provided. The geometric actuation matrix of test mass stage is revisited as an example in Sec. 6.2.2. Diagonalization of the actuators is also discussed.

6.2.1 Sensing matrices

Sensing matrices are matrices that transform displacements from the sensor basis to the desire control basis. Sensing matrix \mathbf{S} in a stage can be described by

$$\mathbf{y} = \mathbf{S}\mathbf{x}, \quad (6.30)$$

where $\mathbf{x} = (x_1, x_2, x_3, \dots, x_n)^\top$ is the array of sensor readouts $x_1, x_2, x_3, \dots, x_n$, $\mathbf{y} = (y_1, y_2, y_3, \dots, y_m)^\top$ is the array of displacements in the each direction of the control basis, and $\mathbf{S} \in \mathbb{R}^{m \times n}$ is the sensing matrix.

As a quick example, Fig. 6.8 shows the sensor directions and control directions of the pre-isolator. In this case, the sensor readout array \mathbf{x} at the pre-isolator would be the three tangential displacements (at intervals of 120°) along the circumference where the LVDTs are located. And, \mathbf{y} would be the three displacements / angular displacements, longitudinal, transverse, and yaw. Geometrically, the sensing matrix is simply

$$\mathbf{S}_{\text{preisolator}} = \begin{bmatrix} \cos \theta & \sin \theta & R \\ \cos \left(\theta + \frac{2}{3}\pi \right) & \sin \left(\theta + \frac{2}{3}\pi \right) & R \\ \cos \left(\theta + \frac{4}{3}\pi \right) & \sin \left(\theta + \frac{4}{3}\pi \right) & R \end{bmatrix}^{-1}, \quad (6.31)$$

where the first sensor is located at an angle of θ from the negative transverse direction, and R is the distance from the axis of rotation of the inverted pendulum to the sensors.

In practice, the geometric sensing matrix may not be sufficient and the control readouts could contain observable residual cross-coupling between different degrees of freedoms. In this case, the sensing matrix needs to be adjusted according to the cross-coupling measured. However, directly adjusting the geometric sensing matrix may be confusing as

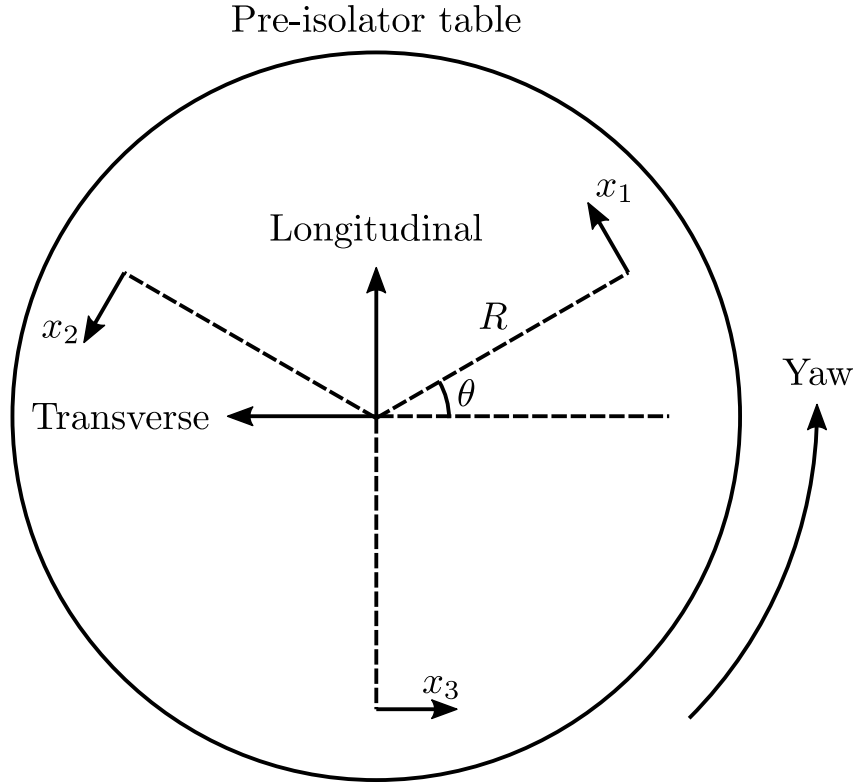


Figure 6.8: Sensor directions and control directions of the pre-isolator table (top view).

it couples geometric information and cross-coupling information in one matrix. Therefore, in reality, the sensing matrix \mathbf{S} is decomposed into two parts, geometric sensing matrix \mathbf{S}_{geo} and diagonalization sensing matrix \mathbf{S}_{diag} ,

$$\mathbf{S} = \mathbf{S}_{\text{diag}} \mathbf{S}_{\text{geo}} . \quad (6.32)$$

Note that the order of operation is important. In this section, the geometric sensing matrix of the test mass is derived and the diagonalization of the sensing matrix is also exemplified.

Sensing matrix of the optical lever

An optical lever is a device that measures small angular displacements of a reflective surface. The word “optical lever” is used in KAGRA to describe the optical displacement sensing setup at the test mass stage, which is composed of an actual optical lever and other devices. The optical displacement sensing setup (optical lever from hereafter) reads longitudinal displacement, pitch, and yaw of the optics. It is composed of a light source, typically a superluminescent diode (SLED), two beam position sensors, typically quadrant photodiodes (QPDs), and some small optics such as lens, mirrors, and beamsplitters.

Consider a simply optical lever setup as shown in Fig. 6.9. The beam strikes the suspended optics at an incidence angle of α and is reflected by the suspended optics. The




Figure 6.9: A rotation in yaw of the optics causes the optical lever beam spot to shift (top view).

suspended optics is rotated in the yaw direction by an amount of θ_Y . At distance r from the suspended optics, the position of the optical lever beam spot is shifted by an amount of $2r\theta_Y$, assuming small angles. A beam position sensor placed here will measure a beam spot displacement of

$$x = 2r\theta_Y, \quad (6.33)$$

which can be used to infer the yaw of the optics, i.e. dividing the displacement readout by $2r$.

Besides rotating in the yaw direction, the suspended optics can also shift in the longitudinal direction. This creates a similar effect compared to that of a yaw rotation in the optics. Fig. 6.10 shows the effect on the optical lever beam spot when the optics is shifted in the longitudinal direction by an amount of x_L . Note that in this case the optical lever is located at the back side of the optics, which is not necessarily the case for all optics. In this case, the beam spot position is shifted by an amount of $(2 \sin \alpha)x_L$ in the same direction that a yaw rotation would cause. Combined with the effect of rotation, a beam position sensor would sense a superposition of yaw and longitudinal displacement

$$x = 2r\theta_Y + (2 \sin \alpha)x_L. \quad (6.34)$$

This means one beam position sensor cannot tell apart yaw and longitudinal displacement.

To remedy this issue, two beam position sensors are used instead and a beamsplitter




Figure 6.10: A longitudinal shift of the optics causes the optical lever beam to shift (top view).

is used to divide the beam into two. There are two options where the other sensor can be placed, 1) at a different distance r' , or 2) behind a lens. In KAGRA, the latter configuration was chosen and it is shown in Fig. 6.11. In this configuration, a lens is placed at r_{lens} away from the suspended optics. The second beam position sensor, reading a beam spot displacement of x_2 , is placed at distance d behind the lens. In this case, the beam spot displacement x_2 can be derived using the ray transfer matrix analysis [40]. Using the ray transfer matrices for beam propagation and thin lens, it gives

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix} \begin{bmatrix} 1 & r_{\text{lens}} \\ 0 & 1 \end{bmatrix} \begin{pmatrix} (2 \sin \alpha) x_L \\ 2\theta_Y \end{pmatrix}, \quad (6.35)$$

where θ_2 is the angle of the beam when it reaches the second beam position sensor, and f is the focal length of the convex lens. This gives the beam spot displacement

$$x_2 = (2 \sin \alpha) \left(1 - \frac{d}{f}\right) x_L + 2 \left[\left(1 - \frac{d}{f}\right) r_{\text{lens}} + d \right] \theta_Y. \quad (6.36)$$

Now, the distance d is the distance between the lens and beam position sensor on the

optical table. It is a design parameter. It is obvious that when

$$d = \frac{r_{\text{lens}} f}{r_{\text{lens}} - f}, \quad (6.37)$$

the beam spot at the sensor plane reads

$$x_2 = -\frac{2f \sin \alpha}{r_{\text{lens}} - f} x_L. \quad (6.38)$$

This means that the beam position measured by the second sensor is only sensitive to longitudinal displacement x_L but not yaw θ_y . For this reason, the second beam position sensor is referred to the length-sensing QPD while the first beam position sensor is called the tilt-sensing QPD².




Figure 6.11: Sensing the beam spot position behind a convex lens. A convex lens is placed at r_{lens} from the suspended optics. The sensor is placed at distance d behind the lens.

This is not yet the full story. Each beam position sensor also measures the vertical displacement of the beam spot so there are in total four readouts. To break it down, the tilt-sensing and length-sensing optical lever path can be analyzed separately. Without loss of generality, consider the optical lever with tilt-sensing QPD shown in Fig. 6.12. The optical lever beam travels along an arbitrary vector in the 3D space and is reflected by the suspended optics. The reflected beam travels a vector of \vec{r} to reach the tilt-sensing QPD. The tilt-sensing QPD reads the beam position in two directions x_{tilt} and y_{tilt} . In reality, the directions of x_{tilt} and y_{tilt} do not necessarily align with the beam displacements caused by a pure yaw (horizontal) or a pure pitch (vertical). But, for simplicity, let us assume that they do align. The beam of the optical lever is projected into a vertical plane and a horizontal plane. And on those planes, the lengths of the reflected beam are r_h and r_v on the horizontal and vertical plane, respectively. The angle of incidences, as viewed on the horizontal and vertical planes, are α_h and α_v , respectively.

Here, without loss of generality, the beam spot at the optics is off-centered horizontally towards the x and y direction by an amount of δ_x and δ_y , respectively. With non-zero δ_x or δ_y , a rotation in yaw or pitch causes the reflection spot at the optics to move longitudinally, simulating a longitudinal displacement of the optics. This must be taken

²The tilt-sensing QPD is technically still coupled to longitudinal displacement according to Eqn. (6.34).




Figure 6.12: Tilt-sensing path of the optical lever. (Not to scale.)

into account in the derivation of the sensing matrix. The tilt-sensing QPD readouts now read

$$x_{\text{tilt}} = (2 \sin \alpha_h) (x_L + \delta_y \theta_P + \delta_x \theta_Y) + 2r_h \theta_Y, \quad (6.39)$$

and

$$y_{\text{tilt}} = (2 \sin \alpha_v) (x_L + \delta_y \theta_P + \delta_x \theta_Y) + 2r_v \theta_P, \quad (6.40)$$

where $\delta_x \theta_Y$ and $\delta_y \theta_P$ are the simulated longitudinal displacements due to rotation in yaw and pitch and an offsetted beam. Again, note that here the assumption is that x_{tilt} and y_{tilt} are aligned with the beam spot displacements caused by a pure yaw and pure pitch, respectively, of the optics.

The length-sensing path of the optical lever is shown in Fig. 6.13. In reality, length-sensing and tilt-sensing share the same light source. The beam is split into two paths by a beamsplitter hence they can be considered as separate paths as shown in Fig. 6.12 and Fig. 6.13. For the length-sensing part of the optical lever, the beam again is reflected from the suspended optics, and travels a path \vec{r}_{len} towards a convex lens with focal length f . The beam then travels a short distance d further to reach the length-sensing QPD, which reads the beam spot position in two directions x_{len} and y_{len} . Again, for simplicity, let us assume that x_{len} and y_{len} are aligned with the directions of the beam displacements caused by a pure yaw and pure pitch, respectively. Following Eqn. (6.36), the readouts




Figure 6.13: Length sensing path of the optical lever. (Not to scale)

x_{len} can be straightforwardly written as

$$x_{\text{len}} = 2 \sin \alpha_h \left(1 - \frac{d_h}{f_h} \right) (x_L + \delta_y \theta_P + \delta_x \theta_Y) + 2 \left[\left(1 - \frac{d_h}{f_h} \right) r_{\text{lens},h} + d_h \right] \theta_Y, \quad (6.41)$$

where d_h is the vector between the lens and the length-sensing QPD projected on the horizontal plane, f_h is the horizontal projection of the focal length on the horizontal plane, $r_{\text{lens},h}$ is the horizontal length between the lens and the reflection spot, and $\delta_x \theta_Y$ is the simulated longitudinal displacement due to a rotation in yaw and a horizontally offsetted beam. Assuming that the angle between the reflected beam and the horizontal plane is ϕ_h , the horizontal projections d_h , f_h , and $r_{\text{lens},h}$ can be written as $d \cos \phi_h$, $f \cos \phi_h$, and $r_{\text{lens}} \cos \phi_h$, respectively. This simplifies Eqn. (6.41) to

$$x_{\text{len}} = 2 \sin \alpha_h \left(1 - \frac{d}{f} \right) (x_L + \delta_y \theta_P + \delta_x \theta_Y) + 2 \cos \phi_h \left[\left(1 - \frac{d}{f} \right) r_{\text{lens}} + d \right] \theta_Y. \quad (6.42)$$

Again, it is easy to see that at a sensor placed at $d = r_{\text{lens}} f / (r_{\text{lens}} - f)$ decouples the readout from the yaw motion θ_Y , assuming small δ_x and δ_y . In this case, the horizontal length-sensing readout can be simplified to

$$x_{\text{len}} = -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} (x_L + \delta_y \theta_P + \delta_x \theta_Y). \quad (6.43)$$

Similarly, without further derivation, the vertical length-sensing readout can be written

as

$$y_{\text{len}} = -\frac{2f \sin \alpha_v}{r_{\text{lens}} - f} (x_L + \delta_y \theta_P + \delta_x \theta_Y). \quad (6.44)$$

At last, combining Eqn. (6.39), (6.40), (6.43), and (6.44) yields

$$\begin{aligned} \begin{pmatrix} x_{\text{tilt}} \\ y_{\text{tilt}} \\ x_{\text{len}} \\ y_{\text{len}} \end{pmatrix} &= \begin{bmatrix} 2 \sin \alpha_h & 0 & 2r_h \\ 2 \sin \alpha_v & 2r_v & 0 \\ -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} & 0 & 0 \\ -\frac{2f \sin \alpha_v}{r_{\text{lens}} - f} & 0 & 0 \end{bmatrix} \begin{pmatrix} x_L + \delta_y \theta_P + \delta_x \theta_Y \\ \theta_P \\ \theta_Y \end{pmatrix} \\ &= \begin{bmatrix} 2 \sin \alpha_h & 0 & 2r_h \\ 2 \sin \alpha_v & 2r_v & 0 \\ -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} & 0 & 0 \\ -\frac{2f \sin \alpha_v}{r_{\text{lens}} - f} & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & \delta_y & \delta_x \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x_L \\ \theta_P \\ \theta_Y \end{pmatrix}. \end{aligned} \quad (6.45)$$

Now, this equation relating the suspended optics displacement or angular displacements to the optical lever readouts assumes that the horizontal and vertical readouts align with the beam displacement vectors caused by pure yaw and pure pitch of the optics, which is not necessarily the case. To alleviate these assumptions, the left-hand side of Eqn. (6.45) must be multiplied by a matrix that takes into account the rotational misalignment of the QPDs, i.e.

$$\begin{bmatrix} \cos \psi_{\text{tilt}} & \sin \psi_{\text{tilt}} & 0 & 0 \\ -\sin \psi_{\text{tilt}} & \cos \psi_{\text{tilt}} & 0 & 0 \\ 0 & 0 & \cos \psi_{\text{len}} & \sin \psi_{\text{len}} \\ 0 & 0 & -\sin \psi_{\text{len}} & \cos \psi_{\text{len}} \end{bmatrix} \begin{pmatrix} x_{\text{tilt}} \\ y_{\text{tilt}} \\ x_{\text{len}} \\ y_{\text{len}} \end{pmatrix} = \begin{bmatrix} 2 \sin \alpha_h & 0 & 2r_h \\ 2 \sin \alpha_v & 2r_v & 0 \\ -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} & 0 & 0 \\ -\frac{2f \sin \alpha_v}{r_{\text{lens}} - f} & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & \delta_y & \delta_x \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x_L \\ \theta_P \\ \theta_Y \end{pmatrix}, \quad (6.46)$$

where ψ_{tilt} denote the angle of misalignment between the tilt-sensing QPD readouts and the pitch-yaw axis, and ψ_{len} denotes that of the length-sensing QPD. Nevertheless, an additional misalignment parameter δ_d can be added to representing the misplacement of the length-sensing QPD such that $d = r_{\text{lens}} f / (r_{\text{lens}} - f) + \delta_d$. But, this is not included in the equation for simplicity.

All optical levers in KAGRA have plane of incidence aligned to either the horizontal plane (type-A and type-Bp suspensions) or the vertical plane (type-B suspensions). This means that sensing matrix can be drastically simplified. For a horizontal optical lever, $r_v = r_h \cos \alpha_h$ and $\alpha_v = 0$. Assuming that optical lever is perfectly aligned, i.e. $\psi_{\text{tilt}} =$

$\psi_{\text{len}} = \delta_x = \delta_y = \delta_d = 0$, the QPD readouts for a horizontal configuration read

$$\begin{pmatrix} x_{\text{tilt}} \\ y_{\text{tilt}} \\ x_{\text{len}} \\ y_{\text{len}} \end{pmatrix} = \begin{bmatrix} 2 \sin \alpha_h & 0 & 2r_h \\ 0 & 2r_h \cos \alpha_h & 0 \\ -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{pmatrix} x_L \\ \theta_P \\ \theta_Y \end{pmatrix}, \quad (6.47)$$

which gives a geometric sensing matrix of

$$\mathbf{S}_{\text{OL,horizontal}} = \begin{bmatrix} 2 \sin \alpha_h & 0 & 2r_h \\ 0 & 2r_h \cos \alpha_h & 0 \\ -\frac{2f \sin \alpha_h}{r_{\text{lens}} - f} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}^+, \quad (6.48)$$

where $[\cdot]^+$ denotes the pseudo inverse of a matrix. As for the vertical optical levers, the geometric sensing matrix is simply

$$\mathbf{S}_{\text{OL,vertical}} = \begin{bmatrix} 0 & 0 & 2r_h \cos \alpha_v \\ 2 \sin \alpha_v & 2r_v & 0 \\ 0 & 0 & 0 \\ -\frac{2f \sin \alpha_v}{r_{\text{lens}} - f} & 0 & 0 \end{bmatrix}^+. \quad (6.49)$$

Diagonalization of sensing matrices

With the plane of incidence aligned either on the horizontal plane or the vertical plane, the geometric sensing matrix is drastically simplified. However, the given geometric sensing matrix does not guarantee that the optical lever readouts are aligned perfectly to the control basis. The geometric sensing matrices were derived with the assumption that the optical levers are perfectly aligned, which is not possible in reality. The same can be said for other sensors in other stages. In experience, cross-couplings between different control degrees of freedom in the order of $10^{-2} - 10^{-1}$ can usually be observed when using the geometric sensing matrix. While the fundamental assumption of the control topology described in Chap. 5 requires each degree of freedom to be able to controlled independently, it is best to reduce cross-couplings in sensor readouts to the greatest extent.

For the optical lever, there are six misalignment parameters, one angle of incidence α_h or α_v , two distances of the reflection spot from the center of the optics δ_x and δ_y , two QPD rotational misalignment angles ψ_{tilt} and ψ_{len} , and one length-sensing QPD misplacement δ_d , to list all of them. Putting them into the geometric sensing matrices Eqn. (6.48) or (6.49) would yield no zero elements in the matrix. With three control readouts, longitudinal displacement x_L , pitch θ_P , and yaw θ_Y , there are six possible ways that the degrees of freedom can be coupled to each other, longitudinal to pitch readout,

longitudinal to yaw readout, pitch to longitudinal readout, pitch to yaw readout, yaw to longitudinal readout, and yaw to pitch readout. Therefore, finding the correct values of the six misalignment variables would, in principle, decouple all degrees of freedom without the use of a diagonalization matrix in the form of Eqn. (6.32). One major advantage of this approach is the preservation of calibration. Although the misalignment parameters can be derived from measurable cross-couplings, the process can be tedious. On the other hand, using a diagonalization matrix to decouple all degrees of freedom is extremely convenient. Assuming that the calibration was correct before implementing the diagonalization matrix, with small non-diagonal values in the diagonalization matrix, the calibration factors are majorly preserved. In any case, post calibration can always be done in several ways, for example, inter-calibrating the diagonalized readouts with the readouts at the pre-isolator. Therefore, in KAGRA, using a diagonalization sensing matrix is the preferred way to achieve a finer control basis alignment of the sensors.

To see how a diagonalization sensing matrix works, let us assume that a geometric sensing matrix \mathbf{S}_{geo} operating on sensor readouts \mathbf{x} gives a coupled control readout $\mathbf{y}' = \mathbf{C}\mathbf{y}$ such that

$$\mathbf{S}_{\text{geo}}\mathbf{x} = \mathbf{y}' = \mathbf{C}\mathbf{y}, \quad (6.50)$$

where $\mathbf{y}' = \{y'_i\}$ is the coupled control readout, $\mathbf{C} = \{C_{ij}\}$ is the coupling matrix and $\mathbf{y} = \{y_i\}$ is the array of desired readout, e.g. $\{x_L, \theta_P, \theta_Y\}$ for the optical lever. Suppose the geometric sensing matrix correctly transforms the sensor readouts to give calibrated control readouts, the coupling matrix can be written as

$$C_{ij} = \begin{cases} 1, & i = j \\ \frac{y'_i}{y_j}, & i \neq j \end{cases}, \quad (6.51)$$

where y'_i/y_j is the cross-coupling from the j^{th} displacement to the i^{th} readout, or simply j to i coupling. If the couplings are measurable, then it is obvious that an additional matrix $\mathbf{S}_{\text{diag}} = \mathbf{C}^{-1}$ operating on the coupled readout would yield the desired control readout,

$$\mathbf{S}_{\text{diag}}\mathbf{S}_{\text{geo}}\mathbf{x} = \mathbf{C}^{-1}\mathbf{C}\mathbf{y} = \mathbf{y}, \quad (6.52)$$

where \mathbf{S}_{diag} is the diagonalization sensing matrix. Therefore, there are two stages of sensing matrices in the control topology. The calibrated sensor readouts first go through the geometric sensing matrix, and then, the diagonalization sensing matrix. One major advantage of using the diagonalization sensing matrix is that it can be obtained multiple times in an iterative process, which improves the quality of diagonalization in each iteration. For consecutive diagonalization, the new diagonalization sensing matrix is simply $\mathbf{S}_{\text{diag}}^{\text{new}} = \mathbf{C}^{-1}\mathbf{S}_{\text{diag}}^{\text{old}}$.

In order to obtain the diagonalization sensing matrix, the cross-couplings y'_i/y_j must

be measured. Here, y'_i is just the i^{th} control readout output from the geometric sensing matrix, which can be measured directly. As for y_j , the j^{th} actual displacement of the suspension, a successful measurement of it relies on one of two things

1. Whether the suspension has a resonance involving only that degree of freedom but not the other degrees of freedom that the diagonalization is attempting to decouple from, or
2. Whether there is a way to actuate the suspension in that degree of freedom but not the degrees of freedom that the diagonalization is attempting to decouple from.

Suppose there exists a resonance in the j^{th} degree of freedom, it needs to be excited so the signal-to-noise ratio is high enough for a good measurement of the coupling. There are two ways this can be done, either by injecting a white noise or an impulse signal to that degree of freedom via the actuators. Here, the actuators are not assumed to be aligned to the control basis yet but it does not affect the results. This is because the frequency spectrum of a white noise or an impulse signal is flat. This means that any actuation coupling is frequency-independent and only the impulse responses of the suspension can be observed, i.e. shapes that look like the transfer functions. At the resonance frequency, the j^{th} readout y'_j is dominated by the resonance of the j^{th} degree of freedom, even if it is cross-coupled to other degrees of freedom. Under this condition, y_j can be estimated directly from the impulse response at the resonance frequency using the y'_j readout, and y'_i can also be read at the same frequency in a simultaneous measurement. On the other hand, if the particular degree of freedom can be actuated purely, then $y_j \approx y'_j$ and y'_i can be measured straightforwardly when the suspension is actuated purely in that degree of freedom.

The SRM suspension at KAGRA is used as an example to demonstrate the use of a diagonalization matrix. Fig. 6.14 shows three measurements of the readouts in the control basis, without a diagonalization sensing matrix, at the test mass stage of the SRM suspension, each corresponding to a white noise actuation in longitudinal, pitch, and yaw direction using the test mass actuators with the geometric actuation matrix (To be discussed in Sec. 6.2.2.). The actuators are therefore assumed to be not aligned to the control basis. Here, the readouts in the control basis are outputs directly from the geometric sensing matrix of the optical lever. As can be seen, the geometric sensing matrix indeed coarsely aligns the sensors to the control basis. This is indicated by the fact that all three readouts have distinct shapes corresponding to different dynamics of the suspension.

In Fig. 6.14, the resonance frequencies in interest are identified from the diagonal readouts, i.e. longitudinal readout with longitudinal actuation (blue line in (a)), pitch readout with pitch actuation (orange line in (b)), and yaw readout with yaw actuation (green line in (c)). The corresponding resonance frequencies are located at 0.658 Hz

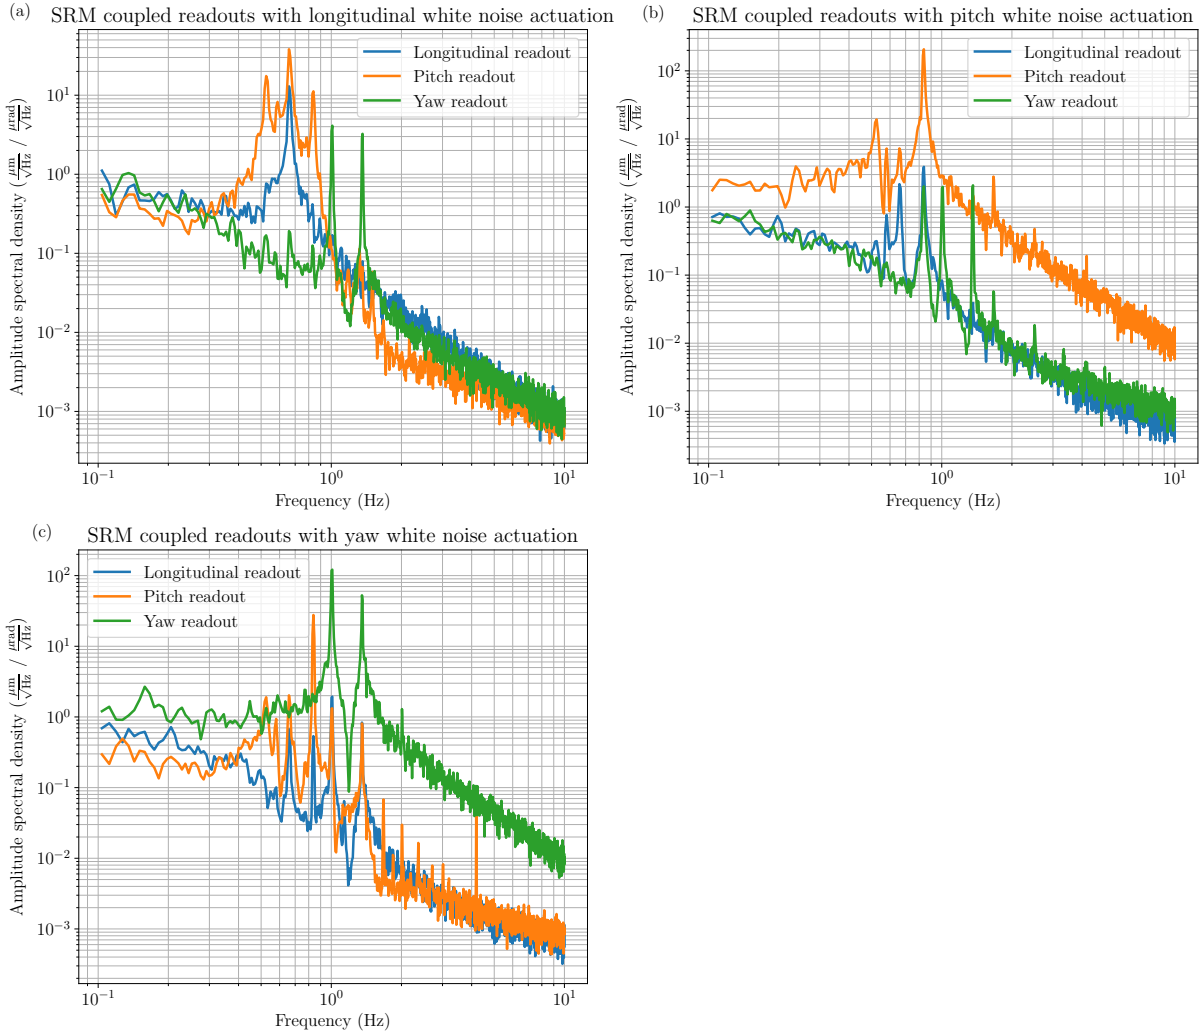


Figure 6.14: SRM test mass stage control readout without diagonalization sensing matrix. (a) Readouts with white noise actuation in the longitudinal direction. (b) Readouts with white noise actuation in the pitch direction. (c) Readouts with white noise actuation in the yaw direction. Blue: Longitudinal readout. Orange: Pitch readout. Green: Yaw readout.

(pendulum mode), 0.838 Hz (Pure pitch mode), and 1.010 Hz / 1.354 Hz (pure yaw mode). Except for the pendulum mode at 0.658 Hz, these resonances are pure in each degree of freedom and should not be observed in readouts of other degrees of freedom. Therefore, readouts at these resonance frequencies can be used to determine the coupling ratios C_{ij} between different degree of freedom. The pendulum mode, which is a resonance mode involving longitudinal and pitch motion of the optics, cannot be used to determine cross-coupling between longitudinal displacement and pitch readout. However, it is still useful for determining longitudinal displacement to yaw coupling as it should not be observed in the yaw readout.

In subplot (a) of Fig. 6.14, a small peak is observable in the yaw readout at 0.658 Hz. However, in this case it is unclear whether this coupling comes from the longitudinal

displacement or pitch at this stage. Therefore, the coupling ratio from longitudinal displacement to yaw readout is undetermined for the first iteration. In subplot (b), the pitch resonance is very clear and it obviously couples to both longitudinal and yaw readouts. The corresponding coupling ratios are -0.018781 and 0.00968755 for longitudinal and yaw readouts, respectively. Note that the values are not the ratio of the two amplitude spectral densities. The magnitude of the ratios are obtained from the ratio between the Fourier transform of the readouts (not shown here), and the signs are obtained from the phase. The coupling ratio is said to be valid only when the phase is close to 0° or $\pm 180^\circ$, which correspond to positive and negative signs, respectively. Any phase in between 0° and $\pm 180^\circ$ means the result is inconclusive and a coupling ratio cannot be obtained. In subplot (c), the yaw resonance is observable at 1.354 Hz. The yaw to longitudinal and pitch readouts couplings are also observable and the ratios are 0.0158224 and -0.0153659 for longitudinal and pitch readouts, respectively.

The longitudinal displacement to pitch readout coupling cannot be determined from any resonances since a pure longitudinal mode is simply non-existence. However, a short experiment involving the pre-isolator can be done to obtain this ratio. The three degrees of freedom of the pre-isolator table are longitudinal, transverse, and yaw. A DC offset, i.e. actuation at 0 Hz, in longitudinal actuation signal of the pre-isolator could shift the would cause a shift of the whole suspension chain in the longitudinal direction. If the pre-isolator actuators are not diagonalized, this would result in a small shift in the transverse direction and a small rotation in the yaw direction. But, a DC shift in the pre-isolator can, in principle, never cause the optics to pitch. Therefore, any shift in the pitch readout at the optics during this experiment can be treated as coupling from longitudinal displacement and can be used compute the coupling ratio from longitudinal displacement to pitch readout.

The experiment was performed and the results are as follows. A static shift in the positive longitudinal direction was applied to the pre-isolator. After that, the longitudinal readout and pitch readout from the optical lever changed by an amount of $185.1934 \mu\text{m}$ and $1.8085 \mu\text{rad}$. This gives a longitudinal to pitch readout coupling of $1.8085/185.1934 = 0.00976547$. Combining this ratio with the values obtained from the resonance measurements, the coupling matrix is

$$\mathbf{C}_1 = \begin{bmatrix} 1 & -0.018781 & 0.0158224 \\ 0.00976547 & 1 & -0.0153659 \\ 0 & 0.00968755 & 1 \end{bmatrix}. \quad (6.53)$$

The inverse of this matrix was applied as the initial diagonalization sensing matrix for the first iteration and the measurement was repeated.

After the implementation of the initial diagonalization sensing matrix, nearly all cross-coupling observed previously becomes undetectable with one exception. The decoupling

of the pitch resonance from the yaw readout reveals a longitudinal displacement to yaw readout coupling at the pendulum mode resonance frequency. Fig. 6.15 shows the readouts




Figure 6.15: SRM initially diagonalized test mass readouts with white noise actuation in the longitudinal direction. Blue: Longitudinal readout. Orange: Pitch readout. Green: Yaw readout.

of the optics displacement with white noise actuation in the longitudinal direction. These readouts are obtained after implementing the inverse of the coupling matrix Eqn. (6.53). As can be seen, the 0.658 Hz peak in the yaw readout becomes far more observable compared to that in subplot (a) of Fig. 6.14. The reason of this is perhaps due to the removal of pitch to yaw coupling, which may have partially canceled this peak in the initial measurement. The 0.658 Hz peak in the yaw readout shown in Fig. 6.15 cannot be a pitch to yaw coupling. If it was due to pitch motion, there would have been three peaks instead of one, as there are three (at least) observable resonances in the pitch readout (orange line as shown in the figure). Therefore, it has to be a longitudinal displacement to yaw readout coupling and the coupling ratio was obtained to be 0.0331021. This gives a second coupling matrix of

$$\mathbf{C}_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0.0331021 & 0 & 1 \end{bmatrix}. \quad (6.54)$$

The final diagonalization sensing matrix can be obtained as $\mathbf{S}_{\text{diag}} = \mathbf{C}_2^{-1}\mathbf{C}_1^{-1}$.

Fig. 6.16 shows the measurements taken after implementation of the diagonalization sensing matrix. As can be seen, the original cross-couplings have been reduced by a great extent. As a verification, the readouts were also taken for an impulse actuation instead of white noise actuation and the results are shown in Fig. 6.17. In this experiment, an impulse signal was sent to one of the actuators at the test mass stage, exciting all

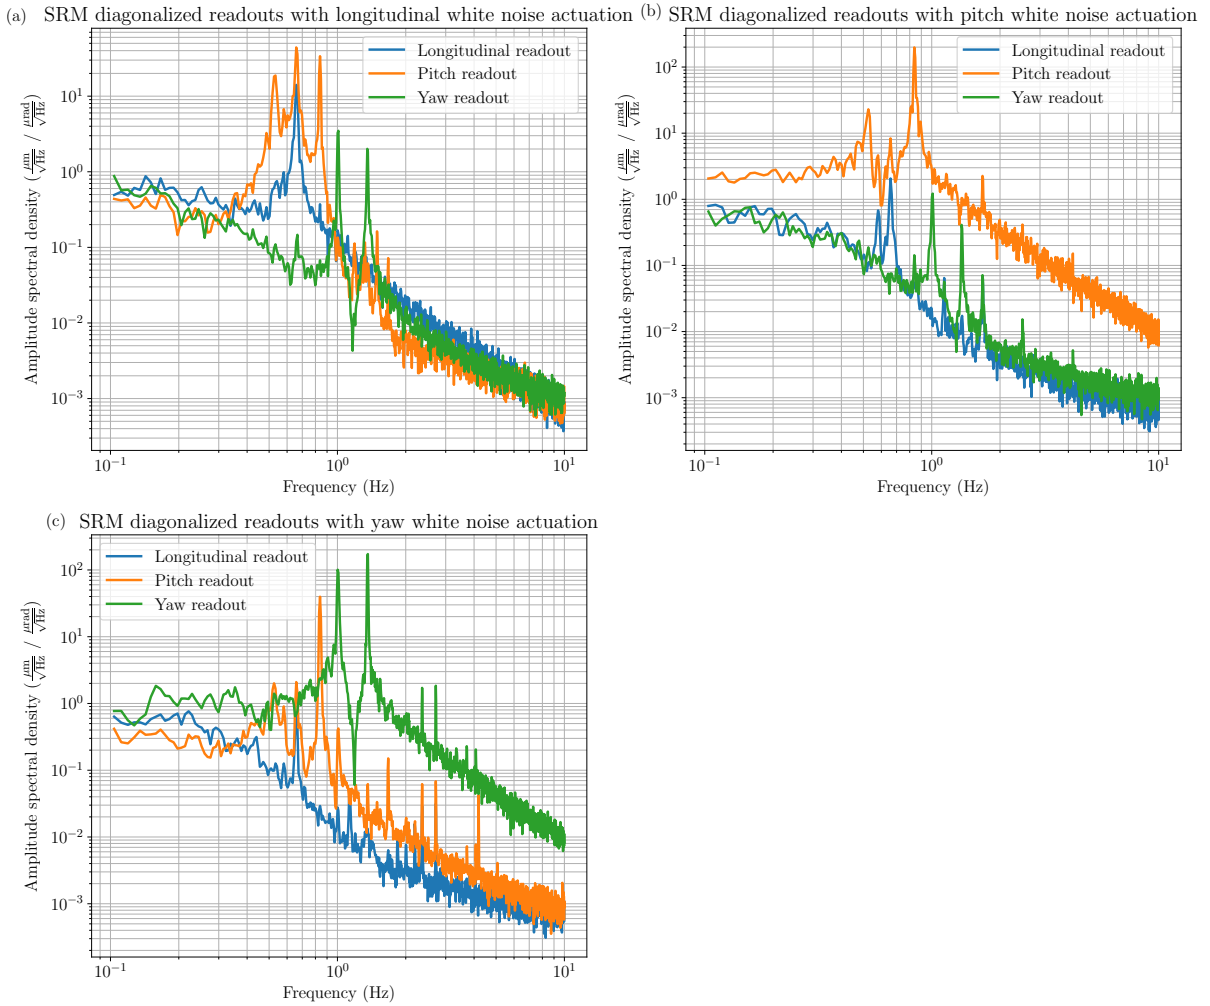


Figure 6.16: SRM test mass stage control readout with diagonalization sensing matrix. (a) Readouts with white noise actuation in the longitudinal direction. (b) Readouts with white noise actuation in the pitch direction. (c) Readouts with white noise actuation in the yaw direction. Blue: Longitudinal readout. Orange: Pitch readout. Green: Yaw readout.

resonances in all degrees of freedom. Unfortunately, the signal-to-noise ratio below 0.8 Hz is worse compared to that of the white noise actuation measurements and therefore cannot be used to verify the results. Above 0.8 Hz, it is clear that the cross-couplings are very minimal. Yaw displacement to pitch readout couplings at 1.010 Hz and 1.354 Hz are still clearly visible but they are well below the 10^{-2} level. With that said, it is worth mentioning that the 10^{-2} cross-coupling level is not a benchmark and there is currently no required specifications for diagonalization. The control performance degradation due to cross-coupling is not well-studied at the moment and this should be investigated in the future.




Figure 6.17: SRM test mass stage control readout with diagonalization matrix with impulse actuation. Blue: Longitudinal readout. Orange: Pitch readout. Green: Yaw readout.

6.2.2 Actuation matrices

Actuation matrices are matrices that align the desired actuators to the control basis. Technically, it is the other way around: It operates on the control actuation signal \mathbf{u} and distribute the control signals to each individual actuators,

$$\mathbf{A}\mathbf{u} = \mathbf{a}, \quad (6.55)$$

where $\mathbf{u} = \{u_1, u_2, u_3, \dots, u_p\}$ is the array of control actuation signals from the, e.g. actuation signal in the longitudinal direction, $\mathbf{a} = \{a_1, a_2, a_3, \dots, a_q\}$, is the array of actuation signals going to each individual actuator, e.g. actuation signal of the first actuator, and $\mathbf{A} \in \mathbb{R}^{q \times p}$ is the actuation matrix aligning the control signals to the actuators.

Like sensing matrices, the actuation matrix maybe decomposed into two matrices, geometric actuation matrix and diagonalization actuation matrix

$$\mathbf{A} = \mathbf{A}_{\text{geo}}\mathbf{A}_{\text{diag}}, \quad (6.56)$$

where \mathbf{A}_{geo} is the geometric actuation matrix and \mathbf{A}_{diag} is the diagonalization actuation matrix. Note that the order of operation is different from that of sensing matrix (6.32) and this is important. The goal of the actuation matrices is to correctly distribute actuation signals to the actuators such that each control signal only actuates the system in the corresponding degree of freedom. In this section, the geometric actuation matrix at the test mass stage is derived as an example and the diagonalization of actuation matrices is

discussed.

Actuation matrix of the test mass actuators

Fig. 6.18 shows the location and direction of actuation at the test mass stage of the main optics suspensions. As shown in the figure, the actuators are located at the four corners of the optics, top-left, bottom-left, bottom-right, and top-right, in the order of ascending actuator number. This particular configuration is exclusive to the type-B and type-Bp suspensions. The type-A suspensions use a diamond configuration where four actuators are placed top and bottom, left and right of the suspended optics.




Figure 6.18: Test mass actuators located at four corners of the test mass.

The positive direction of the actuation is typically defined as the direction of the longitudinal direction. Therefore, a positive actuation signal in any of the actuators pushes the suspended optics in the longitudinal direction. Let the control actuation signals be $\mathbf{u} = [u_L, u_P, u_Y]^T$, where u_L is the actuation in the longitudinal direction, u_P is the actuation in the pitch direction, and u_Y be the actuation in the yaw direction. This gives the first column of the geometric actuation matrix to be $[1, 1, 1, 1]^T$. Note that the scale is arbitrary because there is no calibration factors to be preserved, unlike the geometric sensing matrix.

Actuators 1 and 4 are above the center of mass of the optics and actuators 2 and 3 are below that. Treating the suspended optics as a free mass, positive signals in actuators 1 and 4 give a torque in the pitch direction, and positive signals in actuator 2 and 3 give a torque in the negative pitch direction. Assuming that the actuators are placed with reflection symmetry along the horizontal plane, this gives the second column of the actuation matrix to be $[1, -1, -1, 1]^T$. Similarly, actuators 1 and 2 are on the left side of the optics and actuators 3 and 4 are on the right side. Positive signals in actuators 1

and 2 give a torque in the negative yaw direction while positive signals in actuators 3 and 4 give a torque in the positive direction. This gives the third column of the geometric actuation matrix to be $[-1, -1, 1, 1]^T$. The geometric actuation matrix at the test mass stage is fully defined as

$$\mathbf{A}_{\text{TM}} = \begin{bmatrix} 1 & 1 & -1 \\ 1 & -1 & -1 \\ 1 & -1 & 1 \\ 1 & 1 & 1 \end{bmatrix}. \quad (6.57)$$

The reason why the number 1 is used in the matrix is simply to avoid confusion to the operators. The geometric actuation matrix represents the geometric placement of the actuators and it should be easy to understand from the digital system. At least, this is the general consent at KAGRA. Therefore, there is no reason to use other numbers other than 1 in the geometric actuation matrix.

It should be pointed out that the choice of this matrix is rather unnatural and is different from what is provided in Refs. [43, 44, 29]. One alternative way to approach the scaling is to firstly calibrate the actuators in units of N/counts using the output filters as shown in Fig. 5.1. Then, the actuation signals (output of the actuation matrix, inputs of the output filters) are forces applying to the optics at the locations of the actuators. And then, the control signals (outputs of the control filters, input of the actuation matrix) are forces or torques acting on the optics in the control degrees of freedom. In this case, the geometric actuation matrix then carries the information of normalization factors and geometric information such as the position of the actuators. This is the convention followed by Refs. [43, 44]. However, in reality, unlike sensor calibration, getting the calibration factors for the actuation is very difficult. This is due to the coil-magnet nature of the actuators where actuation efficiency depends on the distance between the magnet and the coil. The actual distance between the magnet and the coil can change within a certain tolerance and is unknown before the installation. For this reason, the actuators cannot be calibrated before installation. And, the actuators cannot be calibrated when they are already installed. Therefore, setting the geometric actuation matrix to retain physical units is not that feasible in reality.

Nevertheless, an alternative way to scale the geometric actuation matrix is to have each column normalized to a certain norm. In Ref. [29], the 1-norm of each is normalized to 1. Another possible choice is to choose the eigenvectors from the singular value decomposition of the matrix (6.57), which naturally have 2-norms normalized to 1. In any case, the scale of the actuation matrix is merely a convention or preference and has no actual affection to the control performance.

Diagonalization of actuation matrices

The geometric actuation matrix is not sufficient to align the control basis to the actuator basis. This is because the geometric actuation matrix was derived with the assumption that the actuated objects are free bodies while they are really constrained by the dynamics of the suspension. Moreover, the geometric actuation matrix does not take into account the variation of the actuators. Therefore, the geometric actuation matrices are often not sufficient in aligning the control basis. However, in contrast to its sensing counterpart, the diagonalization of actuation matrices is not as straightforward and is often not considered. This is because actuation matrix diagonalization is not as important when feedback control in action. Actuation cross-couplings between different degrees of freedom can be treated as disturbances, which are minimized with active control.

At 0 Hz, a non-zero actuation signal applies a static force to the suspension, causing it to shift. The relationship between the control actuation signals and the suspension displacement reads

$$\mathbf{y} = \mathbf{E}\mathbf{u}, \quad (6.58)$$

where \mathbf{E} is the actuation efficiency matrix, \mathbf{u} is the actuation signals in the control basis, and \mathbf{y} is the displacement of the suspension. The off-diagonal elements of the \mathbf{E} matrix represents the cross-coupling and are ideally zero. The diagonal elements are the actuation efficiencies of the actuators. Depending on the situation, they need to be retained after even diagonalization so the scale of the diagonal transfer functions are comparable. \mathbf{E} is physical and cannot be altered. Therefore, any additional control matrices must operate before the actuation efficiency matrix. In order to achieve actuation diagonalization, it is obvious that by modifying control actuation signals from \mathbf{u} to $\mathbf{E}^{-1}\text{diag}(\mathbf{E})\mathbf{u}$, it gives a diagonal actuation efficiency, i.e.

$$\mathbf{E} [\mathbf{E}^{-1}\text{diag}(\mathbf{E})\mathbf{u}] = \text{diag}(\mathbf{E})\mathbf{u} = \mathbf{y}. \quad (6.59)$$

Comparing this with Eqn. (6.55) and (6.56), it is easy to see that by setting the diagonalization actuation matrix to

$$\mathbf{A}_{\text{diag}} = \mathbf{E}^{-1}\text{diag}(\mathbf{E}) \quad (6.60)$$

would achieved the required results.

With a diagonalized sensing matrix, measuring the actuation efficiency matrix \mathbf{E} is straightforward. For this reason, the sensing matrix should be diagonalized before diagonalizing the actuation matrix. The actuation efficiency matrix can be written as $\mathbf{E} = \{E_{ij}\}$ and

$$E_{ij} = \frac{y_i}{u_j}, \quad (6.61)$$

where y_i is the i^{th} displacement in the control basis, and u_j is the j^{th} actuation signal in the control basis. A known static actuation signal u_j can be applied one at a time and all displacements \mathbf{y} can be measured using the diagonalized displacements signals, giving one column of the actuation efficiency matrix \mathbf{E} . Like sensing matrices, the diagonalization actuation matrix can be obtained iteratively by setting the new diagonalization actuation matrix to $\mathbf{A}_{\text{diag}}^{\text{new}} = \mathbf{A}_{\text{diag}}^{\text{old}} (\mathbf{E}^{\text{new}})^{-1} \text{diag}(\mathbf{E}^{\text{new}})$.

The actuation diagonalization described above corresponds to the diagonalization at 0 Hz but the procedure can be done at any one and only one frequency. In such case, the actuation injection would be sinusoidal at a particular frequency. The readback of the actuation efficiency E_{ij} would then be the ratio between between the fourier transform of the displacement signal y_i and the actuation signal u_j at the frequency of the sine wave. Again, this measurement is only said to be value if the phase between y_i and u_j is close to 0° or $\pm 180^\circ$. Any measured phase between 0° and $\pm 180^\circ$ renders the result not useful. At 0 Hz, the phase between the injection and displacement is typically 0° or $\pm 180^\circ$, owing to the spring-mass nature of the suspension. The same can also be said for frequencies much higher than the resonance. Therefore, diagonalization are typically done at these frequencies.

There is a caveat to this approach, that is, the actuators are only diagonalized at one frequency and there is not guarantee that the actuators are also diagonalized at other frequencies. In worst case, the actuation cross-coupling is reduced at one frequency and increased at other frequencies. The actuation efficiency matrix \mathbf{E} in Eqn. (6.58) represents the relationship between the actuation signal and the displacement at one frequency. To generalized the relationship, the actuation efficiency is simply the actuation transfer functions and \mathbf{E} should be replaced by a transfer function matrix, which denotes the input-output relationship of a multiple-input-multiple-output system. Therefore, the proper way to approach actuation diagonalization is to model all transfer functions and compute the frequency-dependent actuation diagonalization matrix using Eqn. (6.60) via an analytic form of the matrix inverse. Alternatively, the measured transfer function matrix, in the form of a complex array, can be inverted numerically to obtain the frequency response of the require frequency-dependent actuation matrix. And then, the each frequency response element in the matrix can be modeled by transfer functions. Either way, the work can be tedious and unnecessary. So, this type of diagonalization has never been implemented at KAGRA due to a tight schedule.

6.3 Modeling and fitting

One major aspect in control system is called system identification, which seeks to obtain a mathematical model a dynamical system from measured data. The goal of system identification is not to fully characterize the dynamical system. Instead, the goal is

to obtain a model that captures the required system dynamics such that the control performance and stability evaluated with the model sufficiently represent those of the real system when control is implemented. This way, the controller can be designed and tested in simulations before actually implementing the control systems.

System identification, by itself, is an ongoing field of research and there are a lot of advanced methods, such as the eigensystem realization algorithm [50], that can identify the full dynamics of the suspension as a multiple-input-multiple-output system. Since the control topology of the active isolation assumes each degrees of freedom to be decoupled, each degrees of freedom can be assumed to be a single-input-single-output system, as shown in Fig. 5.5. In order to design a controller for this, the plant needs to be modeled and it is simply the ratio between the actuation and the displacement. In addition, the spectrum of the disturbance and the noise defined in Fig. 5.5 can optionally be modeled by empirical models or even transfer functions. The former can be used to evaluate the closed-loop performance and noise budget while the latter is useful for \mathcal{H}_∞ optimization, which is to be discussed in Part III. In this section, both topics are covered. Sec. 6.3.1 discusses the modeling of the actuation plant with a focus on fitting transfer functions of the multiple-pendulum suspensions. A real example will also be given to demonstrate the methods discussed. Sec. 6.3.2 discusses the modeling of frequency spectrum, which will be useful in Part III.

6.3.1 Modeling frequency response data with transfer functions

In this subsection, a systematic approach is given to obtain a transfer function model given a measured frequency response of the suspension. The approach can be divided into a few steps:

1. Measurement,
2. Model selection, and
3. Regression.

Notes on frequency response measurement

To model the dynamics of the system, the first step is to obtain a frequency response of the system. Here, the dynamics is referred to the open-loop dynamics of the system, i.e. $P(s)$ in Fig. 5.5. In principle, this is as easy as injecting a known actuation signal u and then obtain a measurement of the system output y . The frequency response of the plant $P(j\omega_i)$ can then be estimated to be

$$P(j\omega_i) \approx \frac{P_{uy}(\omega_i)}{P_{uu}(\omega_i)}, \quad (6.62)$$

where ω_i is the discrete angular frequency data, $P_{uy}(\omega_i)$ is the cross power spectral density between the actuation u and the output y , and $P_{uu}(\omega_i)$ is the power spectral density of the actuation signal u . KAGRA utility `diag / diaggui` and LIGO utility `dtg / dtggui` automatically does this estimation when a transfer function plot³ is requested⁴. Therefore, obtaining a frequency response data is straightforward at KAGRA.

However, there are some caveats that are worth mentioning regarding the measurement of suspension frequency responses. First of all, the result of the modeling is only going to be as good as the measurement data. Therefore, it is important to ensure that the full spectrum of the frequency response data is valid. The frequency response data of a suspension can only be measured using real sensors and when the particular degree of freedom is not under any influence of feedback control. Under this condition, the displacement measurement of the suspension sensors contain sensor noise. And, the displacement is not only excited by the injected actuation but also static external disturbance such as the seismic noise. Therefore, it is important that the injected actuation signal sufficient excite the motion of the suspension such that the displacement due to actuation dominates that of the disturbance and the noise. However, the actuation should not be excessively high to a level where the displacement signal gets distorted due to non-linearity of the sensor as discussed in Sec. 6.1.1. In such a case, higher-order harmonics, which do not represents the dynamics of the system, will occur in the frequency response data. Hence, getting the actuation magnitude (at all frequencies) adequate the crucial first step to successful modeling of the plant.

Model selection

Assuming that the relationship between the actuation signal u and the displacement y is governed by a generic linear differential equation

$$a_0 y + a_1 \frac{dy}{dt} + a_2 \frac{d^2 y}{dt^2} + \dots + a_n \frac{d^n y}{dt^n} = b_0 u + b_1 \frac{du}{dt} + b_2 \frac{d^2 u}{dt^2} + \dots + b_m \frac{d^m u}{dt^m}, \quad (6.63)$$

where a_i and b_i are real coefficients. Taking the Laplace transform if the governing equation gives a transfer function, defined by the ratio of the output and the input,

$$\begin{aligned} (a_0 + a_1 s + a_2 s^2 + \dots + a_n s^n) Y(s) &= (b_0 + b_1 s + b_2 s^2 + \dots + b_m s^m) U(s) \\ \frac{Y(s)}{U(s)} &= \frac{b_0 + b_1 s + b_2 s^2 + \dots + b_m s^m}{a_0 + a_1 s + a_2 s^2 + \dots + a_n s^n}, \end{aligned} \quad (6.64)$$

³Technically, it shouldn't be called a transfer function plot. Instead, it should be called either frequency response or Bode plot.

⁴Obtaining a frequency response is actually easier than obtaining a Fourier transform series using these utilities.

where s is the Laplace variable, $Y(s)$ and $U(s)$ are the Laplace transform of the output y and the actuation signal u , respectively. A generic transfer function model corresponding the physical dynamical system can then be written as

$$\hat{P}(s) = \frac{b_0 + b_1s + b_2s^2 + \cdots + b_ms^m}{a_0 + a_1s + a_2s^2 + \cdots + a_ns^n}, \quad (6.65)$$

and a_i and b_i are the model parameters that we seek to obtain via modeling and fitting. However, the modeling and fitting of the frequency response data can be difficult with a transfer function model in this form. This is due to two reasons:

1. n and m are integers that cannot be obtained via typically fitting methods such as linear regression, and
2. Fitting a_i and b_i lacks numerical stability as the poles of the transfer function model approaches the right-hand side of the complex plane during process of fitting.

Both issues can be avoided by recently developed algorithms such as the adaptive Antoulas-Anderson (AAA) and AAA² algorithms [51, 52].

However, the modeling of mechanical transfer functions of a multiple-pendulum suspension is a simpler task than fitting a generic transfer function. This is because the transfer function model of a multiple-pendulum, which is a coupled oscillator, can be simplified to

$$\hat{P}(s) = k \frac{\prod_{i=1}^{n_z} s^2 + \frac{\omega'_{n,i}}{q'_i} s + \omega'^2_{n,i}}{\prod_{i=1}^{n_p} s^2 + \frac{\omega_{n,i}}{q_i} s + \omega^2_{n,i}}, \quad (6.66)$$

where $k \in (-\infty, \infty)$, $\omega_{n,i} \in (0, \infty)$, $\omega'_{n,i} \in (0, \infty)$, $q_i \in [0.5, \infty)$, and $q'_i \in [0.5, \infty)$ are the model parameters to be found. Note that here n_p and n_z are redefined to be the number of complex-pole pairs and the number of complex-zero pairs, respectively. The element $(s^2 + \omega'_{n,i}s/q'_i + \omega'^2_{n,i})$ corresponds to a pair of complex zeros located at frequency $\omega = \omega'_{n,i}$ and the element $1/(s^2 + \omega_{n,i}s/q_i + \omega^2_{n,i})$ denotes a pair of complex poles located at frequency $\omega = \omega_{n,i}$. From hereafter, a pair of complex poles or complex zeros is referred to a complex pole or a complex zero⁵. In the frequency response, a complex zero corresponds to a notch/antiresonance in the magnitude plot around $\omega = \omega'_{n,i}$ while a complex pole corresponds to a peak/resonance around $\omega = \omega_{n,i}$. This means that the numbers n_p and n_z can be determined from the frequency response measurements, given that the peaks and notches are observable. This solves issue 1.

The stability of the transfer function is determined by the real numbers of the complex poles. The complex poles polynomial can be factorized as

$$\frac{1}{s^2 + \frac{\omega_{n,i}}{q_i} s + \omega^2_{n,i}} = \frac{1}{(s - s_p)(s - s_p^*)}, \quad (6.67)$$

⁵For a physical system, complex zeros or complex poles must come in pairs.

where

$$s_p = \frac{\omega_{n,i}}{2q_i} \left(-1 + j\sqrt{4q_i^2 - 1} \right) \quad (6.68)$$

is one of the complex poles, s_p^* denotes its complex conjugate, which is the other complex pole, and j is the imaginary number. It is obvious that when $0 < \omega_{n,i} < \infty$ and $0.5 \leq q_i < \infty$, the real part of s_p is $-\omega_{n,i}/(2q_i)$, which is always less than zero. Under these constraints, the fitting problem is numerically stable and this solves issue 2, making Eqn. (6.66) a suitable model for fitting the frequency response of a suspension. There are several ways to impose these constraints. For example, the fit can be done via a constrained optimization. Alternatively, instead of using k , $\omega_{n,i}$, $\omega'_{n,i}$, q_i , and q'_i as the model parameters, use $\log k$, $\log \omega_{n,i}$, $\log \omega'_{n,i}$, $\log(q_i - 0.5)$ and $\log(q'_i - 0.5)$ instead. This not only ensures that the numerical stability constraints are imposed, but also scale the parameters properly taking into account the large dynamic range of the frequency response, which is typically plotted with a log-log graph. This may improve the results obtained from numerical optimization. Therefore, the latter approach is the preferred way.

Fitting frequency response data

Once a transfer function model is selected, the final fitting problem can be defined as an optimization problem seeking a minimization:

$$\min_{\theta} \left(\frac{1}{N} \sum_{i=1}^N \left(\log \left| \hat{P}(j\omega_i; \theta) - P(j\omega_i) \right|^2 \right) W(\omega_i) \right), \quad (6.69)$$

where, ω_i is the discrete frequency data of the measurement, $\hat{P}(j\omega_i; \theta)$ is the transfer function model, θ is the array of model parameters that need to be optimized, $P(j\omega_i)$ is the measured frequency response given by Eqn. 6.62, and $W(\omega_i)$ is a weighting function. Similar to geophone calibration described in Sec. 6.1.2, the weighting function is a frequency-dependent function that can be used to weight or filter out unwanted data points. The plant model for fitting the frequency response of a suspension is given by Eqn. 6.66. But, it is more convenient to rewrite Eqn. 6.66 in terms of normalized frequency:

$$\hat{P}(s) = k_{\text{DC}} \frac{\prod_{z=1}^{n_z} \left(\frac{s}{\omega'_{n,i}} \right)^2 + \frac{1}{q'_i} \frac{s}{\omega'_{n,i}} + 1}{\prod_{p=1}^{n_p} \left(\frac{s}{\omega_{n,i}} \right)^2 + \frac{1}{q_i} \frac{s}{\omega_{n,i}} + 1}, \quad (6.70)$$

where k_{DC} denotes the DC gain of the transfer function. This gives the model parameter array $\theta = \{k_{\text{DC}}, \omega_{n,i}, \omega'_{n,i}, q_i, q'_i\}$, which has a length of $2(n_p + n_z) + 1$, where n_p is the number of pairs of complex poles, and n_z is that for the complex zeros. Again, it should be reminded that the optimization works best when by optimizing the logarithm of the model parameters, not the parameters themselves.

The optimization problem (6.69) can be solved in two ways,

1. Using a local optimization algorithm, such as the Nelder-Mead method, with an initial guess, or
2. Using a global optimization algorithm, such as differential evolution, with bounded model parameters.

For the global optimization option, the boundaries of the model parameters for model (6.70) can be obtained as follows. The DC gain k_{DC} refers to the magnitude of the frequency response at 0 Hz. While the frequency response of a couple-oscillator is flat below the resonance frequencies, the boundary of k_{DC} can be set to a close range around the magnitude of the measure frequency response at the lowest frequency $|P(j\omega_1)|$. For example, one can set the boundary to one order of magnitude around the magnitude response at the lowest frequency, i.e. $[0.1|P(j\omega_1)|, 10|P(j\omega_1)|]$. Note that here the assumption is that the measured frequency response adequately captures the low frequency response of the suspension. As for the resonance frequencies $\omega_{n,i}$ and $\omega'_{n,i}$, the boundary can be safely set to the domain of the frequency response, i.e. $[\min(\omega_i), \max(\omega_i)]$. This reasoning is that the number of pole pairs and zero pairs of the model, n_p and n_z , are chosen based on the observed resonances and antiresonances in the Bode plot. This means the resonance frequencies $\omega_{n,i}$ and $\omega'_{n,i}$ must fall within the frequency domain of the measurement. Lastly, for the quality factors q_i and q'_i , the lower bound of the is set to at least 0.5 to maintain the numerical stability. However, the upper bound cannot be obtained easily because it depends on the height/depth of the resonance peaks/notches. For a single oscillator, the quality factor is approximately equal to the peak amplification from the static gain. In practice, the quality factors of mechanical resonances in a suspension are found to be below the order of 10^{-3} . Therefore, $[0.5, 1000]$ is a reasonable initial boundary for the quality factors. The boundaries can always be adjusted when the optimization yields sub-optimal results.

With the boundaries set, the transfer function fitting problem can be solved using a global optimization method such as differential evolution. However, in practice, using global optimization is often a trial-and-error process as the optimization traps in a local minima. This result is that only part of the transfer function is well-modeled. In this case, the optimization needs to be restarted with a different random number generator seed so there is a chance that it gives a better results. But, there is no guarantee that the optimization can yield a satisfactory fit so it could cost a lot of trials to obtain one fit. Therefore, global optimization may not be the best way to solve this problem. Instead, local optimization is used more often in practice and it yields satisfactory results quicker with a user-informed initial guess.

To use local optimization, one needs to initialize the model parameters so the algorithm can start a local search. In the context of transfer function fitting, this corresponds to

a manual fit of the frequency response, which is already a typical practice at KAGRA. Although this is also a trial-and-error process, it is straightforward as the parameters can easily be guessed via a quick visual inspection or even algorithmically obtained. The model parameters that needed to be initialized are the static gain k_{DC} , the resonance frequencies $\omega_{n,i}$ and $\omega'_{n,i}$ and the quality factors q_i and q'_i . First of all, as discussed previously, the static gain k_{DC} corresponds to the magnitude response at 0 Hz. If the measured magnitude response shows a plateau at low frequency, the magnitude of the plateau, or simply the magnitude at the lowest frequency, can be used as an initial estimation of k_{DC} . Then, one resonance frequency and one quality factor constitute a complex pole/zero and they must be guessed in pairs. Estimating all the resonance frequencies and quality factors of a complex pole/zero is a straightforward iterative process as described by Algorithm 2.

As stated in Algorithm 2, the transfer function model $\hat{P}(s)$ is initialized with the static gain k_{DC} , which is obtained previously. In order to obtain the resonance frequencies and the quality factors, the first step is to scan the magnitude response, from low frequency to high frequency, for either a peak or a notch. If a peak is found, then one needs to guess a quality factor q . After that, the measured frequency response $P(j\omega)$ is plotted together with $\hat{P}(j\omega) \times \omega_n^2 / [(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2]$ and compare, where ω_n is the frequency at which the peak was found. If the peak heights match, then the model is changed to $\hat{P}(j\omega) \times \omega_n^2 / [(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2]$ and the resonance frequency and quality factor are marked as two of the initial parameters. Otherwise, adjust the quality factor q and try again. Similarly, the same procedure can be done for a notch. But instead of $\hat{P}(j\omega) \times \omega_n^2 / [(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2]$, $\hat{P}(j\omega) \times [(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2] / \omega_n^2$ is instead compared with the measured response. The process repeats until all peaks and notches, i.e. the whole frequency response, are coarsely modeled. This gives all initial parameters for k_{DC} , $\omega_{n,i}$, $\omega'_{n,i}$, q_i , and q'_i .

Now, Algorithm 2 assumes that the multiplying a pair of complex pole or a pair of complex zero at higher frequency will not change the height or depth of a peak or notch at lower frequency. This holds true only when the frequency of the pair of complex poles/zeros is sufficiently high compared to the frequency of the peak or the notch. When there are multiple peaks and notches localized in a small frequency range, they cannot be individually treated as in Algorithm 2. Instead, one needs to go back and forth to adjust the quality factors to match the shape of the cluster of peaks and notches.

SRM F1 GAS filter frequency response

To exemplify the methods described above, the modeling of frequency response of the SRM F1 GAS filter is given in this section.

The magnitude of the frequency response of the SRM F1 GAS filter degree of freedom is shown in Fig. 6.19. First of all, it should be noted that the frequency response is inaccurately presented at frequencies above ~ 4 Hz. The notch at 60 Hz corresponds

Algorithm 2 Estimating resonance frequencies and Q factors from a frequency response

$\omega_i \leftarrow \{\omega_1, \omega_2, \dots, \omega_N\}$ ▷ The frequency data points.
 $P(j\omega) \leftarrow \{P(j\omega_1), P(j\omega_2), \dots, P(j\omega_N)\}$ ▷ The frequency response data complex series.
 $k_{DC} \leftarrow |P(j\omega_1)|$ ▷ The static gain.
 $n_p \leftarrow$ Number of peaks in $|P(j\omega)|$ ▷ Estimate the number of resonances.
 $n_z \leftarrow$ Number of notches in $|P(j\omega)|$ ▷ Estimate the number of antiresonances.
 $\hat{P}(s) \leftarrow k_{DC}$ ▷ Initialize a transfer function model.
 $\omega_{n,i} \leftarrow \{\}$ ▷ List of resonance frequencies.
 $\omega'_{n,i} \leftarrow \{\}$ ▷ List of antiresonance frequencies.
 $q_i \leftarrow \{\}$ ▷ List of resonance quality factors.
 $q'_i \leftarrow \{\}$ ▷ List of antiresonance quality factors.

for $i = 1, i \leq N, i++$ **do** ▷ Scan from low to high frequency
 if $|P(j\omega_i)|$ is a peak **then** ▷ A peak is found.
 $\omega_n \leftarrow \omega_i$
 Guess q . ▷ Guess a quality factor for this resonance.
 Plot $P(j\omega)$ and $\hat{P}(j\omega) \times \frac{\omega_n^2}{(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2}$ ▷ Inspect the Bode plots.
 if Peak height matches **then**
 Add ω_n and q to the lists $\omega_{n,i}$ and q_i .
 $\hat{P}(j\omega) \leftarrow \hat{P}(j\omega) \times \frac{\omega_n^2}{(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2}$ ▷ Multiply the model with the poles.
 else ▷ Inspection not satisfied.
 Increase or decrease q and repeat inspection.
 end if
 else if $|P(j\omega_j)|$ is a notch **then** ▷ A notch is found instead.
 $\omega_n \leftarrow \omega_i$
 Guess q . ▷ Guess a quality factor for this antiresonance.
 Plot $P(j\omega)$ and $\hat{P}(j\omega) \times \frac{(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2}{\omega_n^2}$ ▷ Inspect the Bode plots.
 if Notch depth matches **then**
 Add ω_n and q to the lists $\omega'_{n,i}$ and q'_i .
 $\hat{P}(j\omega) \leftarrow \hat{P}(j\omega) \times \frac{(j\omega)^2 + \omega_n(j\omega)/q + \omega_n^2}{\omega_n^2}$ ▷ Multiply the model with the zeros.
 else ▷ Inspection not satisfied.
 Increase or decrease q and repeat inspection.
 end if
 end if
end for

to a notch filter at the input filter, which is used to filter out the 60 Hz coupling from the electricity mains. The magnitude response with $\sim f^2$ frequency dependency at above ~ 10 Hz is not part of the mechanical frequency response, which should have a frequency dependency of $\sim f^{-2}$, either. Instead, it is due to magnetic coupling between the coil-magnet actuator and the LVDT coils reading the displacement of the GAS filter. The notch at 8 Hz is a result of the $\sim f^2$ magnetic coupling being added to the $\sim f^{-2}$ mechanical response so it should be ignored also. In order to neglect these artifacts, the weighting function $W(\omega)$ is used to mask the data. In this case, the weighting function




Figure 6.19: The magnitude response of the SRM F1 GAS filter.

is set to 0 above 3 Hz and is set to 1 below that.

The corresponding transfer function model is given by Eqn. (6.70) as discussed. Judging from the number of peaks and notches and ignoring those above 3 Hz, the number of complex-pole pairs is $n_p = 3$, and that of complex-zero pairs is $n_z = 2$. This means there are 11 model parameters (1 DC gain, 5 resonance frequencies, and 5 quality factors) to be fitted. Following Algorithm 2, all parameters are identified in 5 iterations, each corresponding to an identification of a pair of complex pole or a complex zero. The initial parameter values identified are shown in the second column in Table 6.2 and the model corresponding to iteration 1, 3, and 5 are shown in Fig. 6.20. In this example, the identification was done via visual comparison between the magnitude response of the model and the measured data. Here, the model does not need to fit perfectly with the data because of the optimization step that comes after this.

With the initial parameters obtained, the transfer function model can be obtained using a local optimization method. In this case, the model was fitted using `scipy.optimize.minimize(fun, x0, method="Nelder-Mead", options={"adaptive":True, maxiter=len(x0)*1000})` in Python, where `fun` is the cost function and `x0` is the initial parameters. Note that it is the logarithm of the parameters in Table 6.2 being optimized, not the model parameters themselves. The optimization was terminated after 1771 iterations, with a final cost function of -0.1725. In comparison, the cost function is evaluated to -0.1492 with the initial guess. The result is shown in Fig. 6.21 and the fitted parameters are shown in the third column of Table 6.2. As can be seen, the optimized model fits well to the frequency response and is an improvement of the initial fit.

Parameter	Initial guess	Locally optimized	Globally optimized
k_{DC}	3.955×10^{-2}	3.761×10^{-2}	3.761×10^{-2}
$\omega_{n,1}$	$0.27 \times 2\pi$	$0.2771 \times 2\pi$	$0.2771 \times 2\pi$
$\omega_{n,2}$	$0.6 \times 2\pi$	$0.5959 \times 2\pi$	$0.5959 \times 2\pi$
$\omega_{n,3}$	$1.1 \times 2\pi$	$1.078 \times 2\pi$	$1.078 \times 2\pi$
$\omega'_{n,1}$	$0.4 \times 2\pi$	$0.4111 \times 2\pi$	$0.4111 \times 2\pi$
$\omega'_{n,2}$	$0.7 \times 2\pi$	$0.7098 \times 2\pi$	$0.7098 \times 2\pi$
q_1	20	24.85	24.83
q_2	10	17.36	17.38
q_3	10	7.428	7.426
q'_1	20	23.36	23.10
q'_2	50	9.742	9.755

Table 6.2: Model parameters identified for modeling SRM F1 GAS frequency response.




Figure 6.20: Iterations 1 (orange dashed), 3 (green dash-dot), 5 (red dotted) of model parameter initialization for the SRM F1 GAS filter frequency response.

As a comparison, the global optimization method was also used to perform the fitting and the results are shown in Fig. 6.22. The global optimization was performed multiple times using different random seeds. The boundary of the for the global optimization was set using the aforementioned criteria. And, the optimization was achieved in Python using `scipy.optimize.differential_evolution(func, bounds, workers=-1, updating="deferred")`, where `func` is the cost function and `bounds` specifies the boundary of the parameters. As shown in Fig. 6.22, the fit was only successful when the random number seed equals to 4 and this illustrates the aforementioned discussions about using a global method where the fit is done partially. The cost function evaluates to -0.1593, -0.1603, and -0.1725 for seed equals to 2, 3, and 4, respectively. For the successful case,




Figure 6.21: SRM F1 GAS filter frequency response modeling using local optimization. Blue: Measured frequency response. Orange dashed: Initial guess for the local optimization. Green dash-dot: Transfer function model fitted using local optimization.

the evaluated cost function is practically identical to that of the local optimization. The optimized parameters are shown in the fourth column in Table 6.2. As can be seen, the parameters are practically identical to that of the locally optimized ones (with 4 significant figures.).




Figure 6.22: SRM F1 GAS filter frequency response modeling using global optimization. Blue: Measured frequency response. Orange dashed: Model fitted using global optimization with random seed = 2. Green dash-dot: Model fitted using global optimization with random seed = 3. Red dotted: Model fitted using global optimization with random seed = 4.

6.3.2 Frequency spectrum modeling

The modeling of frequency spectrum can be straightforwardly done by minimizing a cost function:

$$\min_{\theta} \left(\frac{1}{N} \sum_{i=1}^N \left[\log \hat{S}(f_i; \theta) - \log S(f_i) \right]^2 W(f_i) \right), \quad (6.71)$$

where $S(f_i)$ is the measured frequency spectrum, f_i is the frequency data points, N is the number of data points, $\hat{S}(f_i; \theta)$ is the model of the frequency spectrum, θ is the array of model parameters, and $W(f_i)$ is the weighting function. In this case, the plant model and the frequency response are replaced with an empirical model of the spectrum and the measured spectrum, respectively. An example of an empirical model is something like

$$\hat{S}(f; a, b, N_a, N_b) = \left[\left(\frac{N_a}{f^a} \right)^2 + \left(\frac{N_b}{f^b} \right)^2 \right]^{\frac{1}{2}}, \quad (6.72)$$

where a , b , N_a , and N_b are the model parameters. An empirical model in this form corresponds to a quadrature sum of two straight line spectrum with two frequency dependencies a and b . It is very useful for describing the amplitude spectral density (ASD) of sensor noises such as the geophone noise and the LVDT noise.

In some applications, it is useful to model the amplitude spectral density of a signal

with the magnitude response of a transfer function. This corresponds to modeling the signal as an output of a system that has a white noise input with unity amplitude spectral density. This is particularly useful in applications such as \mathcal{H}_∞ optimization, which is to be discussed in Part III. In this case, the modeling is not as straightforward as modeling the frequency response of the suspensions since a transfer function model cannot be well defined. This is because the frequency dependency of the ASD of a signal can have fractional order, such as $1/f^{0.5}$, whereas a transfer function model only has integer order, which corresponds to the number of poles and zeros. It is possible to represent a fractional order system with a higher integer order transfer function. However, this only works within a certain frequency range.

Moreover, even if the required order of the transfer function is known, it is still difficult to define a transfer function model. Recall, from Eqn. (6.65), that a generic transfer function model can be defined as

$$\hat{P}(s) = \frac{b_0 + b_1s + b_2s^2 + \cdots + b_ms^m}{a_0 + a_1s + a_2s^2 + \cdots + a_ns^n}, \quad (6.73)$$

where n and m denote the order of the denominator and numerator polynomials, respectively. Again, directly using this as the model for fitting is an ill-conditioned optimization as it gets numerically unstable. Assuming that numerical stability is not an issue, there are still other problems: 1) Using the model for a local optimization directly requires a good initial guess, which is difficult to obtain, and 2) using the model for a global optimization can be difficult as the boundaries of the coefficients a_i and b_i cannot be well defined. The model can be factorized into

$$\hat{P}(s) = k \frac{\prod_{i=1}^{n_{sz}} s - z_i \prod_{i=1}^{n_{cz}} s^2 + \frac{\omega'_{n,i}}{q_i} s + \omega_{n,i}^{\prime 2}}{\prod_{i=1}^{n_{sp}} s - p_i \prod_{i=1}^{n_{cp}} s^2 + \frac{\omega_{n,i}}{q_i} s + \omega_{n,i}^2}, \quad (6.74)$$

where z_i and p_i are the simple zeros and simple poles, $\omega_{n,i}$, $\omega'_{n,i}$, q_i , and q'_i are the resonance frequencies and quality factors of the resonances and the antiresonances as described previously, and n_{sz} , n_{sp} , n_{cz} , and n_{cp} are the numbers of simple zeros, simple poles, pairs of complex zeros, and pairs of complex poles, respectively. From this model, it is obvious that $n_{sz} + 2n_{cz} = m$ and $n_{sp} + 2n_{cp} = n$ but there are no more information that can further be used to deduce the individual numbers of simple zeros and poles and the numbers of complex zeros and poles. This makes model selection difficult for spectrum modeling.

To make frequency spectrum modeling work, a 2-step method is proposed. The first step is to use a zero-pole-gain (ZPK) model to coarsely model the frequency spectrum. The ZPK model is defined as

$$\hat{P}_{\text{ZPK}}(s; z_i, p_i, k) = k \frac{\prod_{i=1}^{n_{sz}} s - z_i}{\prod_{i=1}^{n_{sp}} s - p_i}. \quad (6.75)$$

A model in this form can be used to fit the frequency spectrum using a global optimization algorithm. This is because the required order of the transfer function model can be used to set n_{sz} and n_{sp} directly so the model becomes defined. The values of z_i and p_i are also supposedly bounded within the frequency range of the measurement, or a known frequency band of interest. This makes global optimization possible. However, the ZPK model is not a generic form of a transfer function so modeling the frequency spectrum with a ZPK model is only an intermediate step. The second step is to convert the ZPK model into a polynomial form, using the globally optimized values of simple zeros and poles. This gives initial values of a_i and b_i for the generic transfer function model in Eqn. (6.73), which can then be proceeded for a local optimization. The local optimization should be numerically stable as the ZPK model should provide a sufficiently good fit close to a local minimum.

As mentioned previously, the modeling of a fractional-order frequency spectrum using a transfer function model only works at within a frequency band of interest. This means the frequency spectrum at outside the frequency band is not modeled. This fact can be used to simplify the modeling problem if the magnitude response of the transfer function model is allowed to be flat outside the frequency band of interest. In this case, the numbers of simple zeros and simple poles selected are equal, i.e. $n_{sz} = n_{sp}$, during the intermediate step. The selection of this number is not obvious. In principle, a higher order transfer function could mean a better fit of the frequency spectrum. But, in practice, this usually gives zero-pole cancellation pairs outside the frequency band of interests, indicating that the numbers of zeros and poles are in excess. As a general rule of thumb, the number can be initiated by the minimum integer that is higher than the fractional order of the frequency spectrum. For example, if the frequency spectrum has a $f^{-3.5}$, then the order is initially chosen to be 4. The order can be increased when the fit is not satisfactory, or it can be decreased when zero-pole cancellation starts to occur.

It is also worth mentioning that the transfer function model obtained using local optimization might represent a non-minimum phase system. This means that the transfer function model is stable but its inverse is not, i.e. it has at least one zero with positive real part. In applications where the inverse of the model needs to be invoked, the model must be converted to a minimum phase one. This can be achieved by simply negating the positive real parts of the zeros. This will not change the magnitude response of the transfer function model. Only the phase response is changed, which is not important for modeling a frequency spectrum.

SRM IP longitudinal geophone noise spectrum modeling

To illustrate the frequency spectrum modeling method, an example showing the modeling of a geophone noise spectrum is shown below. Fig. 6.23 shows the SRM pre-isolator geophone noise spectrum (Blue solid) in the longitudinal direction. The empirical model

(Orange dashed), ZPK model (Green dash-dot), and the transfer function model (Red dotted) for this spectrum are also shown. The three geophones at the pre-isolator stage circles around the pre-isolator in a way like the LVDTs as described in Sec. 6.2.1. The effective longitudinal noise spectrum is obtained by multiplying the geophone spectra in the sensor basis by the corresponding longitudinal row elements of the sensing matrix. Then, they are quadrature summed to obtain one spectrum corresponding to the noise spectrum in the longitudinal direction.




Figure 6.23: Geophone noise spectrum. Blue solid: Measured noise spectrum. Orange dashed: Empirical model. Green dash-dot: ZPK model. Red dotted: Transfer function model.

In Fig. 6.23, three models are used to fit the measured noise spectrum, including the empirical model (6.72), the ZPK model (6.75), and the transfer function model (6.73). In all cases, the models are obtained by minimizing the cost function in Eqn. (6.71). For the ZPK model and transfer function model, only the magnitude response is used, instead of the complex-valued frequency series. The weighting function is chosen to be 1 at $f < 0.1$ and $f > 1$, and is 0 between $0.1 < f < 1$ due to the observable residual resonances of the suspension. The empirical model is obtained using a local optimization algorithm same as that used in Sec. 6.3.1. The initial parameters are chosen to be $\theta = \{a, b, N_a, N_b\} = \{10^{-5.46}, 10^{-5.23}, 3.5, 1\}$. These values were obtained using a graphical method (i.e. calipers) on the typical geophone intrinsic noise shown in Ref. [29]. And, the optimized empirical model is shown as orange dashed curve in Fig. 6.23.

The ZPK model and the transfer function model are obtained by fitting the frequency responses of the models to the optimized empirical model, not the measured data. The benefit of obtaining an empirical model beforehand is that it enables further fitting with

log-spaced frequency data, which helps normalizing the significance of data point at different frequencies. It also filters out unwanted data in the measured spectrum. The ZPK model is obtained first and corresponding model is a ZPK model with 4 simple zeros and 4 simple poles. The order 4 is chosen based on the fact that the geophone noise has a frequency dependency of $\sim f^{-3.5}$ at low frequencies. The optimized ZPK model is obtained straightforwardly with a global optimization algorithm same as that used in Sec. 6.3.1. The magnitude response of the ZPK model is plotted in green dash-dot in Fig. 6.23. As can be seen, it follows the trend of the empirical model well but it fluctuates around the empirical model.

The optimized ZPK model is used as the initial guess of the local optimization for fitting the transfer function model. Because the ZPK model is already close to the empirical fit, the local optimization using a transfer function model did not result in an unstable numerical optimization. The frequency response of the transfer function model obtained using local optimization is shown as red dotted line in Fig. 6.23. It overlaps with the empirical model almost perfectly and is a much better fit compared to the ZPK model.

6.4 Coarse alignment and local damping control

To begin the section, it is worth recalling the control block diagram of one single degree of freedom as shown in Fig. 5.5. In this section, the design of the controller $K(s)$ is discussed. In particular, the discussion for damping control and coarse alignment control is provided. The controller discussed in this section serves the purpose for initially controlling and aligning the interferometer. It stabilizes the displacement of the optics and it roughly aligns the optics, putting the interferometer into the operating point so interferometer control can be further engaged. These controllers are required to be simple and robust. The design of the controllers does not take into account the sensing noise. Therefore, the controller discussed here is not guaranteed to be valid for the observation mode where noise performance at higher frequency becomes crucial. With that said, the noise performance is largely determined by the sensing noise, as the open-loop gain is typically low at observation band. Therefore, the controls at degrees of freedom at lower stages are usually disengaged during observation since the control noise are less attenuated passively by the pendulums.

The design of the controller is typically done manually by experts by shaping the open-loop frequency response. They do so by putting required control bandwidth above unity gain and making sure that the system is stable according to simple stability criteria. This is usually done according to personal experience and it leads to inconsistency across different systems. For instance, one may require a stability phase margin of 30° while others may require 45° . Moreover, this way of designing of the controller lacks reproducibility as the ideology is based on personal experience and is not quantified ob-

jectively. The controllers should not be copied over other systems as there are subtle differences, e.g. slightly different resonance frequencies that might prevent that. To solve the issues above, this section seeks to provide an algorithmic way to shape the feedback controllers according to the modeled plant of the suspension.

There are three types of filters discussed in this section. Sec. 6.4.1 discusses the design of a derivative controller for critically damping the resonances of the suspension. Sec. 6.4.2 discusses an augmentation to the damping controller for position control. And, Sec. 6.4.3 discusses additional post-filter techniques that can help attenuating unwanted high frequency noise and improving stability.

6.4.1 Damping control

For simplicity, consider a single pendulum shown in Fig. 5.4. The transfer function from the actuation on the pendulum to the displacement of the pendulum has the form

$$P(s) = k_{\text{DC}} \frac{\omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2}, \quad (6.76)$$

where k_{DC} is the static actuation efficiency, ω_n is the resonance frequency, and q is the quality factor. Now, suppose the disturbance $D(s)$ is the perturbation of the displacement $X(s)$ due to the motion of the platform $X_{\text{platform}}(s)$ where the pendulum is suspended from. The disturbance reads

$$D(s) = \frac{\frac{\omega_n}{q}s + \omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2} X_{\text{platform}}(s), \quad (6.77)$$

where $X_{\text{platform}}(s)$ is the platform displacement. Neglecting sensing noise, the displacement of the pendulum under feedback control reads

$$\begin{aligned} X(s) &= \frac{1}{1 + K(s)P(s)} D(s) \\ &= \frac{1}{1 + K(s)k_{\text{DC}} \frac{\omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2}} \frac{\frac{\omega_n}{q}s + \omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2} X_{\text{platform}}(s) \\ &= \frac{\frac{\omega_n}{q}s + \omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2 + K(s)k_{\text{DC}}\omega_n^2} X_{\text{platform}}(s), \end{aligned} \quad (6.78)$$

where the expression in the first line is given by Eqn. (5.4), and in the second line, $P(s)$ and $D(s)$ are substituted by Eqn. (6.76) and Eqn. (6.77), respectively. It is obvious that when $K(s) = 0$, the denominator is composed of a pair of complex poles, which represents the resonance of the pendulum. The magnitude response has a peak around $\omega = \omega_n$. This amplifies the motion of the platform $X_{\text{platform}}(s)$ and is obviously not desirable.

Simply judging from Eqn. (6.78), it is obvious that when $K(s)$ is large, i.e. $K(s) \rightarrow \infty$, the pendulum displacement $X(s)$ gets decoupled from the motion of the platform.

However, this is not feasible due to two reasons: 1) The displacement will be dominated by sensing noise in reality, and 2) The actuation signals may get saturated, which is something that needs to be avoided as it causes actuation coupling to other degrees of freedom or even causing large oscillations. A compromise is to set $K(s)$ such that the dynamics of the system contains no resonance, hence no oscillation and not amplifying any external disturbance. Although this does not guarantee that the actuation signals will not be saturated, it provides a baseline for setting the controller that can be tuned in practice.

To achieve a closed-loop system with no oscillation, the controller $K(s)$ needs to lower the quality factor in the closed-loop condition. The inverse of the quality factor is a coupling constant to the Laplace variable s in the denominator. This requires the controller to at least contain a factor proportional to the Laplace variable s . And, without changing the resonance frequency of the system, the controller is exactly proportional to s , giving

$$K(s) = k_d s, \quad (6.79)$$

where k_d is the derivative gain. The system is said to be critically damped when the quality factor is 0.5⁶. This is because a quality factor of 0.5 corresponds to a system with a transfer function $1/(s^2 + 2\omega_n s + \omega_n^2)$, which has exactly two simple poles at $s = -\omega_n$, corresponding to two exponential decays with no oscillation. To achieve this, it is easy to see from Eqn. (6.78) that when

$$\frac{\omega_n}{q} + k_d k_{\text{DC}} \omega_n^2 = 2\omega_n, \quad (6.80)$$

the system is critically damped. This gives the critical derivative gain

$$k_d = \frac{2 \left(1 - \frac{1}{q}\right)}{k_{\text{DC}} \omega_n}. \quad (6.81)$$

For degrees of freedom with only one resonance, the critical damping controller (6.79) can be straightforwardly implemented by substituting the derivative gain by Eqn. (6.81). The parameters k_{DC} , ω_n , and q in Eqn. (6.81) are exactly those obtained from system identification as discussed in Sec. 6.3.1.

However, most degrees of freedom in a suspension have more than one resonance, recall that the suspension is a multiple-pendulum. For these systems, the critical derivative gain cannot be easily derived analytically. Instead, a numerical approach is used to obtain the critical derivative gain. In this case, the critical derivative gain is defined such that one of the closed-loop complex-pole pairs turn into a pair of simple poles. To visualize this,

⁶A quality factor of $\sqrt{2}/2$ is also a reasonable choice. This gives frequency response containing no local maximum at the resonance frequency. Therefore, anywhere between 0.5 and $\sqrt{2}/2$ satisfies the no amplification requirement.




Figure 6.24: Root locus plot for SRM F1 GAS filter derivative control. Solid lines: root loci of the closed-loop poles. Crosses: Open-loop poles corresponding to the starting points of the root loci. Circles: Open-loop zeros corresponding to the end points of the root loci.

consider a root locus plot as shown in Fig. 6.24. The root locus plot shows the locations of the poles of a closed-loop system as a parameter varies. In this example, the SRM F1 GAS filter transfer function identified in Sec. 6.3.1 is used as the plant $P(s)$ and the parameter is the derivative gain k_d . In Fig. 6.24, the cross markers represent the poles of the open-loop transfer function $K(s)P(s)$ and the circle markers represent its zeros. At $k_d = 0$, the closed-loop poles are the open-loop poles. As k_d increases, each closed-loop pole follows one of the root locus (branches) and travels towards the open-loop zero. As $k_d \rightarrow \infty$, the closed-loop poles end at the open-loop zeros. As shown in Fig. 6.24, the first branch (blue) and the second branch (orange), corresponding to a pair of poles, meet on the real axis between -6 and -4. The point where they meet corresponds to a break-in point and is where the two complex-pole pairs turn into two simple poles, and there will be no oscillation corresponding to this mode. This is where the k_d should be set for a critically damped system. The other poles, however, never meet on the real axis and go toward the closest complex zeros directly. This is typical for systems in a suspension, which have transfer functions with a relative order of two, i.e. there are two more poles than zeros.

The root locus analysis showed that one pair of complex poles can be turned into a

Algorithm 3 Finding the critical derivative gain using a bisection algorithm.

$P(s) \leftarrow k \frac{\prod_{i=1}^{n_z} s^2 + \frac{\omega'_{n,i}}{q'_i} s + \omega'^2_{n,i}}{\prod_{i=1}^{n_p} s^2 + \frac{\omega_{n,i}}{q_i} s + \omega^2_{n,i}}$ ▷ The identified plant.
 $\underline{k}_d \leftarrow \frac{1}{\max_{\omega} |j\omega P(j\omega)|}$ ▷ The lower limit of the derivative gain
 $K(s) \leftarrow k_d s$ ▷ The controller.
 $C(s) \leftarrow \frac{K(s)P(s)}{1+K(s)P(s)}$ ▷ The closed-loop transfer function.
 $\epsilon \leftarrow 0.1$ ▷ A multiplicative increment. For example, 0.1.
 $\bar{k}_d \leftarrow \underline{k}_d$ ▷ The upper limit of the derivative gain, to be found.

Find the upper limit of the derivative gain:

while C(s) only has complex poles **do**

$\bar{k}_d \leftarrow \bar{k}_d (1 + \epsilon)$ ▷ Increment the derivative gain until it exceeds the critical

$K(s) \leftarrow \bar{k}_d s$

$C(s) \leftarrow \frac{K(s)P(s)}{1+K(s)P(s)}$

end while

Bisection algorithm for the critical damping gain:

tolerance $\leftarrow 10^{-6}$ ▷ An acceptable tolerance for termination. For example, 10^{-6} .

while $\frac{\bar{k}_d - \underline{k}_d}{\underline{k}_d} > \text{tolerance}$ **do**

$k'_d \leftarrow 10^{\frac{\log \bar{k}_d + \log \underline{k}_d}{2}}$ ▷ A trial gain the logarithmic mean of between the boundaries.

$K(s) \leftarrow k'_d s$

$C(s) \leftarrow \frac{K(s)P(s)}{1+K(s)P(s)}$

if C(s) only has complex poles **then**

$\underline{k}_d \leftarrow k'_d$

else

$\bar{k}_d \leftarrow k'_d$

end if

end while

$k_d \leftarrow k'_d$

return k_d

▷ The critical derivative gain.

pair of simple poles, indicating a critically damped condition. However, it is not obvious from the plot that what derivative gain k_d should be set. And even if there is, the critical derivative gain is only good to an accuracy of the grid size of the gain used to plot the root locus. Instead, a numerical approach is proposed to obtain the critical damping gain, as shown in Algorithm 3. In the algorithm, the lower limit \underline{k}_d and upper limit \bar{k}_d of the derivative gain is first obtained. And then, a bisection algorithm is used to tighten these limits until a convergence condition is met.

The lower limit is defined by

$$\underline{k}_d = \frac{1}{\max_{\omega} |j\omega P(j\omega)|}. \quad (6.82)$$

This lower limit is the lowest gain such that the open-loop gain $|K(j\omega)P(j\omega)|$ is below unity at all frequencies, which means there is zero control bandwidth. The critical derivative gain should be higher than this value so it is a reasonable way to set this lower limit. As for the upper limit, it is found by an iteration. The upper limit \bar{k}_d is first set to the lower limit \underline{k}_d . Then, in each step of the iteration, \bar{k}_d gets multiplied by a factor of $(1 + \epsilon)$, where ϵ is a small number and is set to 0.1 in Algorithm 3 as an example. The gain \bar{k}_d is used to test the closed-loop transfer function to see if the gain is high enough such that the two complex poles turned into two simple poles. The process repeats until this happens and the upper limit \bar{k}_d is obtained.

Now, the critical derivative gain is known to fall in the interval $(\underline{k}_d, \bar{k}_d)$. Therefore, a numerical method like the bisection algorithm can be used to find the critical derivative gain. Strictly speaking, the bisection algorithm is an algorithm for finding the root of a function that falls within a certain interval. However, it can be modified to find the critical derivative gain as shown in the second while-loop in Algorithm 3. In each step of the iteration, a trial derivative gain k'_d is set to the logarithmic mean of between $(\underline{k}_d, \bar{k}_d)$. Practically, setting the trial to the half-way between the limits also works but the logarithm is chosen based on the fact that Bode plots are typically plotted in log-scale. The trial gain is then used to evaluate the closed-loop poles. If the closed-loop poles still have complex values, then the lower limit \underline{k}_d is set to the trial gain and the process repeats. Otherwise, the upper limit \bar{k}_d is set to the trial gain and the process repeats. The iteration is terminated when $(\bar{k}_d - \underline{k}_d)/\underline{k}_d$ is smaller than a user-defined tolerance, i.e. when they become close enough.

SRM F1 GAS filter critical damping control

To illustrate, algorithm 3 was used to obtain the critical derivative gain for the SRM F1 GAS filter degree of freedom. The transfer function of the system is given by Eqn. (6.66) and the identified parameters are those given in Table 6.2. Using algorithm 3, the critical derivative gain was found to be 18.96 for the SRM F1 GAS filter degree of freedom. As a comparison, using Eqn. (6.81) directly yields a derivative gain of 29.32, which is slightly overdamping but is actually not a bad estimation of the critical derivative gain in log-scale. The frequency response of the open-loop transfer function $K(s)P(s)$ is shown in Fig. 6.25. Here, the dominant resonance is the one at 0.2771 Hz. As shown in the figure, only the open-loop gain around the resonance is above the unity gain, indicating the feedback action only exists around this frequency but not at other frequencies. Other resonances are also put above the unity gain but this is purely coincidence. It is also worth noting that this system is unconditionally stable since the phase never goes lower than -180° .

Fig. 6.26 shows a comparison between the magnitude response, i.e. the ratio between the displacement and the actuation, of the open-loop system and the closed-loop system.




Figure 6.25: Open-loop transfer function of the critical damping control for SRM F1 GAS degree of freedom. (a): Magnitude. (b) Phase.

The frequency response of the closed-loop system is estimated by $P(s)/(1 + K(s)P(s))$. As can be seen, the closed-loop magnitude response has no amplification above the DC gain.

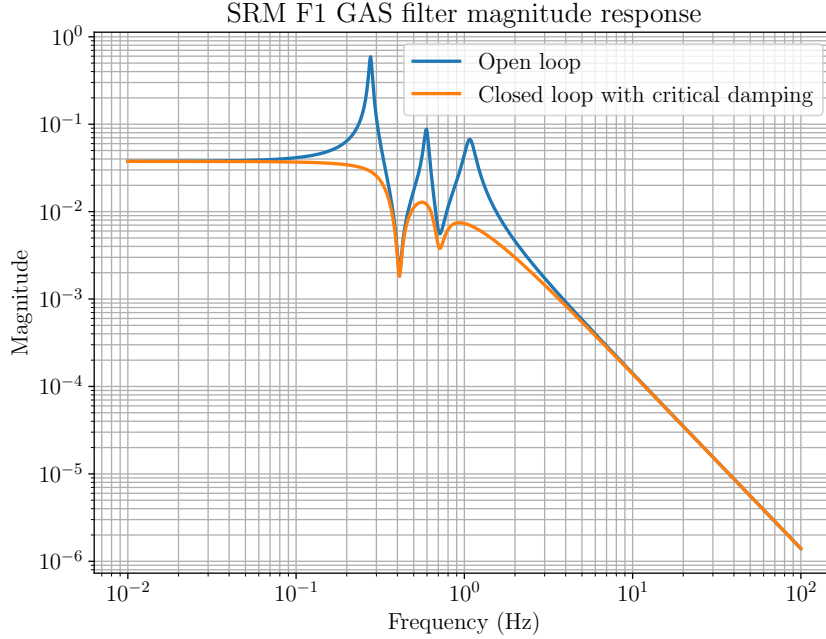


Figure 6.26: SRM F1 GAS filter magnitude response. Blue: Open loop. Orange: Closed loop with critical damping.

6.4.2 Coarse alignment control

The purpose of a position control system is to coarsely align the main optics such that the interferometer is sufficiently aligned. This brings the interferometer to an operating point where control systems using the interferometer sensors can be engaged. Initially, the aligned position of a suspension is found via a manual alignment. The positions are then recorded and set as the setpoint $R(s)$ of the control system as shown in Fig. 5.5. They are static position at DC so the setpoints $R(s)$ are typically a static value, i.e. DC. For the displacement $X(s)$ to reach the setpoint the least requirement is that as time $t \rightarrow \infty$, $[R(s) - X(s)] \rightarrow 0$. Recall the coupling factor of $R(s)$ in Eqn. (5.1), $K(s)P(s)/(1 + K(s)P(s))$, this gives

$$\begin{aligned} R(s) - X(s) &= R(s) - \frac{K(s)P(s)}{1 + K(s)P(s)}R(s) \\ &= \frac{1}{1 + K(s)P(s)}R(s). \end{aligned} \quad (6.83)$$

Note that the disturbance and noise terms in Eqn. (5.1) are ignored in this analysis. According to the final value theorem, as $t \rightarrow \infty$, $R(s) - X(s)$ approaches

$$\lim_{s \rightarrow 0} s[R(s) - X(s)] = \lim_{s \rightarrow 0} s \frac{1}{1 + K(s)P(s)}R(s). \quad (6.84)$$

Let $R(s)$ to be a unit step function $1/s$, then, position control requires

$$\lim_{s \rightarrow 0} \frac{1}{1 + K(s)P(s)} = 0. \quad (6.85)$$

For Eqn. (6.85) to be true, the open-loop transfer function $K(s)P(s)$ must have at least one pole at the origin, i.e. $K(s)P(s)$ must contain a multiplicative factor of $1/s$. While $P(s)$ is the transfer function of the suspension containing only complex zeros and complex poles, position control can only be achieved by having a controller that is proportional to $1/s$. This corresponds to using an integral controller

$$K(s) = \frac{k_i}{s}, \quad (6.86)$$

where k_i is the integral gain.

To see how the integral gain affects the time response, consider the low frequency region where the plant has a flat frequency response. At frequencies below the resonance frequencies, the frequency response of the plant can be estimated as a constant, k_{DC} . In this case, the displacement with integral control can be written as

$$\begin{aligned} X(s) &= \frac{\frac{k_i}{s} k_{DC}}{1 + \frac{k_i}{s} k_{DC}} R(s) \\ &= \frac{k_i k_{DC}}{s + k_i k_{DC}} R(s). \end{aligned} \quad (6.87)$$

Again, letting $R(s)$ to be a unit step function gives

$$\begin{aligned} X(s) &= \frac{k_i k_{DC}}{s + k_i k_{DC}} \frac{1}{s} \\ &= \frac{1}{s} - \frac{1}{s + k_i k_{DC}}, \end{aligned} \quad (6.88)$$

and taking the inverse Laplace transform gives

$$\mathcal{L}^{-1} \{X(s)\} = 1 - e^{-k_i k_{DC} t}. \quad (6.89)$$

Fig. 6.27 shows the time response of this type of system. As can be seen, the displacement is gradually brought to the setpoint following an exponential decay function with a time constant τ . For coarse alignment control, it is desirable for the time constant τ to be as low as possible. This is because the coarse alignment control is used for initial alignment and re-alignment of the interferometer after lock-loss. All these alignment tasks take time off the observation period and accomplishing fast alignment directly translates to reduction of downtime and an improvement in duty cycle.

Now, the suspension control integral control alone is not unconditionally stable, unlike




Figure 6.27: Time response of an integral control system.

the derivative control. This is because the integral term $1/s$ introduces a phase of -90° to the open-loop transfer function. Combined with the -180° phase lag introduced by the resonance of the suspension, it makes integral control possibly unstable for the suspensions. Therefore, one must select the integral gain carefully to avoid instability. For integral control to work alone as a position control system for suspensions, the gain has to be low enough such that the peaks around the resonances in the magnitude response the open-loop transfer function are lower than unity. However, this kind of design lacks efficacy as it would imply the integral gain k_i has to be so low, way lower than the resonance frequency, and hence the time constant $\tau = 1/k_i k_{DC}$ has to be so high. Alternatively, the integral control can work with additional filters such as notch filters and low-pass filters, which can help stabilizing the control. But, this design can be complicated as there are a lot of design freedom.

Fortunately, the integral gain can work together with derivative control such that the derivative action dominates at the resonance while the integral action dominates at low frequency where the frequency response of the plant is flat. Here, a heuristic approach for obtaining a reasonable integral gain is presented. First of all, the derivative gain k_d is selected according to Algorithm 3 in Sec. 6.4.1. After that, record the minimum unity gain frequency $\omega_{\text{ugf},1}$ at $|j\omega_{\text{ugf},1}k_d(j\omega_{\text{ugf},1})P(j\omega_{\text{ugf},1})| = 1$. Now, an integral controller is added in parallel to the derivative controller such that

$$K(s) = \frac{k_i}{s} + k_d s. \quad (6.90)$$

Here, the idea is to have the derivative controller dominating the control bandwidth at

frequencies above $\omega_{\text{ugf},1}$ while the integral control dominates at lower frequencies. One possible realization is to require

$$\left| \frac{k_i}{j\omega_{\text{ugf},1}} k_{\text{DC}} \right| \leq 1. \quad (6.91)$$

The line above assumes a plant with a flat plant at that has a magnitude response of k_{DC} , and sets the integral control unity gain frequency to lower than $\omega_{\text{ugf},1}$. Therefore, for the fastest position control, the integral gain is set to

$$k_i = \frac{\omega_{\text{ugf},1}}{k_{\text{DC}}}, \quad (6.92)$$

where, again, k_{DC} is the DC gain of the plant as shown in Eqn. (6.70). The reason why requirement as shown in Eqn. (6.91) implies that the derivative control dominates at frequencies above $\omega_{\text{ugf},1}$ is due to the fact that the open-loop gain of the derivative control is proportional to ω while that of the integral control is proportional to $1/\omega$. While the critical damping criterion requires the first unity gain frequency $\omega_{\text{ugf},1}$ to be lower than the main resonance frequency, this ensures (for most suspension plants) that the integral control dominates at frequencies below the resonances and the derivative control dominates at frequencies higher than the resonances.

To further improve the transient response of the position control system, a proportional gain can be added so the controller turns into a full proportional-integral-derivative (PID) controller:

$$K(s) = k_p + \frac{k_i}{s} + k_d s, \quad (6.93)$$

where k_p is the proportional gain, k_i is the integral gain, and k_d is the derivative gain. Again, like the integral control, a similar requirement can be set for the proportional gain

$$|k_p P(j\omega_{\text{ugf},1})| \leq 1. \quad (6.94)$$

So for the fastest position control, the proportional control gain is

$$k_p = \frac{1}{|P(j\omega_{\text{ugf},1})|}. \quad (6.95)$$

With that said, the proportional gain is more flexible than the integral gain since proportional control will not lead to instability but only reduce the phase margin, if there is any. Again, all PID gains here are selected based on some heuristic criteria and can be tuned, shall there be other design considerations. Also, the design presented here serves as a baseline that is suitable for most plants in a suspension. Further improvements can always be made on a case-by-case basis.

SRM F1 GAS filter coarse alignment control

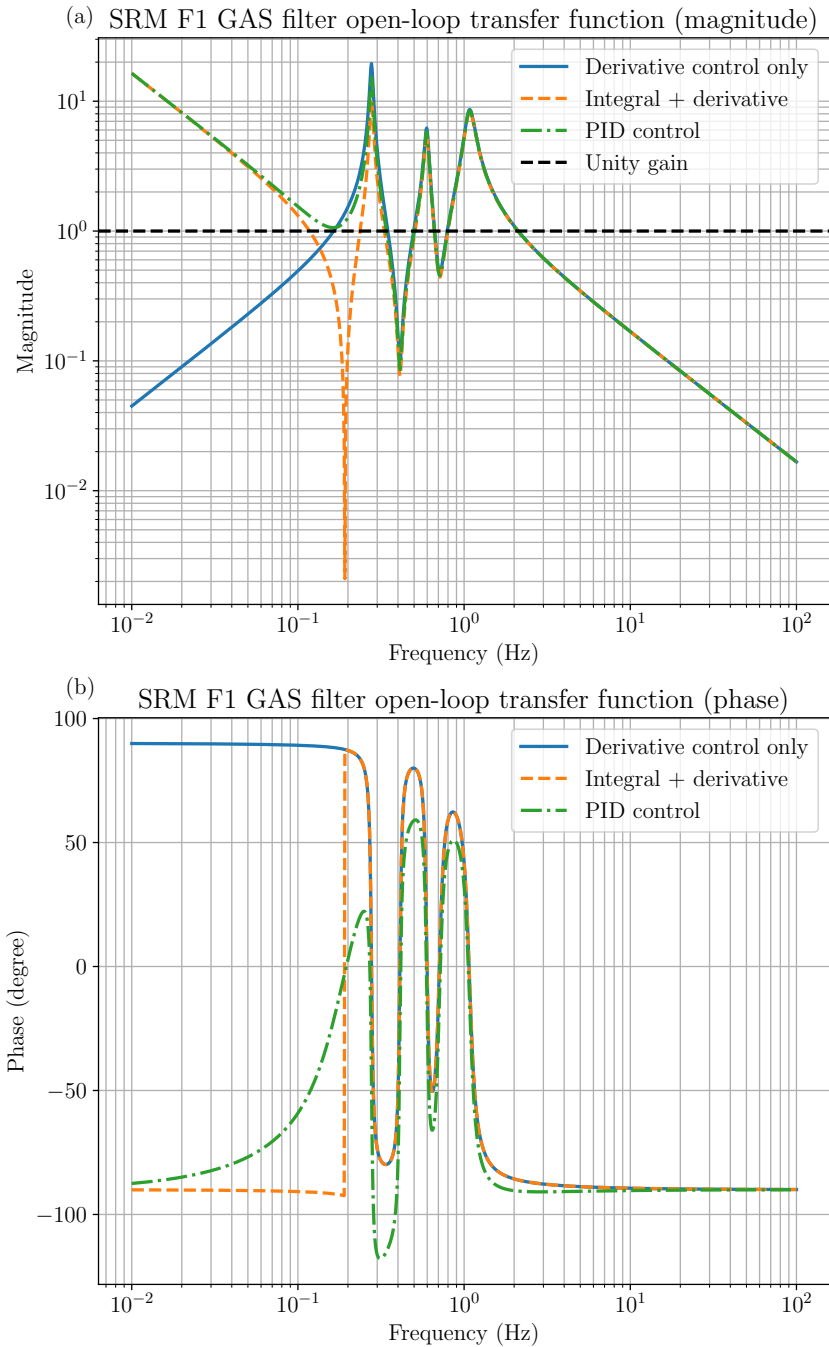


Figure 6.28: Open-loop transfer function of the SRM F1 GAS filter control. (a): Magnitude plot. (b): Phase plot. Blue: Derivative control. Orange dashed: Integral and derivative control. Green dash-dot: PID control.

The integral and proportional gains for the SRM F1 GAS filter degree of freedom were obtained using Eqn. (6.92) and (6.95). Together with the derivative gain obtained in Sec. 6.4.1, three types of control systems are presented here, including derivative control (blue), integral and derivative control (orange dashed), and PID control (green dash-dot). As shown in Table 6.3, the PID gains are 19.58, 27.45, and 18.96 for proportional control,




Figure 6.29: Simulated time response of the SRM F1 GAS filter control. Blue solid: Derivative control. Orange dashed: Integral and derivative control. Green dash-dot: PID control. Red dotted: No control.

integral control, and derivative control respectively. Fig. 6.28 shows the corresponding

k_p	k_i	k_d
19.58	27.45	18.96

Table 6.3: PID gains for the SRM F1 GAS filter control.

open-loop transfer functions $K(s)P(s)$. First of all, from the phase plot (b), it can be seen that the phase response of all open-loop transfer functions are well within $\pm 180^\circ$, indicating all systems are stable. From the magnitude response, it can be seen that a pure integral control with derivative control introduces a notch before first resonance. This correspond to the complex zeros introduced by the controller $k_i/s + k_d s$, which has a pair of complex zeros at $\pm j\sqrt{k_i/k_d}$. This may be fine practically but it implies that the system will not be able to control or suppress any disturbance at that frequency, qualitatively speaking. This issue can be mitigated by the addition of the proportional gain, which puts the open-loop gain at the notch above unity. At higher frequencies, the open-loop gain is almost visually indistinguishable from that of the derivative control, indicating a similar performance. At lower frequencies, the open-loop gain is dominated by the integral control which has a signature $1/f$ frequency dependency.

Fig. 6.29 shows the simulated time response with all three controllers alongside the time response with no control. Since there is no position control in the derivative control and the open loop system, they cannot track the reference signal $R(s)$. Instead, their time response are generated by injecting a step actuation signal with a magnitude of

$1/k_{DC}$, which is the corresponding actuation required to put the displacement at unity. For integral with derivative control and PID control, the time responses are generated by injecting a unit step function to the setpoint $R(s)$, i.e. commanding the system to move the suspension by an amount of unity. This makes the time responses of the three systems comparable.

As shown in Fig. 6.29, all systems have an equilibrium position at $1\ \mu\text{m}$. Compared to the open loop (red dotted), all three controlled system settled around the equilibrium quicker therefore they damped the resonances well. It might seem that the derivative control has superior control performance in terms of rise time and settling time. But, again, the position of the derivative control system is not regulated so these performances are irrelevant. This, however, suggests an interesting control strategy for even faster alignment control, where derivative control and an actuation offset are engaged prior to the position control. For the integral + derivative control (orange dashed) and PID control (green dash-dot), they yielded similar performance. A subtle difference is that the integral + derivative control has a slightly higher overshoot owing to the reduced bandwidth. Nevertheless, they both controlled the position well and can be used for coarse alignment control of the suspension.

6.4.3 Post-filtering

With a well-behaved system, the PID controller in the form of Eqn. (6.93) is typically sufficient for coarse alignment and damping control. However, in most cases, additional filtering is required not just for noise attenuation, but also for stability due to unmodeled dynamics and outlier resonances. These abnormalities are treated with usually two types of filters, low-pass filters and notch filters.

Low-pass filter

Fig. 6.28 and Fig. 6.29 shows the simulated SRM F1 GAS filter control system with the modeled transfer function. No abnormalities were found in these simulations. It might seem that the controllers are ready to be implemented. However, plotting the frequency response of the open-loop transfer function of the PID controller using the measured frequency response (as shown in Fig. 6.19) of the SRM F1 GAS filter degree of freedom yields some interesting results. Fig. 6.30 shows the open-loop transfer function estimated with the measured frequency response (blue) and with the transfer function model. As can be seen, the magnitude responses match very well below the last unity gain frequency at $\sim 2\ \text{Hz}$. However, the magnitude responses above that starts to deviate from each other and the one with measured frequency response started rising drastically above the unity gain at above $\sim 10\ \text{Hz}$. Again, the rising magnitude response is not the mechanical response from the GAS filter. Rather, it is the magnetic coupling between the coil-magnet




Figure 6.30: SRM F1 GAS filter open-loop transfer function. Blue: Open-loop transfer function estimated with measured frequency response. Orange: Open-loop transfer function with transfer function model.

actuators and the LVDT coils (hence it has perfect coherence). While this coupling may not cause instability in theory, this configuration causes the the LVDT coils and actuation signal to saturate, an effect basically indistinguishable from instability. Therefore, the actuation to sensor coupling must be filtered such that the control bandwidth does not extend to the frequency where actuation to sensor coupling dominates.

To achieve this filtering, the a low-pass filter can be added after the control. Or equivalently, the feedback controller can be modified by multiplying the low-pass filter. Now, there are many types of low-pass filters, such as the Chebyshev filter, the Butterworth filter, and the Elliptic filter. For simplicity, the simple low-pass filter $L(s)$

$$L(s; \omega_c) = \left(\frac{\omega_c}{s + \omega_c} \right)^{n_l}, \quad (6.96)$$

where ω_c is the cutoff angular frequency and n_l is the order of the filter, is chosen for further discussion in this section. But, the method discussed here are general for all types of low-pass filter.

The order of the low-pass filter n_l determines how fast the filter rolls off and the cutoff frequency determines at what frequency does the filter begin to roll off. Fig. 6.31 shows examples of low-pass filters up to an order of $n_l = 4$. The selection of the order n_l depends on the purpose of the low-pass filter. For the SRM F1 GAS filter, the minimum order is $n_l = 4$. This is selected is based on the fact that the magnitude coupling introduces unwanted signal that is proportional to f^2 in the sensing readout. The mechanical response

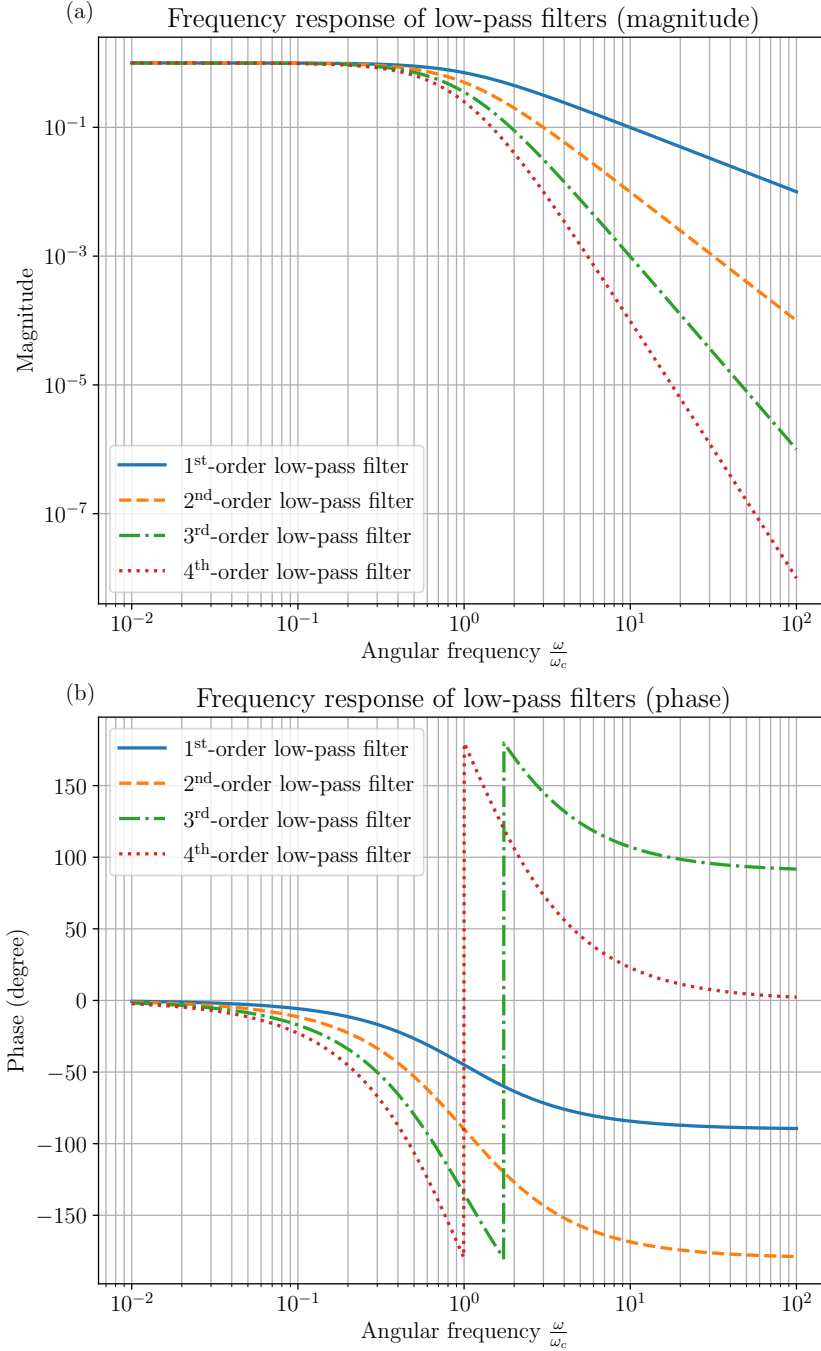


Figure 6.31: Frequency response of low pass filters. (a) Magnitude response. (b) Phase response. Blue solid: 1st-order low-pass filter. Orange dashed: 2nd-order low-pass filter. Green dash-dot: 3rd-order filter. Red dotted: 4th-order low-pass filter. ω_c : cutoff angular frequency.

of the GAS filter has an expected roll-off of f^{-2} , which is that of a spring-mass system. To retain this roll-off, this gives a relative order of 4, which leads to the selection $n_l = 4$.

As for the cutoff frequency ω_c , a choice of around the big notch at ~ 8 Hz seemed reasonable as the magnetic coupling starts to dominate. However, the cutoff frequency ω_c cannot be selected arbitrarily. This is because the low-pass filter introduces negative

phase distortion, as shown in subplot (b) in Fig. 6.31. The low-pass filter shifts the phase by an amount of -90° per n_l , the order of the filter, and the distortion spans over two decades around the cutoff frequency. With the low-pass filter added as part of the controller, this makes instability a possibility as the phase response of the open-loop transfer function gets below -180° . Therefore, the open-loop transfer function must be inspected with additional post-filtering to make sure that, at unity gain frequencies, the phases are above -180° .

Alternatively, since ω_c is a continuous variable, it can be optimized for a targeted phase margin. This way, the cutoff frequency gets as low as possible with guaranteed stability. Algorithm 4 describes a simple 2-step approach to obtain a cutoff frequency such that the open-loop transfer function with the low-pass filter has guaranteed phase margin. The algorithm is similar to Algorithm 3, it first obtain a lower limit and an upper limit of the cutoff frequency, and then, use bisection algorithm to search for the target cutoff frequency.

The first step is to obtain an upper limit of the cutoff frequency $\bar{\omega}_c$ by setting it to some multiple of the highest unity gain frequency of $K(s)P(s)$, the original open-loop transfer function. In Algorithm 4, it is set to 10 times that of the highest unity gain frequency, but in principle, it can be higher. The idea is for $\bar{\omega}_c$ to be high enough such that the phase margin of $K(s)P(s)$ is roughly the same as that of $L(s; \bar{\omega}_c)K(s)P(s)$. Then, the lower limit $\underline{\omega}_c$ is found by an iteration. The lower limit is first initialized by $\underline{\omega}_c = \bar{\omega}_c$. And then, in each step, it is divided by a factor of $1 + \epsilon$, where ϵ is a small number and is set to 0.1 in the algorithm as an example. The minimum phase margin of the open-loop transfer function with the low-pass filter $L(s; \underline{\omega}_c)K(s)P(s)$ is evaluated in each step. When it is lower than the target phase margin ϕ_{target} , the iteration terminates and the lower limit is found. The minimum phase margins corresponding to the lower limit and the upper limit of the cutoff frequency are evaluated as $\underline{\phi}$ and $\bar{\phi}$, respectively.

The target phase margin is believed to lie between $(\underline{\phi}, \bar{\phi})$, so a bisection algorithm can be used to search for the target cutoff frequency. The bisection algorithm tightens the boundaries $(\underline{\omega}_c, \bar{\omega}_c)$ until the phase margins $\underline{\phi}$ and $\bar{\phi}$ are closed enough, such that the fractional difference $(\bar{\phi} - \underline{\phi})/\underline{\phi}$ is smaller than an acceptable tolerance, which is set to 10^{-6} in the algorithm for example. Alternatively, the termination condition can be set such that the cutoff frequencies $\underline{\omega}_c$ and $\bar{\omega}_c$ are close enough. In each step of the bisection algorithm, a trial cutoff frequency ω'_c is set to the logarithmic center of $\underline{\omega}_c$ and $\bar{\omega}_c$. Then, the minimum phase margin of $L(s; \omega'_c)K(s)P(s)$ is evaluated. If it is lower than the target phase margin ϕ_{target} , then the lower limit is tightened by setting it to $\underline{\omega}_c = \omega'_c$ and the corresponding phase margin $\underline{\phi}$ is re-evaluated as the minimum phase margin of $L(s; \underline{\omega}_c)K(s)P(s)$. Otherwise, if the minimum phase margin of $L(s; \omega'_c)K(s)P(s)$ is higher than the target phase margin ϕ_{target} , then the upper limit $\bar{\omega}_c$ is tightened and $\bar{\phi}$ is re-evaluated. The process repeats until, again, the fractional difference $(\bar{\phi} - \underline{\phi})/\underline{\phi}$ is smaller than an acceptable tolerance.

Algorithm 4 Optimizing the low-pass filter cutoff frequency

$P(s) \leftarrow k \frac{\prod_{i=1}^{n_z} s^2 + \frac{\omega'_{n,i}}{q} s + \omega_{n,i}^2}{\prod_{i=1}^{n_p} s^2 + \frac{\omega_{n,i}}{q} s + \omega_{n,i}^2}$ ▷ The transfer function model.
 $K(s) \leftarrow k_p + \frac{k_i}{s} + k_d s$ ▷ The designed controller (Not necessarily PID).
 $\omega_{\text{ugf},i} \leftarrow \text{get_unity_gain_frequencies}(K(s)P(s))$ ▷ List of unity gain frequencies.
 $\bar{\omega}_c \leftarrow 10 \max(\omega_{\text{ugf},i})$ ▷ The upper limit of the cutoff. E.g. 10 times max UGF.
 $L(s; \omega_c) = \left(\frac{\omega_c}{s + \omega_c} \right)^{n_l}$ ▷ The low-pass filter.
 $\epsilon = 0.1$ ▷ A multiplicative decrement. For example, 0.1.
 $\phi_{\text{target}} \leftarrow 45$ ▷ The target phase margin. For example, 45°.
 $\underline{\omega}_c \leftarrow \bar{\omega}_c$ ▷ Initialize a lower limit of the cutoff frequency.

Decrement the lower limit until the phase margin is lower than the target:

while Min. phase margin of $L(s; \underline{\omega}_c)K(s)P(s) < \phi_{\text{target}}$ **do**

$\omega_c \leftarrow \frac{\omega_c}{1 + \epsilon}$

end while

Bisection algorithm to find the target cutoff frequency:

tolerance $\leftarrow 10^{-6}$ ▷ An acceptable tolerance for termination. For example, 10^{-6} .

$\bar{\phi} \leftarrow$ Min. phase margin of $L(s; \bar{\omega}_c)K(s)P(s)$ ▷ Phase margin with $\omega_c = \bar{\omega}_c$.

$\underline{\phi} \leftarrow$ Min. phase margin of $L(s; \underline{\omega}_c)K(s)P(s)$ ▷ Phase margin with $\omega_c = \underline{\omega}_c$.

▷ The target phase margin is between $(\underline{\phi}, \bar{\phi})$

while $\frac{\bar{\phi} - \underline{\phi}}{\underline{\phi}} > \text{tolerance}$ **do**

$\omega'_c = 10^{\frac{\log \bar{\omega}_c + \log \underline{\omega}_c}{2}}$ ▷ Logarithmic center between the lower and upper limits.

if Min. phase margin of $L(s; \omega'_c)K(s)P(s) < \phi_{\text{target}}$ **then** ▷ Phase margin too low.

$\underline{\omega}_c \leftarrow \omega'_c$ ▷ Tighten the lower limit.

$\underline{\phi} \leftarrow$ Min. phase margin of $L(s; \underline{\omega}_c)K(s)P(s)$ ▷ Re-evaluate $\underline{\phi}$.

else ▷ Phase margin can be lowered.

$\bar{\omega}_c \leftarrow \omega'_c$ ▷ Tighten the upper limit.

$\bar{\phi} \leftarrow$ Min. phase margin of $L(s; \bar{\omega}_c)K(s)P(s)$ ▷ Re-evaluate $\bar{\phi}$.

end if

end while

$\omega_c \leftarrow \omega'_c$ ▷ The target cutoff frequency

return ω_c

And, in the end, the target cut-off frequency ω_c is obtained.

While Algorithm 4 works to provide the low-passed open-loop transfer function with a target phase margin, the controller with the low-pass filter may not be robust. This is because the unity gain frequency may happen to be at frequencies where the phase response has a steep negative slope falling below -180° , which is common for plants with multiple resonances. The SRM F1 GAS filter is an example of this category. As shown in subplot (b) of Fig. 6.28, the phase response swings drastically with steep slopes and it is entirely possible that the unity gain frequencies happen to be on one of the steep slopes. This would mean that if there is a slight deviation in the modeled plant and the actual frequency response, the phase response at those frequencies changes by a large amount

with a possibility to go below -180° . To prevent this from happening, during the evaluation of the phase margins, instead of using the plant $P(s)$ directly, a reduced plant

$$P'(s) = k_{\text{DC}} \frac{\omega_{n,1}^2}{s^2 + \frac{\omega_{n,1}}{q_1}s + \omega_{n,1}^2}, \quad (6.97)$$

can be used to evaluate the phase response, where $\omega_{n,1}$ and q_1 are the resonance frequency and the quality factor of the lowest resonance mode of the original plant $P(s)$ and k_{DC} is the DC gain of the original plant $P(s)$. The phase response of $P'(s)$ is similar to that of $P(s)$ but it does not get bumped up by the additional complex zeros and poles. Therefore, the phase response of $P'(s)$ serves as a pessimistic estimation as it is always lower or equal to that of $P(s)$. With that said, this is not sufficient as the minimum phase margin of $L(s; \omega_c)K(s)P(s)$ might not be the same as that of $L(s; \omega_c)K(s)P'(s)$. Therefore, it is only safe to execute the bisection algorithm in Algorithm 4 twice, one with $P(s)$ and one with $P'(s)$. After that, the higher cutoff frequency is chosen to be the final one. Nevertheless, it is always a good practice to perform a visual inspection of the phase response to ensure the unity gain frequencies do not lie on the falling edges (for those that fall below -180°) of the phase response.

Post-filtering for the SRM F1 GAS filter controller

Using Algorithm 4, and with n_l , the low-pass filter was optimized for the SRM F1 GAS filter degree of freedom. The cutoff frequency was found to be $\omega_c = 71.04 \text{ rad/s}$, i.e. at 11.31 Hz. The final controller is the PID controller obtained in Sec. 6.4.2 with the low-pass filter. The magnitude response of the control filters is shown in Fig. 6.32 as a reference.

The corresponding open-loop transfer functions are plotted in Fig. 6.33. First of all, from subplot (a), it can be seen that the actuation to sensor coupling at high frequencies has been attenuated to below the unity gain. Second of all, from subplot (b), the reduced transfer function model $P'(s)$ (green dash-dot) indeed denotes a worst-case estimation of the phase response of the transfer function model $P(s)$ (orange dashed). The phase response of $P'(s)$ is used in Algorithm 4 to evaluate the phase margins. As can be seen, at the unity gain frequency (black dotted), the phase responses of the measured frequency response, the transfer function model, and the reduced transfer function model cross the target phase margin (black dash-dot). And, all phase margins are above the target phase margin.

At last, the time response of the post-filtered PID controller is compared with that of the PID controller. The SRM F1 GAS filter system is simulated with the two controllers to track a setpoint represented by a unit step function. The time responses are shown in Fig. 6.34. As can be seen, there is practically no difference in time response performance between the two controllers. The only difference is that now the post-filtered PID




Figure 6.32: SRM F1 GAS filter control filters. Blue solid: PID controller. Orange dashed: Low-pass post-filter. Green dash-dot: PID controller with low-pass post-filtering.

controller can be implemented into digital control system of the KAGRA suspension.

Notch filter

A notch filter can be represented by a transfer function with one pair of complex poles and another pair of complex zeros at the same frequency but with a higher Q factor. One with a pair of complex zeros with a lower Q factor is called a peak filter that has a peak rather than a notch filter. There are many ways to express a notch filter. The way the Foton utility (a control utility used in LIGO and KAGRA) defines it is

$$K_{\text{notch}}(s; \omega_{\text{notch}}, q, d) = \frac{s^2 + \frac{\omega_{\text{notch}}}{qd} s + \omega_{\text{notch}}^2}{s^2 + \frac{2\omega_{\text{notch}}}{q} s + \omega_{\text{notch}}^2}, \quad (6.98)$$

where ω_{notch} is the notch angular frequency, q and d are two other parameters. The q parameter behaves like the quality factor but is not quite the quality factor. It controls the sharpness of the notch. The d parameter is attenuation at the notch frequency so it represents the depth of the notch filter. In Fig. 6.35, the magnitude response of example notch filters with different q and d are shown. As can be seen, the attenuation of the notch filter at the notch frequency ω_{notch} is exactly $1/d$, and the sharpness increases with q .

The notch filter is useful in cases where peak-like structures in the signals or transfer functions need to be removed. One typical example would be the vertical degrees of freedom of the Type-B suspension bottom filter (BF) and the intermediate mass (IM).

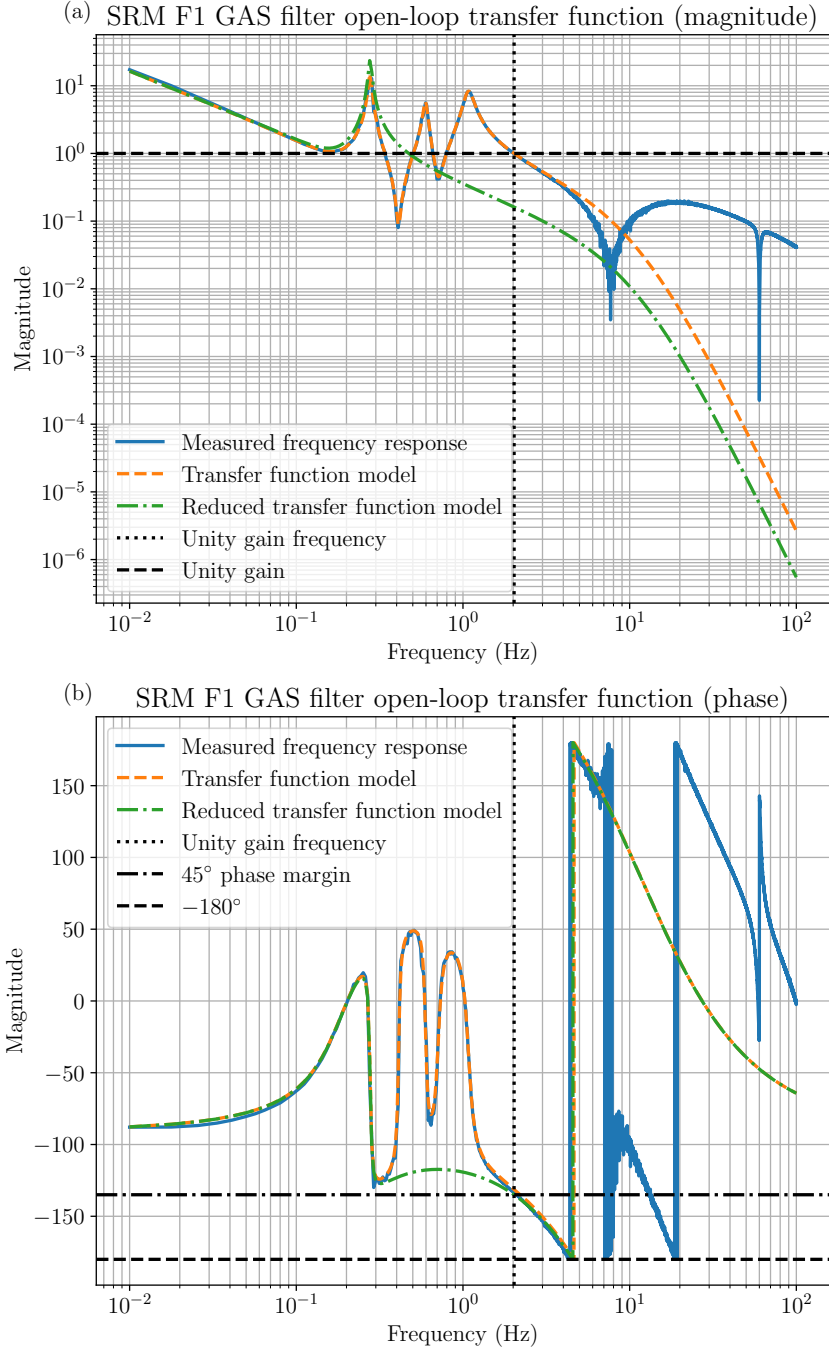


Figure 6.33: SRM F1 GAS filter open-loop transfer function with low-pass post-filtering. (a): Magnitude response. (b): Phase response. Blue solid: Open-loop transfer function estimated with measured frequency response. Orange dashed: Open-loop transfer function estimated with transfer function model. Green dash-dot: Open-loop transfer function estimated with the reduced transfer function model. Black dotted: Unit gain frequency at which the phase margin is lowest. Black dashed in (a): Unity gain. Black dash-dot in (b): Target phase margin. Black dashed in (b): -180° .

The magnitude responses of those degrees of freedom are shown in Fig. 6.36. Again, from the measured frequency response of the BF GAS filter (blue solid), the $\propto f^2$ actuation to sensor coupling due to magnetic coupling is clearly visibly at high frequency. And, this




Figure 6.34: Simulated time response of the SRM F1 GAS filter control. Blue solid: PID control. Orange dashed: PID controller post-filtered with low-pass filter.




Figure 6.35: Magnitude responses of notch filters. Blue solid: Notch filter with $q = 1, d = 10$. Orange dashed: Notch filter with $q = 10, d = 10$. Green dash-dot: Notch filter with $q = 1, d = 100$. Red dotted: Notch filter with $q = 10, d = 100$.

is not modeled as part of transfer function. Together with it, there is a resonance peak at 11.4 Hz that is not modeled as well due to the increased complexity as it falls into the frequency where actuation to sensor coupling dominates.

Assuming that the unmodeled resonance is of no issue, the transfer function model




Figure 6.36: SRM BF GAS filter and IM vertical magnitude response. Blue solid: Measured frequency response of the BF GAS filter. Orange dashed: BF GAS filter transfer function model. Green dash-dot: IM vertical transfer function model.

(orange dashed in Fig. 6.36) is used to design a PID controller and a low-pass post-filter according to Sec. 6.4.2 and low-pass filter subsection in Sec. 6.4.3. The open-loop transfer function is shown in blue solid and orange dashed in Fig. 6.37, with the former estimated with the measured frequency response and the latter estimated with the transfer function model. As shown in Fig. 6.37 the open-loop transfer function estimated with the transfer function model (orange dashed) showed no issues whatsoever, i.e. phase margin target is met and actuation to sensor coupling is well filtered. However, the one estimated with the measured frequency response (blue solid) reveals a serious problem. In subplot (a), at ~ 11.4 Hz, the neglected peak goes above unity gain. And in subplot (b), at that frequency, the phase response is clearly below -180° (starting to fall below -180° at ~ 3 Hz), indicating that the system is unstable. Therefore, the PID controller with low-pass filter is not implementable in reality. To solve this issue, the controller must be post-filtered by a notch filter to keep the system stable.

In principle, the notch filter only needs to attenuate the peak in the open-loop gain such that it goes sufficiently below the unity gain. However, this requirement is loosely defined as there are multiple values of q and d that can achieve this. Moreover, the peak height from the measurement is not always accurate due to measurement error and resolution. Therefore, putting the peak below the unity gain is not necessarily sufficient to ensure stability. Instead, a perfect pole cancellation approach is suggested here as a conservative approach.

Suppose the complex poles in the plant associated with the peak that needs to be




Figure 6.37: SRM BF GAS filter open-loop transfer function. (a): Magnitude response. (b): Phase response. Blue solid: PID with low-pass post-filtering open-loop transfer function estimated with the measured frequency response. Orange dashed: PID with low-pass post-filtering open-loop transfer function estimated with the transfer function model. Green dash-dot: PID with low-pass and notch post-filtering open-loop transfer function estimated with the measured frequency response.

filtered is

$$P_{\text{peak}}(s) = k_{\text{DC}} \frac{\omega_p^2}{s^2 + \frac{\omega_p}{q_p} s + \omega_p^2}, \quad (6.99)$$

where k_{DC} is the DC gain of the transfer function $P_{\text{peak}}(s)$, ω_p is the angular frequency

where the peak is located and q_p is the associated quality factor. This transfer function can be obtained by decomposing the suspension plant into superposition of complex-pole pairs. To exactly cancel this peak, the notch filter needs to have a pair of complex zeros with $\omega_{\text{notch}} = \omega_p$ and $qd = q_p$. The notch filter should not change the frequency-dependency at higher frequency so it must restore the f^{-2} roll-off originally supplied by the complex poles in the plant. This requires the notch filter to contain a double simple pole at $-\omega_p$. This gives the notch filter

$$K_{\text{notch}}(s) = \frac{s^2 + \frac{\omega_p}{q_p}s + \omega_p^2}{(s + \omega_p)^2}. \quad (6.100)$$

This gives the notch filter parameters $\omega_{\text{notch}} = \omega_p$, $q = 1$, and $d = q_p$.

In theory, this above method works well to obtain a notch filter that perfectly cancels the peak. But such low value of q sometimes results in an ill-definition in the Foton utility, which would result in an error. Instead, consider notch filter that has a depth roughly equal to the peak height. The peak of the plant (6.99) locates at $\omega \approx \omega_p$, and the associated height is $\sim k_{\text{DC}}q_p$. Then, the notch filter needs to change that height to k_{DC} , which requires the complex poles of the notch filter to have a quality factor of 1. While the complex zeros of the notch filter needs to cancel the original complex poles of the plant, this gives

$$K_{\text{notch}}(s) = \frac{s^2 + \frac{\omega_p}{q_p}s + \omega_p^2}{s^2 + \omega_p s + \omega_p}. \quad (6.101)$$

Comparing this with Eqn. (6.98), the notch parameters are then $q = 2$, $d = q_p$ and $\omega_{\text{notch}} = \omega_p$ are obtained. It is worth mentioning that the peak height $k_{\text{DC}}q_p$ is only an approximation, so the notch filter does not exactly put the peak at the DC gain. But, this formulation works most of the time in the Foton utility so it is preferred. Otherwise, the former approach with $q = 1$ can always be implemented in Foton by manually entering the complex poles and zeros instead of using the notch filter definition in Foton.

Coming back to the SRM BF GAS filter control, since the peak at 11.4 Hz is not modeled, the transfer function model (orange dashed line in Fig. 6.36) is not useful for obtaining $P_{\text{peak}}(s)$, which is required to define the notch filter. Instead, consider the transfer function of the vertical degree of freedom of the intermediate mass, which is shown as green dash-dotted line in Fig. 6.36. The 11.4 Hz peak is also present. This should be of no surprise since the BF GAS filter degree of freedom is, in principle, measuring the vertical motion of the intermediate mass as it hangs from the BF GAS filter keystone. Therefore, the vertical transfer function model of the intermediate mass can be used to obtain the required notch filter.

After the obtaining the notch filter, the low-pass filter needs to be optimized using Algorithm 4 but $K(s)$ is not the PID controller, but the PID controller with the notch filter applied. This is because the notch filter will also change the phase response open-loop

transfer function so applying the notch filter directly without re-optimizing the low-pass filter would put the phase margin off the target. This means the low-pass filter with the notch filter has a different cutoff frequency compared to that without the notch filter. The open-loop transfer function using the final controller (PID with low-pass and notch filters) is shown as green dash-dot in Fig. 6.37. As can be seen, the peak around 11.4 Hz is no longer observable and this makes a stable closed-loop system. And finally, this concludes the long Chapter 6.

Part III

\mathcal{H}_∞ optimization of control filters in
active vibration isolation systems

Chapter 7

Active Isolation Control Filters

The pre-isolator of the Type-A and Type-B suspensions at KAGRA is the first stage in the suspension chain. It is the interface between the ground and the rest of the suspension that holds the optics. By virtue of the anti-spring effect, the horizontal resonance frequencies of the inverted pendulum suspending the pre-isolator table can be reduced to below 0.1 Hz. This means that the pre-isolator provides passive seismic noise attenuation above 0.1 Hz and the level of attenuation grows as $\sim f^2$ above the resonance frequency. Hence, the pre-isolator not only provide a $\sim 10^4$ times attenuation above the detection band at 10 Hz, but also provide some level of seismic isolation at the frequency band where the secondary microseism dominates, which is around 0.1 Hz – 0.5 Hz. For example, Fig. 7.1 shows the magnitude response (from actuation to displacement) of the SRM pre-isolator in the longitudinal direction. As can be seen, the resonance frequencies of the lowest horizontal modes are well below 0.1 Hz. From the figure, the DC gain is $\sim 10^{-1}$ and at 0.2 Hz, the magnitude is $\sim 10^{-2}$. Assuming that the transfer function from the ground to pre-isolator displacement has roughly the same shape as shown in the figure, this gives roughly 10 times passive attenuation at the secondary microseism (simply referred to microseism hereafter).

The secondary microseism is a background hum in the seismic noise and it appears as a peak in the seismic noise spectrum. It dominates the motion of the ground. Fig. 7.2 shows the amplitude spectral density of the seismic noise measured in the KAGRA mine along the Y-arm direction. Here, the mean, the 10th percentile, the 50th percentile and the 90th percentile spectra are shown. Also, note that at frequencies below 0.1 Hz, the seismometer noise dominates so the spectra do not represent the real seismic noise at those frequencies. As can be seen, at above 0.1 Hz, the seismic noise spectrum is indeed dominated by a peak corresponding to the secondary microseism. The seismic noise level can vary from time to time, depending on the weather. As shown in Fig. 7.2, the microseismic peak can go as low as $\sim 0.1 \mu\text{m}/\sqrt{\text{Hz}}$ at the 10th percentile and it can go as high as $\sim 1 \mu\text{m}/\sqrt{\text{Hz}}$ at the 90th percentile.

While the pre-isolator of the Type-A and Type-B provides some attenuation to the

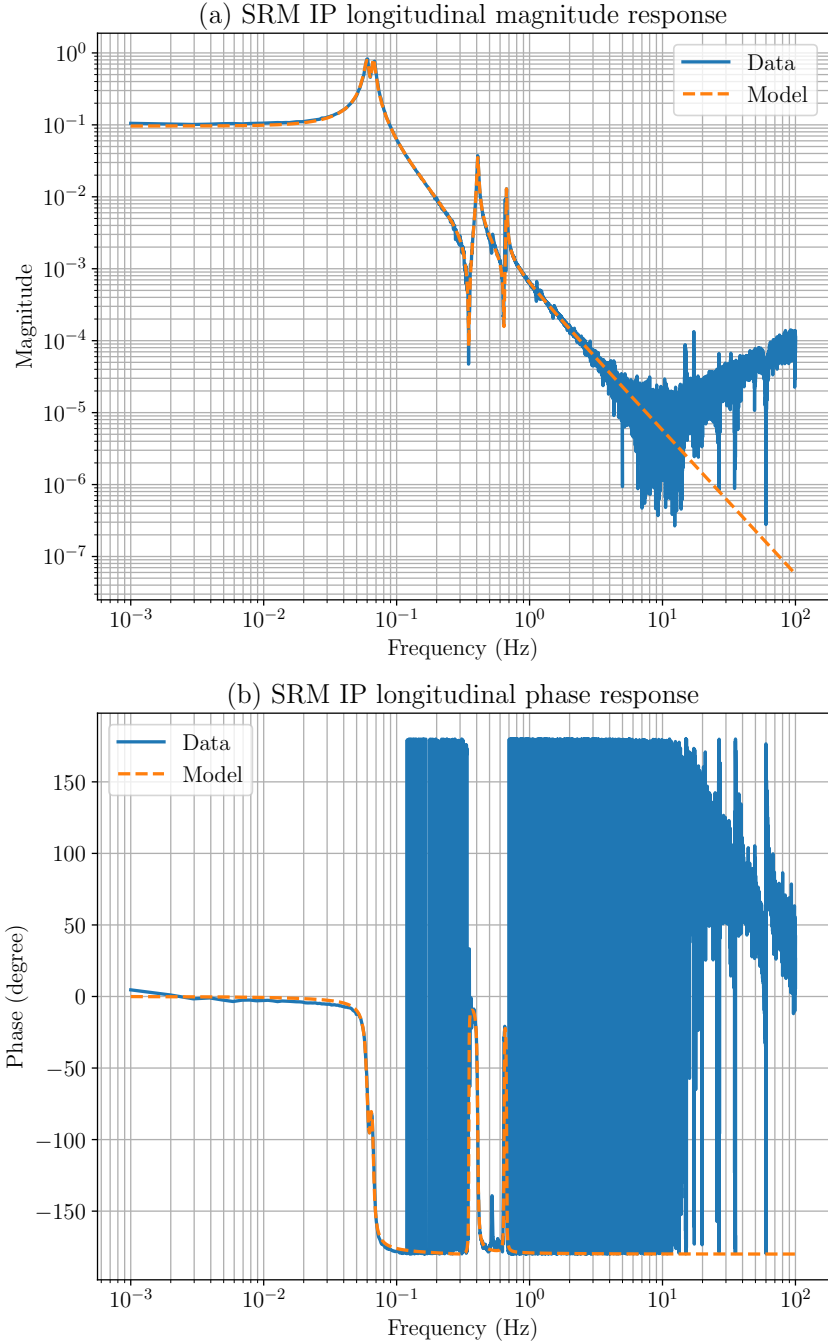


Figure 7.1: SRM IP longitudinal magnitude response.

microseism, the resonances of the suspension are still being excited by the seismic noise at other frequencies. In turn, the resonant motion of the suspension become the dominating factor in the residual motion of the suspension. To suppress the resonances, the use of feedback control, in particular, derivative control, is already discussed in Sec. 6.4. An algorithmic way to obtain the critical derivative gain is also provided. However, the handling of sensing noise has not been discussed. At the pre-isolator stage, the sensors in question are the linear variable differential transformers (LVDTs). While they work well for damping the resonances, there is another issue: they measure relative displacements




Figure 7.2: Seismic noise at KAGRA [53]. Blue solid: Mean spectrum. Orange dashed: 10th percentile. Green dash-dot: 50th percentile. Red dotted: 90th percentile.

rather than inertial motion. This means that the LVDT measurement is coupled to seismic noise. Fig. 7.3 shows the displacement spectrum of the SRM pre-isolator table in the longitudinal direction measured by the LVDT (orange dashed) and the geophone (green dash-dot) under a free swing condition. The seismic noise (blue) is also plotted for a reference. As can be seen, the microseism at around 0.2 Hz is measured almost perfectly by the LVDT. This is not present in the geophone readout, which means the peak correspond to a real seismic coupling but not real motion of the pre-isolator table.

Recall the displacement of a controlled platform $X(s)$, Eqn. (5.4), is a superposition of disturbance $D(s)$ and sensing noise $N(s)$. To restate Eqn. (5.4), the displacement of a feedback controlled object is

$$X(s) = \frac{1}{1 + K(s)P(s)}D(s) - \frac{K(s)P(s)}{1 + K(s)P(s)}N(s), \quad (7.1)$$

where $X(s)$ is the displacement of an object that needs to be controlled, $K(s)$ is the controller, $P(s)$ is the actuation plant, $D(s)$ is the disturbance, and $N(s)$ is the sensing noise. The sensing noise is inevitably injected into to the displacement of the controlled platform where the open-loop gain $|K(j\omega)P(j\omega)|$ is not low enough. With the LVDT coupled to the seismic noise, this means that the secondary microseism is also inevitably re-injected to the motion of the pre-isolator table via the feedback action. And, this puts the situation into a dilemma.

There are three ways to resolve this,

1. Sensor fusion,




Figure 7.3: SRM IP longitudinal displacement spectrum under free swing. Blue solid: Seismic noise measured by a seismometer. Orange dashed: LVDT readout. Green dash-dot: Geophone readout.

2. Sensor correction, and
3. Feedback control optimization.

They are briefly discussed and reviewed in Sec. 7.1, 7.2, and 7.3, respectively. Each of the approach corresponds to an optimization of a control filter(s) in the control system. They are the complementary filters, the sensor correction filter, and the feedback controller, respectively. They can be used independently to improve control performance or they can be used in conjunction. The optimization of these control filters is discussed in later sections after the introduction of the \mathcal{H}_∞ method in Chapter 8, which comes shortly after this chapter.

7.1 Sensor fusion using complementary filters

The idea of sensor fusion is to combine two or more types of sensors with different noise characteristic into one “super sensor” that has an overall better noise performance. There are more than one way achieve sensor fusion. In LIGO, Virgo, and KAGRA, the readouts from different types of sensors are digitally filtered by low-pass filters, band-pass filters, or high-pass filters [26, 27, 29]. The idea of filtering is to pick the sensor readout from frequency region where the sensor has good noise performance and ignore those where the noise performance is poor. The filtered readouts from different types of sensors are then summed into one complete super sensor readout. In order to obtain the original signal

that the sensors all commonly measure, the filters need to be complementary, meaning that their transfer functions are summed to unity. Suppose $H_i(s)$ is the transfer function of the i^{th} filter and there are n sensors, the following constraint

$$\sum_i^n H_i(s) = 1, \quad (7.2)$$

has to be satisfied so the original signal is not distorted in the super sensor readout, assuming that the sensors are all well inter-calibrated. For this reason, these filters are called complementary filters.

In KAGRA, the pre-isolators of the Type-A and Type-B are equipped with two types of sensors, LVDT and geophone¹, as mentioned previously. They both measure the motion of the pre-isolator table. But, in the case of the LVDT, it also measures the motion of the ground since its readout is relative. As a result, the presences of the secondary microseism in the LVDT readout makes it undesirable to be used at frequencies higher than ~ 0.1 Hz. On the other hand, the geophone measures inertial motion of the pre-isolator table rather than relative motion. This means the microseism is not coupled to the geophone readout, making it plausible to be used at frequencies higher than ~ 0.1 Hz. However, inertial measurement units, including geophones, are intrinsically AC-coupled, meaning that they have poor low-frequency performance. In the case of KAGRA geophones, it so happens that the geophones start to become noisier than the LVDTs also at frequencies lower than ~ 0.1 Hz. The LVDT readout requires a low-pass filter while the geophone readout required high-pass filtering. This makes the LVDT and geophone a perfect complement for each other. And, the question is how should these filters be designed?

To state the sensor fusion problem formally, consider a two-sensor configuration as shown in Fig. 7.4. In Fig. 7.4, $X(s)$ is the signal that the two sensors are commonly measuring. The two sensor readouts are $Y_1(s)$ and $Y_2(s)$. They read

$$Y_1(s) = X(s) + N_1(s) \quad (7.3)$$

and

$$Y_2(s) = X(s) + N_2(s), \quad (7.4)$$

where $N_1(s)$ is the sensing noise of the sensor 1 and $N_2(s)$ is the sensing noise in sensor 2. $H_1(s)$ and $H_2(s)$ are the complementary filters that filter the readouts $Y_1(s)$ and $Y_2(s)$, respectively. The filtered readouts are summed to become a super sensor readout $Y_{\text{super}}(s)$,

¹Some Type-A suspensions are equipped with a new type of folded pendulum accelerometer [54]. Development and testing are still on-going.

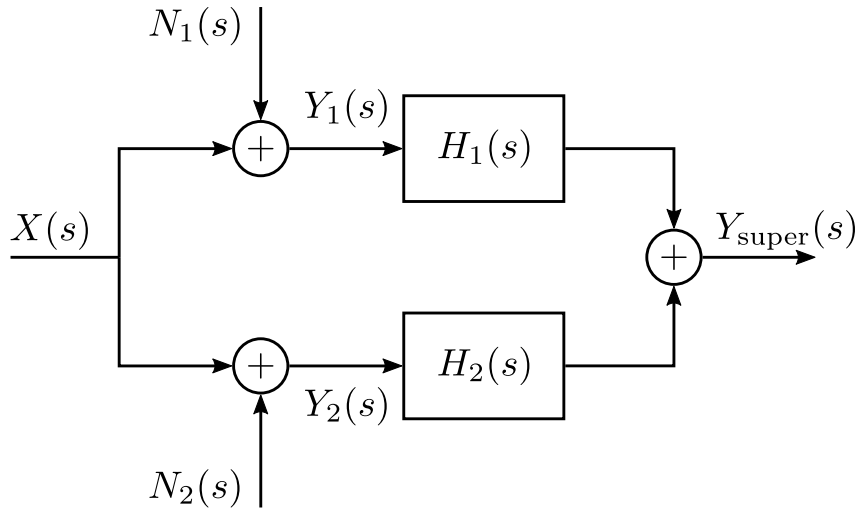


Figure 7.4: Two-sensor sensor fusion configuration.

and it reads

$$\begin{aligned} Y_{\text{super}}(s) &= H_1(s)Y_1(s) + H_2(s)Y_2(s) \\ &= [H_1(s) + H_2(s)] X(s) + [H_1(s)N_1(s) + H_2(s)N_2(s)] . \end{aligned} \quad (7.5)$$

Now, the complementary filters need to satisfy the complementary condition as stated in Eqn. (7.2), so

$$H_1(s) + H_2(s) = 1 . \quad (7.6)$$

This gives

$$Y_{\text{super}}(s) = X(s) + N_{\text{super}}(s) , \quad (7.7)$$

where the super sensor noise $N_{\text{super}}(s)$ is defined as

$$N_{\text{super}}(s) \equiv H_1(s)N_1(s) + H_2(s)N_2(s) . \quad (7.8)$$

And here, the super sensor noise $N_{\text{super}}(s)$ is the quantity that is sought to be minimized (in some way) via an optimization of the complementary filters $H_1(s)$ and $H_2(s)$ under the constraints stated in Eqn. (7.6).

The transfer function of the complementary filters can take any form as long as the complementary condition is met. This makes the design of the complementary filters challenging as the optimization is not a straightforward minimization over certain parameters. In Ref. [29], a pair of complementary filters is proposed for the sensor fusion of the KAGRA LVDTs and geophones. The complementary filters are given as

$$H_1(s; \omega_b) = \frac{35\omega_b^4 s^3 + 21\omega_b^5 s^2 + 7\omega_b^6 s + \omega_b^7}{(s + \omega_b)^7} \quad (7.9)$$

$$H_2(s; \omega_b) = \frac{s^7 + 7\omega_b s^6 + 21\omega_b^2 s^5 + 35\omega_b^3 s^4}{(s + \omega_b)^7}. \quad (7.10)$$

where ω_b is the angular frequency where the magnitude response of the two filters are equal, i.e the blending frequency. The magnitude response of these complementary filters are shown in Fig. 7.5 as an example. As can be seen, the both complementary filters




Figure 7.5: Magnitude response of example complementary filters. Blue solid: Low-pass filter. Orange dashed: High-pass filter.

have a 4th-order roll-off. According to Ref. [29], this filter was chosen according to the frequency-dependency of the geophone noise at low frequencies, which is proportional to $\sim f^{-3.5}$. While it seems reasonable to choose a high-pass filter with a 4th-order for the geophone, a 4th-order low-pass filter seems excessive for filtering the LVDT noise, which is white at high frequency. Following the same argument, a pair of 1st-order low-pass filter and 4th-order high-pass filter should be a more suitable choice of complementary filters for filtering the f^0 noise of the LVDT and $f^{-3.5}$ noise of the geophone.

Ref. [55] generalizes the complementary filters as stated in Eqn. (7.9) and (7.10). It provides a class of complementary filter pairs that is described by

$$H_1(s; a; n_l, n_h) = \frac{\sum_{i=0}^{n_h-1} \binom{n_l + n_h - 1}{i} a^{(n_l+n_h-1-i)} s^i}{(s + a)^{n_l+n_h-1}} \quad (7.11)$$

and

$$H_2(s; a; n_l, n_h) = \frac{\sum_{i=n_h}^{n_l+n_h-1} \binom{n_l + n_h - 1}{i} a^{(n_l+n_h-1-i)} s^i}{(s + a)^{n_l+n_h-1}}, \quad (7.12)$$

where

$$\binom{n_l + n_h - 1}{i} = \frac{(n_l + n_h - 1)!}{i!(n_l + n_h - 1 - i)!}, \quad (7.13)$$

are the binomial coefficients of the polynomial $(s + a)^{n_l + n_h - 1}$, n_l is the order of roll-off of the low-pass filter, n_h is the order of roll-off of the high-pass filter, and a is a real-value variable. It is easy to see that the complementary filters described by Eqn. (7.9) and (7.10) falls into this category, with $n_l = n_h = 4$. Therefore, Eqn. (7.11) and (7.12) can be thought as a generalization of Eqn. (7.9) and (7.10) with any arbitrary order of roll-off for the complementary filters. This makes designing a pair of 1st-order low-pass filter and 4th-order high-pass filter possible for the sensor fusion of LVDT and geophone by selecting $n_l = 1$ and $n_h = 4$.

If the amount of noise of the super sensor can be expressed with a cost function, for example, the root-mean-square value of the noise, then the design of the complementary filters can be straightforwardly obtained by minimizing the cost function over the parameter a . The class of complementary filters presented are simple and easy to design. However, this is not the optimal approach. The class of complementary filters represented in Eqn. (7.11) and (7.12) is not a general representation of all complementary filters. For example, $[H_1(s; a; 4, 4) + H_1(s; a; 1, 4)]/2$ and $[H_2(s; a; 4, 4) + H_2(s; a; 1, 4)]/2$ are also a pair of complementary filters but they cannot be represented by Eqn. (7.11) and (7.12). And, there are infinitely many ways to represent complementary filters. Using Eqn. (7.11) and Eqn. (7.12) are essentially putting additional constraints on the optimization problem.

It is worth mentioning that the complementary filters in Eqn. (7.11) and (7.12) only work well at filtering the asymptotic behavior of the sensor noise at the ends of the spectrum. As shown in Fig. 7.3, in the middle of the spectrum, the LVDT readout is coupled to the secondary microseism, which is clearly visible at around 0.2 Hz. The spectral shape a sensor noise does not follow an asymptote in general and contain different features at different frequencies. This necessarily means that the complementary filters in Eqn. (7.11) and (7.12) will not handle those features well, making them sub-optimal. Refs. [26, 27] show examples of complementary filters that contain notch features that are designed to especially attenuate the seismic noise in relative sensors. However, the transfer functions of these complementary filters are manually designed according to practical experience and are very hard to reproduce for other systems. They are not optimized either since they are manually tuned. As will be seen in later sections, these complementary filters can be optimized with no information other than the sensor noises themselves and these special filtering features will naturally appear as a result of optimization.

7.2 Sensor correction for relative displacement sensors

Sensor correction is a control strategy used in LIGO to convert a measured relative displacement between a sensed platform and a reference platform to an inertial readout measuring majorly the motion of the sensed platform [26]. This is achieved by subtracting the relative readout by an inertial readout of the reference platform, thereby “correcting” the relative sensor readout. The sensor correction scheme is distinguished from the sensor fusion scheme, which uses an inertial readout of the sensed platform instead. In KAGRA, there are several seismometers measuring the ground motion at proximities to the main optics. They can be used to remove the seismic noise component in the relative sensors, i.e. the LVDTs, at the pre-isolator stage. Again, here it is assumed that the relative sensors and seismometers are well inter-calibrated.

The block diagram of the sensor correction scheme at the pre-isolator stage is shown in Fig. 7.6. In the figure, $X(s)$ denotes the displacement of the pre-isolator table and $X_g(s)$ is




Figure 7.6: Sensor correction scheme at the pre-isolator stage.

the ground displacement. The relative sensor reads a relative displacement $X(s) - X_g(s)$. It contains a self-noise $N_{\text{rel}}(s)$ so the relative readout $Y_{\text{rel}}(s)$ is

$$Y_{\text{rel}}(s) = X(s) - X_g(s) + N_{\text{rel}}(s). \quad (7.14)$$

The ground motion $X_g(s)$ is measured by a seismometer that has a self-noise of $N_{\text{seis}}(s)$. The seismometer readout $Y_{\text{seis}}(s)$ then simply reads

$$Y_{\text{seis}}(s) = X_g(s) + N_{\text{seis}}(s). \quad (7.15)$$

Now, the ground component $X_g(s)$ in Eqn. (7.14) is an unwanted signal. It can be removed by adding the seismometer readout (7.15) to the the relative readout (7.14) via a sensor

correction filter $H_{sc}(s)$. This gives a sensor corrected readout $Y_{sc}(s)$ and it reads

$$\begin{aligned} Y_{sc}(s) &= Y_{rel}(s) + H_{sc}(s)Y_{seis}(s) \\ &= X(s) - X_g(s) + N_{rel}(s) + H_{sc}(s)[X_g(s) + N_{seis}(s)] \\ &= X(s) + N_{rel}(s) + H_{sc}(s)N_{seis}(s) - [1 - H_{sc}(s)]X_g(s). \end{aligned} \quad (7.16)$$

The sensor correction readout can be rewritten as

$$Y_{sc}(s) = X(s) + N_{rel}(s) + N_{sc}(s), \quad (7.17)$$

where $N_{sc}(s)$ is the sensor correction noise defined as

$$N_{sc}(s) \equiv H_{sc}N_{seis}(s) - [1 - H_{sc}(s)]X_g(s). \quad (7.18)$$

Here, the sensor correction readout in Eqn. (7.17) has one signal term $X(s)$ and two noise terms $N_{rel}(s)$ and $N_{sc}(s)$. The relative displacement sensor noise $N_{rel}(s)$ is not minimizable by means of sensor correction filtering. Therefore, it remains as an ambient noise in the sensor correction readout. This is an important note for sensor correction filter optimization, as is indicated in later sections. On the other hand, the sensor correction noise $N_{sc}(s)$ is minimizable by optimizing the sensor correction filter $H_{sc}(s)$. In fact, the sensor correction noise in the form of Eqn. (7.18) is very similar to super sensor noise in the form of Eqn. (7.8), assuming that one of the complementary filters can be represented as $H_2(s) = 1 - H_1(s)$ due to the complementary condition. For this reason, the sensor correction problem is very similar to that of the sensor fusion and can be solved in a similar way. However, even if this is acknowledged in Ref. [56], the sensor correction problem and sensor fusion problem have seemingly be treated as separate problems.

The relative readout cannot be subtracted by the seismometer readout directly without a sensor correction filter. This is because the seismometer is an inertial sensor like the geophone. The seismometers at KAGRA have a slightly different noise profile than the geophones but they are still intrinsically AC-coupled. Therefore, they have poor noise performance at low frequencies. For this reason, the sensor correction filter needs to be a high-pass filter that attenuates the low-frequency noise of the seismometer. Besides, the sensor correction filter also need to have a pass-band at above ~ 0.1 Hz such that the secondary microseism gets removed from the relative readout. At the pass-band, the magnitude response of the sensor correction filter needs to be as closed to unity as possible and the phase needs to be as close to zero as possible. While this might seem reasonable, this is not the case when the ambient noise is put into consideration, i.e. the seismic noise attenuation needs not to be infinity. The seismic noise attenuation only needs to be just enough such that none of the seismic noise, relative sensor noise, and seismometer noise dominates in the sensor correction readout.

The first design of a sensor correction filter for LIGO is given in Ref. [56]. In Ref. [56], the sensor correction filter is designed according to the the following criteria:

1. Below 0.008 Hz, the magnitude response of the sensor correction filter needs to be lower or equal to 8×10^{-4} .
2. Between 0.008 Hz and 0.04 Hz, the roll-off of the high-pass filter should be f^3 .
3. Between 0.04 Hz and 0.01 Hz, the magnitude response of the sensor correction filter is allowed to have an amplification of at most 3.
4. Above 0.01 Hz, the transmissivity of seismic noise should be as close to zero as possible. Equivalently, this requires the magnitude response of the sensor correction filter to be as close to one as possible with phase as close to zero as possible. If the optimized filter gives seismic transmissivity greater than 0.1, other requirement need to be loosened.

The sensor correction filter satisfying these requirements was replicated in Ref. [57] and is shown as blue solid in Fig. 7.7 as an example. While it is possible to to satisfy all the requirements stated, the sensor correction filter satisfying these requirement is not necessarily the optimal sensor correction filter. This is due to the fact that the requirements were set based on heuristics, which maybe arbitrary and sub-optimal. For example, in the forth criterion, there is no reason why the seismic transmissivity should be less than one and it certainly shouldn't be set to as low as possible. This is because the sensor correction readout is limited by the relative sensor noise, which is ambient anyway.

The default sensor correction filter in KAGRA is simply a 2nd-order elliptic high-pass filter as shown as orange dashed line in Fig. 7.7 [31]. The order of roll-off was set to cancel the f^{-2} noise of the seismometer noise at lower frequency. And, the pass-band was set to be above 0.06 Hz, which is based on the fact that the primary microseism at around 0.06 Hz is observable in the seismometer readout as shown in Fig. 7.2. This forces the noise amplification band to be lower than that frequency where the seismometer noise is high. As a consequence, this sensor correction scheme with this filter injects more noise to the sensor correction readout than it suppresses. With position control system engaged, the pre-isolator table, along with the whole suspension chain, would drift at ~ 0.013 Hz with an excess amount that is non-negligible, causing difficulties to lock-acquisition of the interferometer [58]. Therefore, this sensor correction filter is not used currently.

Ref. [59] gives the first attempt at optimizing the sensor correction filter to minimize the sensor correction noise using a particle swarm optimization algorithm. While it was shown that the method can reduce the root-mean-square (RMS) value of the sensor correction noise in velocity unit, there are some caveats. First of all, the RMS value in velocity is only one of many performance indices for seismic isolation. Therefore, minimizing the RMS value in velocity does not necessarily lead to a good seismic isolation filter. Solely




Figure 7.7: Replicated LIGO sensor correction filter and KAGRA sensor correction filter. Blue solid: First LIGO sensor correction filter (replicated). Orange dashed: KAGRA sensor correction filter.

using an RMS value as a cost function is not optimal either. This is because the RMS value is a value integrated over the whole full frequency spectrum, which makes broadband features more dominant than peak features in the cost function. The RMS value is also usually dominated by one feature that is the highest in the spectrum, such as the micro-seism. Features with smaller values have very low weighting in the cost function and are usually not treated. Second of all, the implementation of an artificial frequency-dependent weighting function in the cost function does not help either as it distorted the original purpose of the optimization, which is a reduction in RMS velocity. Most importantly, the method optimizes a transfer function with a pre-defined structure. As mentioned in Sec. 7.1, this puts unnecessary constraints on the optimization problem, which makes the final product sub-optimal. Nevertheless, the work claimed that the improvement in velocity noise performance is most significant between 0.4 Hz and 0.9 Hz where the amplitude spectral densities of the noise is clearly below 10^{-9} m/sec/ $\sqrt{\text{Hz}}$. The ambient noise from the capacitive sensors, the relative sensors used at LIGO, have a noise of slightly below 10^{-9} m/ $\sqrt{\text{Hz}}$ [60], which translates to slightly higher than 10^{-9} m/sec/ $\sqrt{\text{Hz}}$ in velocity. Therefore, the improvement may not be observable in reality.

7.3 Feedback control filters

The concept of feedback control for coarse alignment and damping of the suspensions has been discussed previously in Sec. 6.4. The presented method allows one to design a

feedback controller composes of a PID controller and optionally with low-pass and notch post-filtering. The feedback control filter is simple, easy-to-use, and is readily available for implementation to the actual active vibration isolation systems in gravitational-wave detectors. However, the presented method is based on a critical damping criterion, which may a conservative criterion. The controller shaped this way is not optimal in any sense. The method is also based on a pre-defined controller with a PID structure, which is not the only way a controller can be shaped. For example, the controller can contain a component representing the inverse of the plant that completely cancels the dynamics of the plant. In this case, other multiplicative components in the controller represents the required control bandwidth, which can be designed as if the plant is unity. Therefore, as can be seen, the PID structure is not the most general. It would introduce unnecessary constraints to the underlying optimization problem and alternative must be seek for a truly optimized system.

Still, in current literature, documented feedback controllers for active isolation systems in current gravitational-wave detectors are largely based on the PID structure as described in Sec. 6.4, even though the design method may be different [30, 29, 31, 27, 33, 28, 26, 61]. While performances of the controllers are claimed, most of the results are not reproducible since the actual analytic forms of the controllers are usually not documented and the design methodologies are not well-defined. Also, in many KAGRA literature [30, 29, 31, 33], the decay time constants of the resonances are used as a performance index for controller tuning. While the decay time indicates how the control system reacts to a step disturbance, it is not sufficient to prove the efficacy of an active isolation controller as the decay time only indicates the transient response, not the steady state. The RMS values of the optics displacement are sometimes measured also. This maybe a good indicator as the lock stability of the interferometer has a strong correlation to low RMS values. However, the results are not obtained from “controlled experiments” since the external disturbances are neither measured nor injected. This means that the results may differ as the environment changes.

The optimization of a feedback control is to make the optimal trade-off between the disturbance rejection and noise attenuation. As indicated in Eqn. (7.1), the displacement of a controlled platform $X(s)$ is a superposition of disturbance $D(s)$ and noise $N(s)$. The corresponding coupling terms for the disturbance $D(s)$ and the noise $N(s)$ are $1/[1 + K(s)P(s)]$ and $K(s)P(s)/[1 + K(s)P(s)]$, respectively. And, the goal of the feedback control problem is to design a controller $K(s)$ such that the controlled displacement $X(s)$ is minimized in some sense. Roughly speaking, the open-loop gain $K(s)P(s)$ needs to be high when $D(s)$ is greater than $N(s)$ and needs to be low (lower than 1) when $D(s)$ is smaller than $N(s)$. This makes $X(s) \approx -N(s)$ when $N(s)$ is the smaller term and $X(s) \approx D(s)$ when $D(s)$ is the smaller term, i.e. $X(s)$ becomes close to the lower limit.

In control theory, the coupling terms $1/[1+K(s)P(s)]$ and $K(s)P(s)/[1+K(s)P(s)]$ are

termed sensitivity and complementary sensitivity, respectively. In the two terms, the only design freedom is the controller $K(s)$. And, it should be no surprise that the two terms are complementary, i.e. the transfer functions of the two terms are summed to unity. This makes the feedback control problem effectively a sensor fusion problem. Hence, the underlying problem is identical to that of Eqn. (7.8) and (7.18). The sensitivities $1/[1 + K(s)P(s)]$ and $K(s)P(s)/[1 + K(s)P(s)]$ are analogues to the complementary filters $H_1(s)$ and $H_2(s)$. And, the disturbance $D(s)$ and noise $N(s)$ are analogues to the two sensing noises $N_1(s)$ and $N_2(s)$. Therefore, similar to sensor correction, the feedback control problem can also be solved like a complementary filter problem.

Chapter 8

\mathcal{H}_∞ Optimization of Control Filters

H-infinity (\mathcal{H}_∞) method [62] is a relative new and advanced method of optimization controllers. In this chapter, the \mathcal{H}_∞ optimization of control filters, including complementary filters, sensor correction filters, and feedback controllers, of the active isolation systems is discussed. This chapter is divided into four sections. Sec. 8.1 introduces the idea of \mathcal{H}_∞ control and its formalism. Sec. 8.2 explains how the filter shaping problems can be formulated as an \mathcal{H}_∞ problem. Examples will also be given whenever appropriate. Sec. 8.3 shows experimental results using the KAGRA SRM suspension. And, the \mathcal{H}_∞ filters are compared to the current control filters used in KAGRA.

8.1 \mathcal{H}_∞ method

8.1.1 Formalism

\mathcal{H}_∞ method refers to a controller synthesis process where the \mathcal{H}_∞ norm of a generalized plant is minimized. To specify the \mathcal{H}_∞ problem, a few things need to be defined. To begin with, consider a generalized plant representation shown in Fig. 8.1 In the figure,




Figure 8.1: Generalized plant representation.

the generalized plant $\mathbf{P}(s)$ has two inputs w and u , and two outputs z and y . Note that the generalized plant $\mathbf{P}(s)$ under the context of \mathcal{H}_∞ method is, in general, not the same as the plant $P(s)$ discussed in previous sections, which represents a physical dynamical system. Here, w represents the exogenous inputs, u represents the manipulated inputs, z represents the error signals, and y represents the measured outputs. The inputs and outputs are vector quantities so the \mathcal{H}_∞ method is generally applicable to systems with multiple inputs and multiple outputs. The measured outputs y are the inputs of a feedback regulator $\mathbf{K}(s)$, which outputs the manipulated signals u . The feedback regulator $\mathbf{K}(s)$ is the quantity that is sought to be optimized. Again, $\mathbf{K}(s)$ is generally not the same as the feedback controller $K(s)$ mentioned in previous sections.

The generalized plant $\mathbf{P}(s)$ specifies the interconnection between the inputs $[w, u]^\top$ and the outputs $[z, y]^\top$. Under the open-loop condition, the input-output relation reads

$$\begin{aligned} \begin{bmatrix} z \\ y \end{bmatrix} &= \mathbf{P}(s) \begin{bmatrix} w \\ u \end{bmatrix} \\ &= \begin{bmatrix} P_{11}(s) & P_{12}(s) \\ P_{21}(s) & P_{22}(s) \end{bmatrix} \begin{bmatrix} w \\ u \end{bmatrix}, \end{aligned} \quad (8.1)$$

where $P_{11}(s)$, $P_{12}(s)$, $P_{21}(s)$, and $P_{22}(s)$ are the partitioned elements of $\mathbf{P}(s)$ and are also transfer function matrices in general. In the closed-loop condition, the manipulated inputs read

$$u = \mathbf{K}(s)y. \quad (8.2)$$

Rewriting y gives

$$\begin{aligned} y &= P_{21}(s)w + P_{22}(s)u \\ &= P_{21}(s)w + P_{22}(s)\mathbf{K}(s)y \\ &= [1 - P_{22}(s)\mathbf{K}(s)]^{-1} P_{21}(s)w \end{aligned} \quad (8.3)$$

Substituting Eqn. (8.2) and (8.3) into Eqn. (8.1) yields an closed-loop error signal

$$\begin{aligned} z &= P_{11}(s)w + P_{12}(s)u = P_{11}(s)w + P_{12}(s)\mathbf{K}(s)y \\ &= P_{11}(s)w + P_{12}(s)\mathbf{K}(s) [1 - P_{22}(s)\mathbf{K}(s)]^{-1} P_{21}(s)w \\ &= \{P_{11}(s) + P_{12}(s)\mathbf{K}(s) [1 - P_{22}(s)\mathbf{K}(s)]^{-1} P_{21}(s)\} w, \end{aligned} \quad (8.4)$$

where

$$F_l(\mathbf{P}, \mathbf{K}) = P_{11}(s) + P_{12}(s)\mathbf{K}(s) [1 - P_{22}(s)\mathbf{K}(s)]^{-1} P_{21}(s) \quad (8.5)$$

is called the lower linear fractional transformation and it represents the closed-loop transfer function matrix from w to z .

To avoid the tediousness of writing the linear fractional transformation, let us define

$$\mathbf{G}(s) \equiv F_l(\mathbf{P}, \mathbf{K}) \quad (8.6)$$

as the closed-loop transfer matrix such that $z = \mathbf{G}(s)w$. Now, the \mathcal{H}_∞ norm is defined as

$$\|\mathbf{G}(s)\|_\infty \equiv \max_\omega \bar{\sigma}(\mathbf{G}(j\omega)), \quad (8.7)$$

where $\bar{\sigma}(\mathbf{G}(j\omega))$ is the maximum singular value of $\mathbf{G}(j\omega)$ ¹. The maximum singular value is a function of frequency and it can be expressed as

$$\bar{\sigma}(\mathbf{G}(j\omega)) = \max_i \sqrt{\lambda_i(\mathbf{G}(j\omega)^H \mathbf{G}(j\omega))}, \quad (8.8)$$

where $\lambda_i(\mathbf{G}(j\omega)^H \mathbf{G}(j\omega))$ denotes the i^{th} eigenvalue of the matrix $\mathbf{G}(j\omega)^H \mathbf{G}(j\omega)$ and $\mathbf{G}(j\omega)^H$ is the conjugate transpose of $\mathbf{G}(j\omega)$. The \mathcal{H}_∞ problem is then to find a stabilizing controller $\mathbf{K}_\infty(s)$ that minimizes Eqn. (8.7), i.e. the \mathcal{H}_∞ norm.

Now, the \mathcal{H}_∞ norm can be quite abstract as its interpretation depends on the definition of the plant in question. For a system with single input and single output, such as the plant $P(s)$ representing the dynamics of a single degree of freedom of a suspension, the \mathcal{H}_∞ norm is simply the maximum gain of the magnitude response. In other words, the \mathcal{H}_∞ norm is the maximum magnification of a transfer function. Now, for a general multiple-input-multiple-output system, there is no direct way of plotting the magnitude response like the Bode plot since the plant would be a transfer function matrix. Instead, at each frequency, the singular values of the transfer function matrix can be evaluated. The maximum singular value is also the induced 2-norm of the transfer function matrix and the induced 2-norm can be written as

$$\bar{\sigma}(\mathbf{G}(j\omega)) = \max_{\|w\|_2 \neq 0} \frac{\|\mathbf{G}(j\omega)w\|_2}{\|w\|_2} = \max_{\|w\|_2=1} \|\mathbf{G}(j\omega)w\|_2, \quad (8.9)$$

where $\|w\|_2$ denotes the vector 2-norm and is equal to $\sqrt{\sum_i^n w_i^2}$ and w_i is the i^{th} element of the input vector w and n is the number of inputs. Note that $\|\mathbf{G}(j\omega)w\|_2$ is also a vector 2-norm and it is equal to $\sqrt{\sum_i^m z_i^2}$, where z_i is the element of the output vector z and m is the number of outputs. While the quantity $\sqrt{\sum_i^m z_i^2}$ is a measure of the size of the output, it can be seen that the maximum singular value is a measure of the maximum ‘‘amplification’’ of the transfer function matrix subjected to an input w in the worst direction. Consequently, the \mathcal{H}_∞ norm is then a measure of the maximum ‘‘amplification’’ of the transfer function matrix when subjected to an input in the worst direction

¹Strictly speaking, the ‘‘max’’ (the maximum value) should be replaced by ‘‘sup’’ (the supremum, the least upper bound). This is because there exists cases where the maximum can only be approached as $\omega \rightarrow \infty$, which cannot be realistically achieved [63]. For practical purposes, there is no difference between ‘‘max’’ and ‘‘sup’’ so ‘‘max’’ is kept here for the purpose of simplicity.

at the worst frequency. The fact that the \mathcal{H}_∞ is a measure of the peak magnitude of a transfer function allows one to interpret it as the magnitude of the transfer function relative to some specified bounds by introducing weighting functions in the generalized plant $\mathbf{P}(s)$ [63]. For this reason, the generalized plant $\mathbf{P}(s)$ is not necessarily an interpretation of a real system. Rather, the generalized plant implies cost function with specified performance. This shall become important later when considering optimal filter design in Sec. 8.2.

8.1.2 Solving the \mathcal{H}_∞ problem

The \mathcal{H}_∞ problem is to find all stabilizing controllers $\mathbf{K}(s)$ that minimizes the \mathcal{H}_∞ norm $\|F_l(\mathbf{P}, \mathbf{K})\|_\infty$. There are many ways to approach \mathcal{H}_∞ optimal controllers, including Youla–Kucera parametrization approaches [64, 65], Riccati-based approaches [66, 67], and linear matrix inequality (LMI) approaches [68]. These methods are available in commercial packages such as the Robust Toolbox in MATLAB and Control in Python [69]. The problems presented in later sections including Sec. 8.2 and 8.3 are solved using the `control.hinfscn()` function provided by the Python Control package [69]. The \mathcal{H}_∞ synthesis function is a wrapper for the SLICOT library [70] (which stands for Subroutine Library In Control Theory), which provides Fortran 77 implementations of numerical algorithms for computations in control theory [71]. Of all the routines in SLICOT, the synthesis subroutine SB10AD is used to find the \mathcal{H}_∞ optimal controller according to a specified generalized plant. The subroutine is based on a Riccati approach as described in Ref. [67]. A summary of the Riccati approach is given in Ref. [63] and it is re-elaborated in this section.

The idea of the Riccati approach is to first find all stabilizing controllers $\mathbf{K}(s)$ such that

$$\|F_l(\mathbf{P}, \mathbf{K})\|_\infty < \gamma, \quad (8.10)$$

where γ is a specified value such that

$$\gamma > \gamma_{\min} \quad (8.11)$$

and γ_{\min} is the minimum possible value of $\|F_l(\mathbf{P}, \mathbf{K})\|_\infty$. This corresponds to finding a sub-optimal solution to the \mathcal{H}_∞ problem, which is simpler than finding the optimal solution directly [63]. Then, this process is iterated with decreasing value of γ until no solution can be found. This indicates that the optimal solution is approached as $\gamma \rightarrow \gamma_{\min}$.

To use the Riccati approach, it is more convenient to specify the generalized plant in state-space representation rather than transfer matrix representation. Consider the generalized plant in Fig. 8.1, the system is described by

$$\dot{x} = Ax + B_1w + B_2u, \quad (8.12)$$

$$z = C_1x + D_{11}w + D_{12}u, \quad (8.13)$$

and

$$y = C_2x + D_{21}w + D_{22}u, \quad (8.14)$$

where $A, B_1, B_2, C_1, C_2, D_{11}, D_{12}, D_{21}, D_{22}$ are matrices of the state-space model, x is the internal states such as the displacement and velocity (not the same as the displacement in previous sections), y is the measurement, z is the error signals, w is the exogenous inputs, and u is the control variables. The state-space representation of the generalized plant is then given by

$$\mathbf{P} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{array} \right]. \quad (8.15)$$

The following assumptions about the plant are typically made

1. (A, B_2) is stabilizable and (A, C_2) is detectable. This is required for the existence of a stabilizing controller.
2. D_{12} and D_{21} have full rank. This ensures the controller is proper and hence realizable.
3. $\begin{bmatrix} A - j\omega I & B_2 \\ C_1 & D_{12} \end{bmatrix}$ has full column rank for all ω , and
4. $\begin{bmatrix} A - j\omega I & B_1 \\ C_2 & D_{21} \end{bmatrix}$ has full row rank for all ω . Combined with Assumption 3, this ensures the optimal controller doesn't cancel poles or zeros on the imaginary axis which would result in instability.
5. $D_{11} = D_{22} = 0$. This simplifies the formulas of the algorithm.
6. $D_{12} = \begin{bmatrix} 0 \\ I \end{bmatrix}$ and $D_{21} = \begin{bmatrix} 0 & I \end{bmatrix}$. For simplicity. This is achievable by a scaling of u and y and a unitary transformation of z and w .
7. $D_{12}^T C_1 = 0$ and $B_1 D_{21}^T = 0$. A common assumption in optimal control such as Linear-Quadratic-Gaussian (LQG). The former means there is no cross term between the states x and the manipulated inputs u in the cost function. And, the latter assumption means that the plant disturbance and measurement noise are uncorrelated.
8. (A, B_1) is stabilizable (A, C_1) is detectable. If Assumption 7 holds, Assumptions 3 and 4 are replaced by this.

Assumptions 1 and 8 are related to concepts of controllability and observability, which is not discussed here. For testing of stabilizability and detectability, refer to tests like the Popov-Belevitch-Hautus test [72]. The major results are shown here. The Hautus lemma for stabilizability states that a pair (A, B) is stabilizable if

$$\text{rank}\left(\begin{bmatrix} \lambda_i(A)I - A & B \end{bmatrix}\right) = n \quad (8.16)$$

for all eigenvalues $\lambda_i(A)$ with $\Re(\lambda_i(A)) \geq 0$, where n is the number of states such that $x \in \mathbb{R}^n$. The Hautus lemma for detectability states that a pair (A, C) is detectable if

$$\text{rank}\left(\begin{bmatrix} \lambda_i(A)I - A \\ C \end{bmatrix}\right) = n \quad (8.17)$$

for all eigenvalues $\lambda_i(A)$ with $\Re(\lambda_i(A)) \geq 0$.

Now, under these assumptions, there exists a stabilizing controller such that $\|F_l(\mathbf{P}, \mathbf{K})\|_\infty < \gamma$ under three conditions:

- (i) $X_\infty \geq 0$ is a solution to the algebraic Riccati equation

$$A^\top X_\infty + X_\infty A + C_1^\top C_1 + X_\infty (\gamma^{-2} B_1 B_1^\top - B_2 B_2^\top) X_\infty = 0 \quad (8.18)$$

such that the eigenvalues of $[A + (\gamma^{-2} B_1 B_1^\top - B_2 B_2^\top) X_\infty]$ have real parts lower than 0.

- (ii) $Y_\infty \geq 0$ is a solution to the algebraic Riccati equation

$$A Y_\infty + Y_\infty A^\top + B_1 B_1^\top + Y_\infty (\gamma^{-2} C_1^\top C_1 - C_2^\top C_2) Y_\infty = 0 \quad (8.19)$$

such that the eigenvalues of $[A + Y_\infty (\gamma^{-2} C_1^\top C_1 - C_2^\top C_2)]$ have real parts lower than 0.

- (iii) The spectral radius of $X_\infty Y_\infty$ is less than γ^2 , i.e.

$$\rho(X_\infty Y_\infty) = \max(\lambda_i(X_\infty Y_\infty)) < \gamma^2, \quad (8.20)$$

where $\lambda_i(X_\infty Y_\infty)$ denotes the eigenvalues of $X_\infty Y_\infty$.

All such stabilizing controllers can be written as a lower linear fractional transformation $\mathbf{K}(s) = F_l(K_c, Q)$, where

$$K_c = \left[\begin{array}{c|cc} A_\infty & -Z_\infty L_\infty & Z_\infty B_2 \\ \hline F_\infty & 0 & I \\ -C_2 & I & 0 \end{array} \right], \quad (8.21)$$

$F_\infty = -B_2^T X_\infty$, $L_\infty = -Y_\infty C_2^T$, $Z_\infty = (I - \gamma^{-2} Y_\infty X_\infty)^{-1}$, $A_\infty = A + \gamma^{-2} B_1 B_1^T X_\infty + B_2 F_\infty + Z_\infty L_\infty C_2$, and $Q(s)$ is any stable proper transfer function such that $\|Q(s)\|_\infty < \gamma$. For $Q(s) = 0$, we get

$$\mathbf{K}(s) = K_{c11}(s) = -F_\infty (sI - A_\infty)^{-1} Z_\infty L_\infty, \quad (8.22)$$

which is called the central controller and it has the same states as the generalized plant $\mathbf{P}(s)$. To approach the optimal controller, the value of γ is needed to be updated in each iteration. The value is updated based on a modified bisection algorithm called γ -iteration, where the procedure continues until the smallest γ value that passed and the largest γ value that failed is close within an acceptable tolerance [73].

8.2 Filter design problem as an \mathcal{H}_∞ problem

To use \mathcal{H}_∞ optimization, the control filter optimization problems must be expressed as a generalized plant $\mathbf{P}(s)$ as shown in Fig. 8.1. After that, the optimal control filters can be obtained numerically using \mathcal{H}_∞ synthesis using available software such as MATLAB (which is proprietary) or Python Control [69] (which is open source and is used in this work). In this section, the sensor fusion problem is first expressed as an \mathcal{H}_∞ problem in Sec. 8.2.1. The sensor correction and feedback control problems are similar to the sensor fusion problem, i.e. they are optimal trade-off problems with two conflicting objectives as indicated by Eqn. (7.1), (7.8), and (7.18). They can be solved like a sensor fusion problem. A few remarks regarding sensor correction and feedback control problems are given in Sec. 8.2.2 and 8.2.3, respectively.

8.2.1 Sensor fusion problem

The method discussed here is largely inspired by the work in Ref. [57], which provides the \mathcal{H}_∞ framework for shaping complementary filters according to frequency-dependent constraints. The method is also documented in Ref. [74]. And, this subsection provides a re-elaboration of the work.

To formulate the sensor fusion as an \mathcal{H}_∞ problem, the block diagram as shown in Fig. 7.4 must be modified to impose the complementary condition $H_1(s) + H_2(s) = 1$. With this constraint, the problem is really about optimizing one filter, $H_1(s)$ or $H_2(s)$, and the other is automatically constrained as $1 - H_2(s)$ or $1 - H_1(s)$, respectively. Now, it was shown that the sensor correction problem can be seen as a sensor fusion problem. While the sensor correction configuration has only one filter to be optimized, this suggests a modification the sensor correction block diagram in Fig. 7.6 for sensor fusion. By comparing Eqn. (7.8) and (7.17), $X(s)$ and $N_{\text{rel}}(s)$ are the extra terms so they can be removed from the block diagram. Matching the remaining terms suggests replacement

of terms from $N_{\text{seis}}(s)$ to $N_1(s)$, $X_g(s)$ to $-N_2(s)$, and $H_{\text{sc}}(s)$ to $H_1(s)$. Here, $H_1(s)$ is selected for optimization and $H_2(s)$ is obtained via $1 - H_1(s)$ later. Applying these changes gives a modified sensor fusion block diagram shown in Fig. 8.2. In this case, the

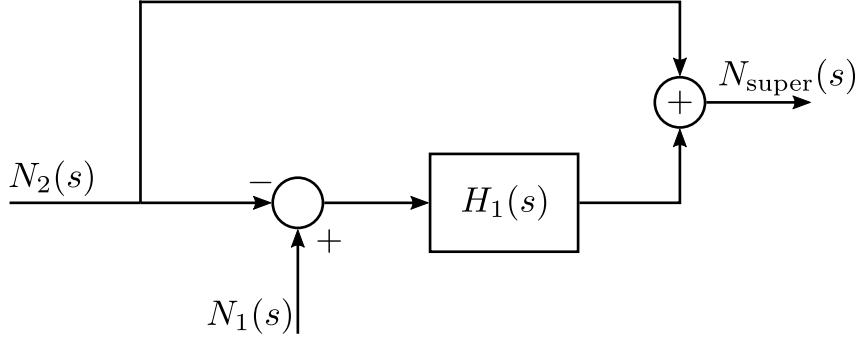


Figure 8.2: Modified sensor fusion configuration from Fig. 7.6.

super sensor noise reads

$$\begin{aligned} N_{\text{super}}(s) &= N_2(s) + (N_1(s) - N_2(s)) H_1(s) \\ &= H_1(s)N_1(s) + [1 - H_1(s)] N_2(s), \end{aligned} \quad (8.23)$$

which is exactly the same as Eqn. (7.8) with the complementary condition imposed.

Now, recall the definition of the \mathcal{H}_∞ norm in Eqn. (8.9), the inputs w of the generalized plant are assumed to be normalized. This requires the sensor noises $N_1(s)$ and $N_2(s)$ to be represented as frequency-dependent weightings in the plant. To achieve this, two transfer function models $\hat{N}_1(s)$ and $\hat{N}_2(s)$ need to be shaped and fitted to the amplitude spectral densities of the noises $S_{N_1}(\omega)$ and $S_{N_2}(\omega)$ and a method and an example for such modeling is already provided in Sec. 6.3.2. This gives $|\hat{N}_1(j\omega)| \approx S_{N_1}(\omega)$ and $|\hat{N}_2(j\omega)| \approx S_{N_2}(\omega)$. In addition, two performance weightings $W_1(s)$ and $W_2(s)$, placed right after each input, can be used to specify the required frequency-dependent attenuation of the sensor noises $N_1(s)$ and $N_2(s)$ respectively. Adding the weights and noise models to the block diagram in Fig. 8.2 and rearranging gives a generalized plant representation as shown in Fig. 8.3. Here, the normalized exogenous inputs are $w = [w_1, w_2]^\top$, the error signal is simply a scalar z , the regulator is $\mathbf{K}(s) = [H_1(s)]$ and the plant is

$$\mathbf{P}(s) = \begin{bmatrix} 0 & W_2(s)\hat{N}_2(s) & 1 \\ W_1(s)\hat{N}_1(s) & -W_2(s)\hat{N}_2(s) & 0 \end{bmatrix}. \quad (8.24)$$

This gives a closed-loop transfer matrix from w to z of

$$\mathbf{G}(s) = \begin{bmatrix} H_1(s)W_1(s)\hat{N}_1(s) & [1 - H_1(s)]W_2(s)\hat{N}_2(s) \end{bmatrix}. \quad (8.25)$$

With the closed-loop transfer matrix derived, it is a good opportunity to take a detour and briefly discuss \mathcal{H}_2 optimization instead. The generalized plant in Fig. 8.3 is perfectly




Figure 8.3: Generalized plant representation for sensor fusion.

compatible with \mathcal{H}_2 control. And, the \mathcal{H}_2 norm of the closed-loop transfer matrix is

$$\begin{aligned} \|\mathbf{G}(s)\|_2 &= \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{tr} [\mathbf{G}(j\omega)\mathbf{G}(j\omega)^H] d\omega} \\ &= \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \left| H_1(j\omega)W_1(j\omega)\hat{N}_1(j\omega) \right|^2 + \left| [1 - H_1(j\omega)] W_2(j\omega)\hat{N}_2(j\omega) \right|^2 d\omega}. \end{aligned} \quad (8.26)$$

Clearly, under the special case $W_1(s) = W_2(s) = 1$, the \mathcal{H}_2 norm is, with a factor of 2π , equivalent to the expected RMS value of the super sensor noise. This is worth noting because, in some applications, the expected RMS is required to be minimized. In this case, \mathcal{H}_2 synthesis would be superior to methods that use numerical optimization techniques to design filters with fixed structure such as that in Ref. [59].

As for the \mathcal{H}_∞ case, the \mathcal{H}_∞ norm here is simply

$$\begin{aligned} \|\mathbf{G}(s)\|_\infty &= \max_{\omega} \bar{\sigma}(\mathbf{G}(j\omega)) \\ &= \max_{\omega} \left(\max_i \sqrt{\lambda_i(\mathbf{G}(j\omega)^H \mathbf{G}(j\omega))} \right) = \max_{\omega} \left(\max_i \sqrt{\lambda_i(\mathbf{G}(j\omega)\mathbf{G}(j\omega)^H)} \right) \\ &= \max_{\omega} \sqrt{\left| H_1(j\omega)W_1(j\omega)\hat{N}_1(j\omega) \right|^2 + \left| [1 - H_1(j\omega)] W_2(j\omega)\hat{N}_2(j\omega) \right|^2}. \end{aligned} \quad (8.27)$$

With the special case $W_1(s) = W_2(s) = 1$, the cost function is simply the peak amplitude spectral density of the super sensor noise, which is the quadrature sum of of the filtered noise $H_1(s)N_1(s)$ and $[1 - H_1(s)]N_2(s)$. Minimizing such cost function may not be in-

interesting because the peak amplitude spectral density is dominated by the low-frequency noise of the inertial sensor and the noise features with much lower order of magnitude are ignored. For active isolation systems in gravitational-wave detectors, the high-frequency noise is also important even if they are orders of magnitude smaller than that at lower frequencies.

Instead of the absolute magnitude of the noise spectrum, another useful quantity to consider would be the amplification relative to the lower bound of the super sensor noise. Eqn. (8.27) has a component in the form of a quadrature sum $\sqrt{a^2 + b^2}$, where a and b are real numbers. When a and b has different values, the quadrature sum is typically dominated either a or b so it can be approximated by $\max\{a, b\}$, with a maximum factor of $\sqrt{2}$ error in the case of $a = b$, which is insignificant in log-scale. In the case of sensor fusion, the maximum error locates at the cross-over frequency of the two sensing noises, which only occurs at one frequency. At other frequencies, the difference in the magnitude of the sensing noises is typically measured by orders of magnitude (logarithmic scale). Therefore, at most frequencies, the \mathcal{H}_∞ norm can be approximated as

$$\|\mathbf{G}(s)\|_\infty \approx \max_{\omega} \left(\left| H_1(j\omega)W_1(j\omega)\hat{N}_1(j\omega) \right| \right) \quad (8.28)$$

for $|H_1(j\omega)W_1(j\omega)\hat{N}_1(j\omega)| > |[1 - H_1(j\omega)]W_2(j\omega)\hat{N}_2(j\omega)|$ and

$$\|\mathbf{G}(s)\|_\infty \approx \max_{\omega} \left(\left| [1 - H_1(j\omega)]W_2(j\omega)\hat{N}_2(j\omega) \right| \right) \quad (8.29)$$

for $|H_1(j\omega)W_1(j\omega)\hat{N}_1(j\omega)| < |[1 - H_1(j\omega)]W_2(j\omega)\hat{N}_2(j\omega)|$. Suppose the \mathcal{H}_∞ synthesis computes the complementary filter $H_1(s)$ such that $\|\mathbf{G}(s)\|_\infty = \gamma$, this yields an important result similar to that of a mixed-sensitivity \mathcal{H}_∞ control [63] where the optimal complementary filters $H_1(s)$ and $[1 - H_1(s)]$ gives

$$\left| H_1(j\omega)\hat{N}_1(j\omega) \right| \leq \gamma |W_1(j\omega)|^{-1} \quad (8.30)$$

and

$$\left| [1 - H_1(j\omega)]\hat{N}_2(j\omega) \right| \leq \gamma |W_2(j\omega)|^{-1} . \quad (8.31)$$

Here, $H_1(s)\hat{N}_1(s)$ and $[1 - H_1(j\omega)]\hat{N}_2(s)$ are the filtered sensing noises, and $\gamma|W_1(j\omega)|^{-1}$ and $\gamma|W_2(j\omega)|^{-1}$ are the upper bounds of the filtered noises. This indicates that the weighting functions $W_1(s)$ and $W_2(s)$ can be used to specify inverse of the target spectral shape of the two filtered sensing noises $H_1(s)N_1(s)$ and $[1 - H_1(s)]N_2(s)$. And, γ value indicates the deviation of the super sensor noise from the specifications. Ideally, the super sensor noise should approach the lower bound, which is given by $\min\{|\hat{N}_1(j\omega)|, |\hat{N}_2(j\omega)|\}$. To achieve this, the attenuation of one sensing noise $N_1(s)$ (or $N_2(s)$) is optimal when it is suppressed to the level of the other sensing noise $N_2(s)$ (or $N_1(s)$). Further attenuation

is unnecessary as the super sensor noise would be dominated by either one of them. This gives the optimal weightings

$$W_1(s) = \frac{1}{\hat{N}_2(s)} \quad (8.32)$$

and

$$W_2(s) = \frac{1}{\hat{N}_1(s)}. \quad (8.33)$$

As can be seen, the weightings only depend on the sensing noises themselves. This is one major advantage of the approach as it eliminates the design freedom the control filter, which could give rise to inconsistency and sub-optimality.

Alternatively, if there exists requirements for the sensing noises, then the weighting functions can be set to the inverse of the requirements, which is straightforward. However, the requirement for the sensing noise of an active isolation system is usually not specified. Rather, the control noise in terms of the displacement of the optics is usually specified. The control noise depends on other parts of the control systems, such as the controller, which has not been optimized yet. Therefore, the sensing noise requirements cannot be specified directly from the control noise requirements.

It is worth noting that the weighting functions and the noise models in the generalized plant need to be stable and proper. This is required by the general assumptions as stated in Sec. 8.1.2 [63]. This restricts the transfer functions of the noise models $\hat{N}_1(s)$ and $\hat{N}_2(s)$ need to be stable and proper. In addition, there are two more restrictions for the noise models if the weights are set according to Eqn. (8.32) and (8.33) since the reciprocals are invoked. That is, they need to be minimum phase, meaning that they must contain no zeros on the right-half plane. Also, the order of the numerator polynomial needs to be the same as that of the denominator polynomial. This means that the transfer functions $\hat{N}_1(s)$ and $\hat{N}_2(s)$ can only have flat frequency responses at low and high frequencies, i.e. outside the frequency band of interests.

\mathcal{H}_∞ (and \mathcal{H}_2) sensor fusion of typical LVDT and geophone

To demonstrate the \mathcal{H}_∞ sensor fusion method, a sensor fusion configuration combining the KAGRA typical LVDT and geophone are considered here. The typical noise spectrum of the LVDT is

$$S_{\text{LVDT}}(f) = \left[\left(\frac{10^{-2.07}}{f^{0.5}} \right)^2 + \left(\frac{10^{-2.3}}{f^0} \right)^2 \right]^{\frac{1}{2}} \frac{\mu\text{m}}{\sqrt{\text{Hz}}} \quad (8.34)$$

and that of the geophone is

$$S_{\text{geophone}}(f) = \left[\left(\frac{10^{-5.46}}{f^{3.5}} \right)^2 + \left(\frac{10^{-5.23}}{f} \right)^2 \right]^{\frac{1}{2}} \frac{\mu\text{m}}{\sqrt{\text{Hz}}}, \quad (8.35)$$

where $S_{\text{LVDT}}(f)$ and $S_{\text{geophone}}(f)$ are the amplitude spectral densities of the LVDT and geophone noises respectively. These are obtained from the empirical noise model described by Eqn. (6.72) and the values are obtained from Ref. [29] using a graphical method.




Figure 8.4: Typical noise spectra and modeling of LVDT and geophone. Blue solid: LVDT noise. Orange dashed: LVDT noise transfer function model. Green dash-dot: Geophone noise. Red dotted: Geophone noise transfer function model.

These noise spectra are modeled as transfer functions to be used as $\hat{N}_1(s)$ and $\hat{N}_2(s)$ in Fig. 8.3. The noise spectra and the models are shown in Fig. 8.4. Here, the method described in Sec. 6.3.2 is used to model the LVDT and geophone noise spectra. The noise spectra are modeled as ZPK models as an intermediate step using a global optimization algorithm. The orders of the ZPK model are chosen to be 3 and 4 for LVDT noise and geophone noise, respectively. The ZPK models are converted into transfer function models, which are used as the initial guess for transfer function modeling using a local optimization algorithm. Here, the frequency series of the noise spectra are generated in log-space between 10^{-4} Hz and 10^2 Hz. But, for the actual modeling, the noise spectra below 10^{-3} Hz and above 10^1 Hz are padded with the edge values of the noise spectra, as shown in Fig. 8.4. This is to ensure that the noise models and their inverses are proper so that can be used as weights $W_1(s)$ and $W_2(s)$ in Fig. 8.3. After fitting the transfer function models, all unstable zeros and poles, if exist, are converted stable ones by negating their real parts. As can be seen, the magnitude of the transfer function models fit the noise spectra well within the frequency band of interest and are flat outside the frequency band of interest.

With the noise models obtained, the complementary filters (blue solid and orange dashed) are optimized using the \mathcal{H}_∞ method as described above (Sec. 8.2.1). The final




Figure 8.5: Complementary filters for sensor fusion of KAGRA typical LVDT and geophone. Blue solid: \mathcal{H}_∞ low-pass filter for LVDT. Orange dashed: \mathcal{H}_∞ high-pass filter for geophone. Green dash-dot: \mathcal{H}_2 low-pass filter for LVDT. Red dotted: \mathcal{H}_2 high-pass filter for geophone.

\mathcal{H}_∞ norm in this case is $\gamma = 4.82$. The complementary filters for the sensor fusion of the LVDT and geophone are shown in Fig. 8.5. Note that the actual complementary filter synthesized is the LVDT low-pass filter (blue solid line). The geophone high-pass filter (orange dashed line) is obtained via the complementary condition $H_2(s) = 1 - H_1(s)$. Both complementary filters have an order of 13.

The obtained complementary filters are used to predict the super sensor noise, which is a quadrature sum of the two filtered noises $H_1(s)N_1(s)$ and $H_2(s)N_2(s)$, assuming that they are uncorrelated. The result is denoted as \mathcal{H}_∞ super sensor noise (Green dash-dot line) in Fig. 8.6. The noise models for LVDT and geophone are also shown in Fig. 8.6 as blue solid and orange dashed curves, respectively. As can be seen, the super sensor noise clearly attains the shape of the lower limit, indicating that both complementary filters have critically roll-offed the sensor noises. The ratio between the super sensor noise and the lower bound is averaged around 5.06, which is well within a factor of $\sqrt{2}$ of the \mathcal{H}_∞ norm (which is the aforementioned $\gamma = 4.82$).

Nevertheless, the same problem is solved using \mathcal{H}_2 synthesis with $W_1(s) = W_2(s) = 1$ as an illustration. This problem corresponds to minimizing the RMS value of the super sensor noise. The \mathcal{H}_2 complementary filters (green dash-dot and red dotted) are also shown in Fig. 8.5 and the super sensor noise (red dotted) is shown in Fig. 8.6. As can be seen, the \mathcal{H}_2 complementary filters look very different from the \mathcal{H}_∞ filters. The blending frequency of the \mathcal{H}_2 complementary filters is at around 0.3 Hz whereas that of the \mathcal{H}_∞ filters is around the cross-over frequency of the two sensor noises, i.e. ~ 0.07 Hz. The




Figure 8.6: Sensor noise spectrum. Blue solid: LVDT noise. Orange dashed: Geophone noise. Green dash-dot: \mathcal{H}_∞ super sensor noise.

\mathcal{H}_2 super sensor noise is similar to the LVDT noise, especially at frequencies where the geophone noise dominates. This is due to the fact that \mathcal{H}_2 optimization tries to ignore the geophone signals as the geophone noise has a much greater RMS value compared to that of the LVDT. In this case, the \mathcal{H}_2 super sensor has not benefited from the excellent noise performance of the geophone at high frequency. \mathcal{H}_2 synthesis with non-unity weights is not illustrated here as the weightings do not carry special meaning as in the \mathcal{H}_∞ case. This explains why the \mathcal{H}_∞ method is so unique.

8.2.2 Sensor correction problem

For sensor correction, the goal is to shape a sensor correction filter to attenuate the seismometer noise at lower frequency such that it can be used to correct relative sensors by partially removing the secondary microseism coupling at frequencies above 0.1 Hz. The sensor correction filter $H_{sc}(s)$ is, conventionally, a high-pass filter representing the attenuation of the seismometer noise $N_{seis}(s)$. On the other hand, its complementary counterpart $[1 - H_{sc}(s)]$, which is a low-pass filter, represents the transmissivity of the seismic noise $X_g(s)$. The underlying problem can almost perfectly described by a sensor fusion problem as described in Sec. 8.2.1. But, there is one subtle difference, that is, the presences of an ambient noise, which is the intrinsic noise $N_{rel}(s)$ of the relative sensor that the scheme is trying to correct. The lower limit of the sensing noise of the corrected sensor is actually $\max\{\min\{|\hat{N}_{seis}(j\omega)|, |\hat{X}_g(j\omega)|\}, |\hat{N}_{rel}(j\omega)|\}$, where $\hat{N}_{seis}(s)$, $\hat{X}_g(s)$, $\hat{N}_{rel}(s)$ are the transfer function models of the seismometer noise, seismic noise, and relative sensor

noise, respectively. In contrast, the lower limit of the sensing noise in the sensor fusion configuration is the simply minimum of the two sensor noises, $\min\{|\hat{N}_1(j\omega)|, |\hat{N}_2(j\omega)|\}$.

Contrary to the conventional belief where the seismic noise transmissivity is believed to be optimal when it is minimized, the seismic noise transmissivity does not have to be zero. This means that the optimal sensor correction filter is not necessary flat at unity at the pass band. When the seismic noise is suppressed to a level lower than the relative sensor, the corrected sensor will have a noise dominated by the relative sensor noise. Therefore, additional attenuation is not observable and does not improve the noise performance where it is expected. While the seismometer noise is typically lower than the relative sensor noise, the lower limit $\max\{\min\{|\hat{N}_{\text{seis}}(j\omega)|, |\hat{X}_g(j\omega)|\}, |\hat{N}_{\text{rel}}(j\omega)|\}$ evaluates to simply the relative sensor noise $|\hat{N}_{\text{rel}}(j\omega)|$. And, this sets the weighting function for the seismic noise path to the inverse of the relative sensor noise.

The seismic noise $X_g(s)$ can be difficult to estimate at frequencies lower than, say, 0.1 Hz since the seismometer itself is used to measure the seismic noise, which is noise-dominated at low frequencies. At such frequencies, there is no seismic isolation at all and the position control actually tries to fix the suspension on the ground. Therefore, in this sense, sensor correction is different from sensor fusion since the sensor correction filter is not optimal when it suppresses the seismometer noise to the level of the seismic noise, which would cause the suspension to drift relative to the ground. In the case of sensor correction, it is safe to say that the target attenuation of the seismometer noise is such that it approaches the relative sensor noise, not the seismic noise. Any suppression lower than that is unnecessary since the noise of the corrected sensor will be dominated by the relative sensor noise anyway. This also sets the weighting function for the seismometer noise path to the inverse of the relative sensor noise model.

In short summary, the generalized plant representation for sensor fusion in Fig. 8.3 is compatible with sensor correction and it can be used to representation the cost function for optimizing sensor correction filter. In this case, several variables need to be replaced. The sensor noises $\hat{N}_1(s)$ and $\hat{N}_2(s)$ are replaced by the seismic noise $\hat{X}_g(s)$ and seismoemter noise $\hat{N}_{\text{seis}}(s)$, respectively. The weights $W_1(s)$ and $W_2(s)$ are set to the reciprocal of the relative sensor noise, i.e. $W_1(s) = W_2(s) = 1/N_{\text{rel}}(s)$. Note that, in KAGRA, the relative sensor noise is equivalent to the LVDT noise ($N_1(s)$ in the example given in Sec. 8.2.1). It is worth noting that setting equal weights is equivalent to having a single weight for the error signal z in Fig. 8.3, which corresponds to the target specification the noise of the corrected sensor. Nevertheless, the filter $H_1(s)$ is replaced by the seismic noise transmissivity $1 - H_{\text{sc}}(s)$ and the optimized sensor correction filter is obtained via the complementary condition. The demonstration of the optimization of the sensor correction filter is given in the result section Sec. 8.3.

8.2.3 Feedback control problem

For active isolation systems, the goal of the feedback control is to design an optimal controller $K(s)$ that suppresses the system disturbance $D(s)$ while attenuating the sensing noise $N(s)$, which are conflicting objectives. The displacement of an active isolation system is what needs to be minimized. It is described by Eqn. (7.1) and is perfectly analogous to the noise of the super sensor, Eqn. (7.8), in a sensor fusion configuration. This means that the generalized plant in Fig. 8.3 can also be used to optimized feedback controllers that seek optimal trade-off between disturbance rejection and noise attenuation. In this case, the sensor noises $\hat{N}_1(s)$ and $\hat{N}_2(s)$ are replaced $-\hat{N}(s)$ and disturbance $\hat{D}(s)$, respectively, where $\hat{N}(s)$ is the sensing noise model and $\hat{D}(s)$ is the disturbance model. As for the weights, $W_1(s)$ and $W_2(s)$, they are set to the inverse of the disturbance and the inverse of the sensing noise models, respectively. \mathcal{H}_∞ synthesis then gives the optimal complementary sensitivity $T(s) \equiv K(s)P(s)/[1 + K(s)P(s)]$ and the optimal controller can be obtained via

$$K(s) = \frac{1}{P(s)} \frac{T(s)}{1 - T(s)}, \quad (8.36)$$

where $T(s)$ is the complementary sensitivity function optimized by the \mathcal{H}_∞ method and $P(s)$ is the mechanical transfer function from actuation to displacement of the suspension. Alternatively, the disturbance model $\hat{D}(s)$ and the noise model $\hat{N}(s)$ can be swapped, in which case the \mathcal{H}_∞ method optimizes the sensitivity function $S(s) \equiv 1/[1 + K(s)P(s)]$ and the optimal controller can be obtained via

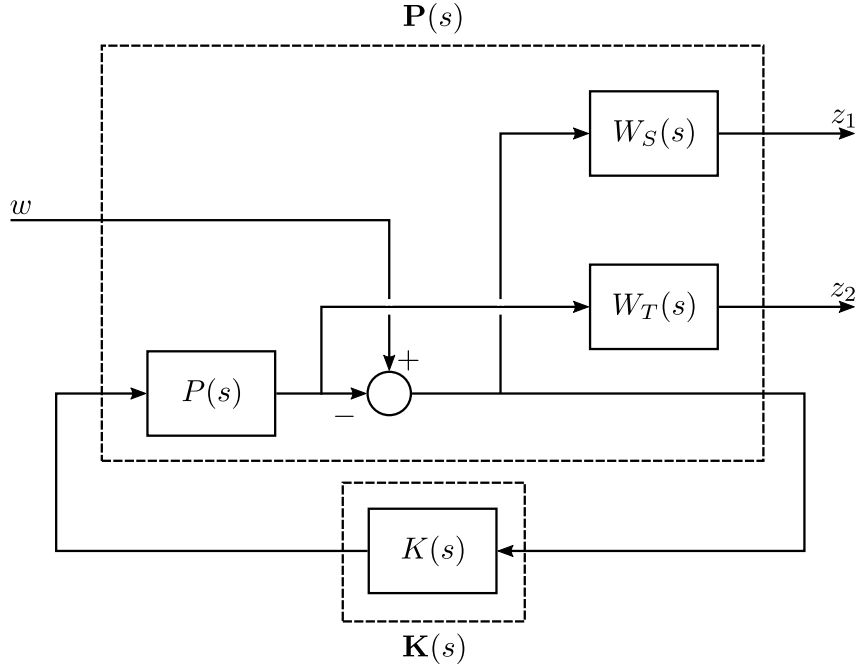
$$K(s) = \frac{1}{P(s)} \frac{1 - S(s)}{S(s)}. \quad (8.37)$$

Treating the feedback control problem as a sensor fusion problem indeed allows the controller to be optimized like complementary filters. But, the feedback control problem is so important and popular in control theory that it deserves a standalone discussion.

In fact, the \mathcal{H}_∞ method is developed based on a robust stabilization of a multivariable feedback system and the \mathcal{H}_∞ optimization of controller that guarantees control performance is well studied [63]. The complementary filter shaping method proposed in Ref. [57] is, in fact, equivalent to the S/T mixed-sensitivity \mathcal{H}_∞ control in Ref. [63]. The generalized plant for an S/T mixed-sensitivity problem is shown in Fig. 8.7. In the mixed-sensitivity problem, the goal is to shape the sensitivity functions $S(s)$ and $T(s)$ (or other sensitivity functions) to achieve desired closed-loop performance. In Fig. 8.7, the \mathcal{H}_∞ norm of the plant is

$$\left\| \begin{bmatrix} W_S(s)S(s) \\ W_T(s)T(s) \end{bmatrix} \right\|_\infty, \quad (8.38)$$

where $W_S(s)$ and $W_T(s)$ are some weighting functions. To restate, the sensitivity functions


 Figure 8.7: S/T mixed-sensitivity generalized plant.

$S(s)$ and $T(s)$ are defined as

$$S(s) \equiv \frac{1}{1 + K(s)P(s)} \quad (8.39)$$

and

$$T(s) \equiv \frac{K(s)P(s)}{1 + K(s)P(s)}. \quad (8.40)$$

With the sensitivity functions defined, the displacement of a feedback controlled suspension in Eqn. (7.1) can be rewritten as

$$X(s) = S(s)D(s) - T(s)N(s). \quad (8.41)$$

Again, the controller in the generalized plant in Fig. 8.7 is such that

$$\left\| \begin{bmatrix} W_S S(s) \\ W_T T(s) \end{bmatrix} \right\|_\infty < \gamma. \quad (8.42)$$

This corresponds to having

$$|S(j\omega)| < \gamma |W_S(j\omega)|^{-1}, \quad (8.43)$$

and

$$|T(j\omega)| < \gamma |W_T(j\omega)|^{-1}, \quad (8.44)$$

where γ is the minimum \mathcal{H}_∞ norm, and the right hand side of the inequalities denotes the frequency-dependent upper bound of the sensitivity functions. Hence, the weighting functions $W_S(s)$ and $W_T(s)$ can be used to specified the desired loop shape of the sensitivity

functions $S(s)$ and $T(s)$. While $S(s)$ and $T(s)$ are the coupling terms of the disturbance $D(s)$ and sensing noise $N(s)$ in Eqn. (8.41), the mixed-sensitivity loop-shaping technique effectively allows one to shape the frequency response of disturbance rejection and noise attenuation.

The mixed-sensitivity setting is equivalent as the complementary filter shaping technique as discussed in Ref. [57] with the sensitivities substituted by the complementary filters. One problem with this configuration is that there exist design freedoms, which are the weighting functions. Classically, this is the way mixed-sensitivity problems are meant to be tackled, assuming that the control designer knows what are the desired loop shape of the sensitivities. But, if there are no specific requirements on the control performance, the weighting functions can be ambiguous and it is not obvious what the optimal weights are, simply by inspecting the mixed-sensitivity generalized plant. By setting up the sensor fusion problem in Fig. 8.3, this problem has been solved by introducing noise models $\hat{N}_1(s)$ and \hat{N}_2 (which are actually weights) and performance weightings $W_1(s)$ and $W_2(s)$ into the plant. This is similar to the signal-based \mathcal{H}_∞ control in Ref. [63] where multiple weights are used to characterized the frequency content of input signals and closed-loop performance of a feedback system. In such a case, the selection of weights are obvious and intuitive, as discussed in Sec. 8.2.2.

Like sensor fusion, a equivalent generalized plant can be defined for the feedback control system in Fig. 5.5. The generalized plant is shown in Fig. 8.8. The approach is




Figure 8.8: Generalized plant representation for feedback control system.

actually identical to the that described in the beginning of this section where the sensor fusion plant is used as an analogy. In this case, the actuation plant $P(s)$ is embedded into

the generalized plant $\mathbf{P}(s)$ and the \mathcal{H}_∞ method optimizes the controller $K(s)$ directly. The generalized plant in this case reads

$$\mathbf{P}(s) = \begin{bmatrix} W_D(s)\hat{D}(s) & 0 & P(s) \\ -W_D(s)\hat{D}(s) & -W_N(s)\hat{N}(s) & -P(s) \end{bmatrix}, \quad (8.45)$$

where $W_D(s)$ and $W_N(s)$ are weighting functions that denote the target noise level of the disturbance $D(s)$ and noise $N(s)$, and again, $\hat{D}(s)$ and $\hat{N}(s)$ are the transfer function models of the disturbance and the sensing noise, respectively. The closed-loop transfer matrix is

$$\mathbf{G}(s) = \begin{bmatrix} \frac{1}{1+K(s)P(s)}W_D(s)\hat{D}(s) & -\frac{K(s)P(s)}{1+K(s)P(s)}W_N(s)\hat{N}(s) \end{bmatrix}. \quad (8.46)$$

The cost function is identical that in the sensor fusion case with the variables substituted accordingly. A similar analysis would show that the weighting functions are optimal when $W_D(s) = 1/\hat{N}(s)$ and $W_N(s) = 1/\hat{D}(s)$ such that the optimal controller gives an closed-loop displacement that is close to the lower limit at all frequencies.

It is worth mentioning that so far the feedback control problem has been treated as a disturbance rejection and noise attenuation problem. Indeed, under this setting, feedback controllers in active isolation systems can be optimized to achieve seismic isolation while minimizing control noise that could contaminate the gravitational-wave detector readout. However, recall in Sec. 6.4, one purpose of the controller is to coarsely align the optics by regulating the displacement to a reference position $R(s)$. The position setpoint is a static value at DC. In Sec. 6.4.2, it is discussed that the corresponding controller that can achieve position control is an integral controller k_i/s , which could be missing in the \mathcal{H}_∞ -optimal controller for the generalized plant in Fig. 8.8.

There are several ways to remedy the lack on integral action in the \mathcal{H}_∞ controller. First of all, if the \mathcal{H}_∞ controller has a flat response at with 0° phase at low frequency, the controller can be modified by multiplying a pre-compensator $(s + \omega_0)/s$. The magnitude response of the resulting controller is bend with a slope of $1/\omega$ at frequencies lower than ω_0 , which corresponds to an integral action in the control. Roughly speaking, the integration time constant is inverse of the unity gain frequency of low frequency asymptote of the open-loop transfer function (the bent section corresponding to the intergral action) Therefore, the ω_0 can be chosen to meet the time constant requirement. If the open-loop transfer function with the \mathcal{H}_∞ controller does not have a flat response with 0° phase at low frequency, the control designer must design a pre-compensator that converts the loop shape such that the integral term dominates at low frequency. However, this requires experience and the selection of the pre-compensator can vary a case by case and is not further discussed here.

Alternatively, one can choose to include the position control objective in the \mathcal{H}_∞ problem such that the optimization yields a controller that provides the required loop

gain at low frequency. The S/T mixed-sensitivity setting in Fig. 8.7 can also be used to solve, not just disturbance rejection and noise attenuation problems, but also tracking and noise attenuation problems [63]. This is because the coupling term of the reference $R(s)$ in Eqn. (5.1) is also the complementary sensitivity $T(s) = K(s)P(s)/[1 + K(s)P(s)]$. In order for the control system to track $R(s)$ at low frequency, the complementary sensitivity needs to be as close to unity as possible. While the complementary sensitivity weighting $W_T(s)$ may be used already to specify noise attenuation requirement, the sensitivity weight $W_S(s)$ can be used instead. As the complementary sensitivity approaches unity $T(s) \rightarrow 1$, the sensitivity approaches zero $S(s) \rightarrow 0$ as they are complementary. This means that for the low frequency tracking performance due to an integral action can be specified by setting the sensitivity function with an term proportional to the integrator $1/s$. However, as mentioned previously, the weightings in the generalized plant must be all be stable. This means that the integrator must be instead approximated with $1/(s + \epsilon)$, where $\epsilon \ll 1$ [63]. This results in a loop shape that practically behave like integrator at low frequency but is flat at frequencies below $\omega = \epsilon$. Similarly, in the signal-based setting in Fig. 8.8, the product of the weightings $W_D(s)\hat{D}(s)$, which is equivalent to $W_S(s)$, can be used to specify the tracking requirement by introducing a pre-compensator $(s + \omega_0)/(s + \epsilon)$. Yet another alternative is to introduce an additional path in Fig. 8.8 corresponding to the reference $R(s)$ with a dedicated weight $W_R(s)$. However, this greatly complicates the design of the weighting functions and therefore is not discussed here.

The \mathcal{H}_∞ controller can be difficult to implement in reality since the central controller attains the same number of states of the generalized plant. While high-order transfer functions are required to capture the frequency content of the disturbance and sensing noise, this could result in high-order controllers, which is impractical [63]. Therefore, actual implementation of \mathcal{H}_∞ controllers may need post-processing such as model reduction. To exemplify the \mathcal{H}_∞ control method for active isolation systems in gravitational-wave detectors, it is worth illustrating the method with a simpler system such as a simple pendulum as follows.

\mathcal{H}_∞ control of a simple pendulum

In this subsection, the signal-based \mathcal{H}_∞ control is used to optimize the controller with position control² for a simple pendulum. The \mathcal{H}_∞ control is compared to the PID control method described in Sec. 6.4. Consider a simple pendulum, such as that shown in Fig. 5.4, that has actuation plant described by

$$P(s) = \frac{\omega_n^2}{s^2 + \frac{\omega_n}{q}s + \omega_n^2}. \quad (8.47)$$

²Without position control, the result would be similar to that of the sensor fusion in Sec. 8.2.1, which is not worth reiterating.

Here, the resonance frequency ω_n is chosen to be $2\pi \times 0.08$ rad/sec and the quality factor is chosen to be $q = 10$, which are both realistic values for inverted pendulums for KAGRA Type-A or Type-B suspensions. For simplicity, suppose the platform suspending the pendulum is subjected to a random motion described by a unit white noise, this gives a disturbance model of $\hat{D}(s) = P(s)$. This can be realistic as the upper stage of the pendulum is over-damped so the displacement of the upper stage approaches the sensor noise, which can be white in reality. And, for simplicity, let us also assume that the measurement noise of the pendulum displacement is also a unity white noise, this gives a noise model of $\hat{N}(s) = 1$. The spectral shapes of the disturbance (blue solid) and noise (orange dashed) are shown in Fig. 8.9.

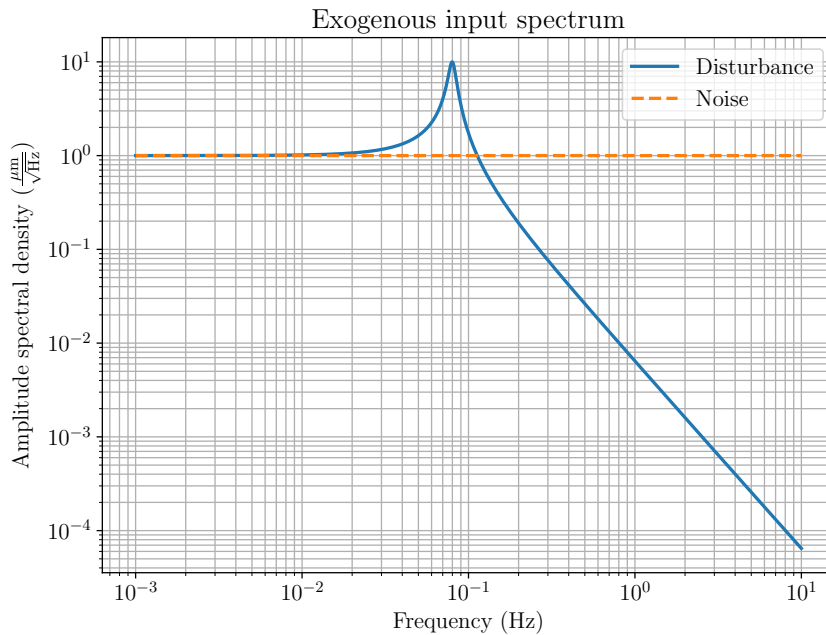


Figure 8.9: Spectrum of the exogenous inputs. Blue solid: Disturbance $\hat{D}(s)$. Orange dashed: Noise $\hat{N}(s)$.

It is optimal when the disturbance is suppressed to the level of the sensing noise and vice versa. This gives optimal weightings $W_D(s) = 1/\hat{N}(s)$ and $W_N(s) = 1/\hat{D}(s)$ as discussed. However, the choice of this disturbance model is not valid for the weight since the reciprocal is not a proper transfer function. For the purpose of synthesis, the disturbance model is actually chosen to be $\hat{D}(s) = P(s)(s/\omega_d + 1)^2$, where $\omega_d = 2\pi \times 100$ rad/sec is arbitrarily chosen to be practically high enough while not altering the dynamics around the frequency band of interest. Now, to achieve position control, a multiplicative factor $(s + \omega_0)/(s + \epsilon)$ is added to the disturbance weight such that $W_D(s) = (s + \omega_0)/[\hat{N}(s)(s + \epsilon)]$. The choice of the values ω_0 and ϵ is described later. With the weights, noise models, and actuation plant, the generalized plant for a signal-based \mathcal{H}_∞ control in Fig. 8.8 is fully defined and the \mathcal{H}_∞ -optimal controller is synthesized. The minimum \mathcal{H}_∞ norm in this case is $\gamma = 1.776$.

A PID controller with a low-pass filter is also tuned according to the plant $P(s)$ using methods in Sec. 6.4. For a better comparison, the order ($= 1$) of the low-pass is chosen such that it matches the roll-off of the \mathcal{H}_∞ controller. The low-pass is optimized for a 45° phase margin. Similarly, the choice of the ω_0 value ($= 0.3$) is chosen to roughly match the integral action in the PID controller. And, ϵ is arbitrarily chosen to be 10^{-6} , which is practically low enough to approximate an integrator.




Figure 8.10: (a) and (c): Open-loop transfer function (magnitude and phase). Blue solid: PID control. Orange dashed: \mathcal{H}_∞ control. (b) and (d): Controller (magnitude and phase). Blue solid: PID controller. Orange dashed: \mathcal{H}_∞ controller.

The PID and the \mathcal{H}_∞ controller along with the corresponding open-loop transfer functions $K(s)P(s)$ are shown in Fig. 8.10. As can be seen, the PID and the \mathcal{H}_∞ controllers share some similarity, with the PID control attaining a higher gain above ~ 0.03 Hz. This difference arose from the fact that the \mathcal{H}_∞ method takes into account the relative magnitude and make optimal trade-off between the disturbance and noise. In contrast, the PID tuning procedure does not take these exogenous inputs into account and only optimized for critical damping. In both cases, the gain margins are ∞ . This phase margins of the

PID control and the \mathcal{H}_∞ control are 45° (which is designed) and 49.93° , respectively. In this case, both systems have similar stability properties.




Figure 8.11: Closed-loop displacement spectrum. Blue solid: PID control. Orange dashed: \mathcal{H}_∞ control. Black dash-dot: Disturbance. Black dotted: Noise.

The predicted closed-loop displacement spectra for the PID control (blue solid) and \mathcal{H}_∞ control (orange dashed) are shown in Fig. 8.11. The raw disturbance (black dash-dot) and the measurement noise (black dotted) are also plotted as references. As can be seen, both controller suppressed the peak of the disturbance reasonably well. As expected, the displacement spectrum of the \mathcal{H}_∞ control follows the the lower limit with a maximum amplification roughly equal to the \mathcal{H}_∞ norm. This has important implications for active isolation in gravitational-wave detectors:

1. The passive isolation performance (denoted by the shape of the disturbance) is not ruined by the \mathcal{H}_∞ control as it pushes the control noise towards the passively isolated displacement, and
2. The disturbance is actively suppressed towards the sensing noise, which maximizes the potential of the sensors.

This is exactly what is required for active isolation in gravitational-wave detectors. On the other hand, in the case of PID control, the roll-off of the sensing noise is similar to that of the \mathcal{H}_∞ control as the order of the low-pass filter is informed by the \mathcal{H}_∞ controller. However, even so, the noise attenuation is not as good and the closed-loop displacement of the PID control appears to be amplified at high frequency compared to the \mathcal{H}_∞ controller.

Nevertheless, in the beginning of the section, it is mentioned that the \mathcal{H}_∞ needs to be able to achieve alignment control (hence the modified weight $W_D(s) = (s + \omega_0)/[\hat{N}(s)(s +$

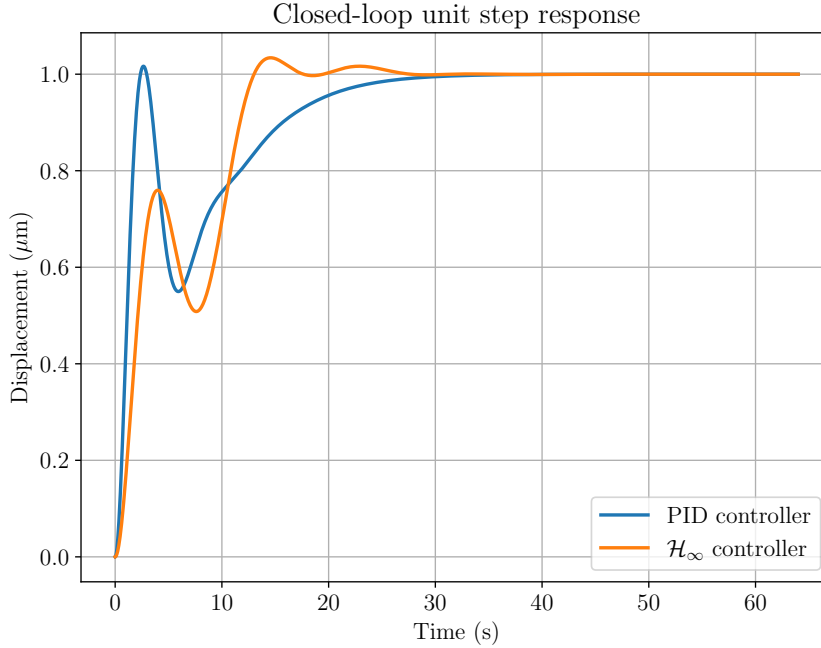


Figure 8.12: Closed-loop unit step response. Blue solid: PID control. Orange dashed: \mathcal{H}_∞ control.

ϵ)). This introduces integral action in the \mathcal{H}_∞ control at low frequency as shown in Fig. 8.10. Fig. 8.12 shows the simulated closed-loop unit step response of the PID control and \mathcal{H}_∞ control. As can be seen, both systems settle at the setpoint at around $t = 30$ s. This shows the \mathcal{H}_∞ control can be used for coarse alignment like a PID controller.

8.3 \mathcal{H}_∞ optimization of the SRM pre-isolator control filters

The \mathcal{H}_∞ methods discussed in Sec. 8.2 are tested with the KAGRA SRM suspension. The \mathcal{H}_∞ optimization of several configurations are presented here. In Sec. 8.3.2, the complementary filters for the sensor fusion of the LVDT and geophone at the preisolator are presented. The sensor correction of the pre-isolator LVDT using the seismometer placed near the IXV is shown in Sec. 8.3.3. With the corrected sensor, the complementary filters need to be re-optimized and the configuration with both sensor correction and sensor fusion is discussed in Sec. 8.3.4. For all sensor configurations above, the open-loop performances are evaluated and compared, individually, against the KAGRA super sensor or (and) the KAGRA sensor correction, which are both described in Sec. 8.3.1. The closed-loop performance is given in in Sec. 8.3.5 along with a short summary regarding the sensor configurations. Finally, in Sec. 8.3.6, \mathcal{H}_∞ control of the pre-isolator is demonstrated. For all configurations, unless otherwise specified, the corresponding degree of freedom is the longitudinal direction of the pre-isolator table, which is suspended by the inverted

pendulum. In Sec. 8.3.1, some basic information regarding the SRM pre-isolator is given, including the transfer function, default control filters, and sensor noises.

8.3.1 The SRM pre-isolator

The plant and the baseline PID controller

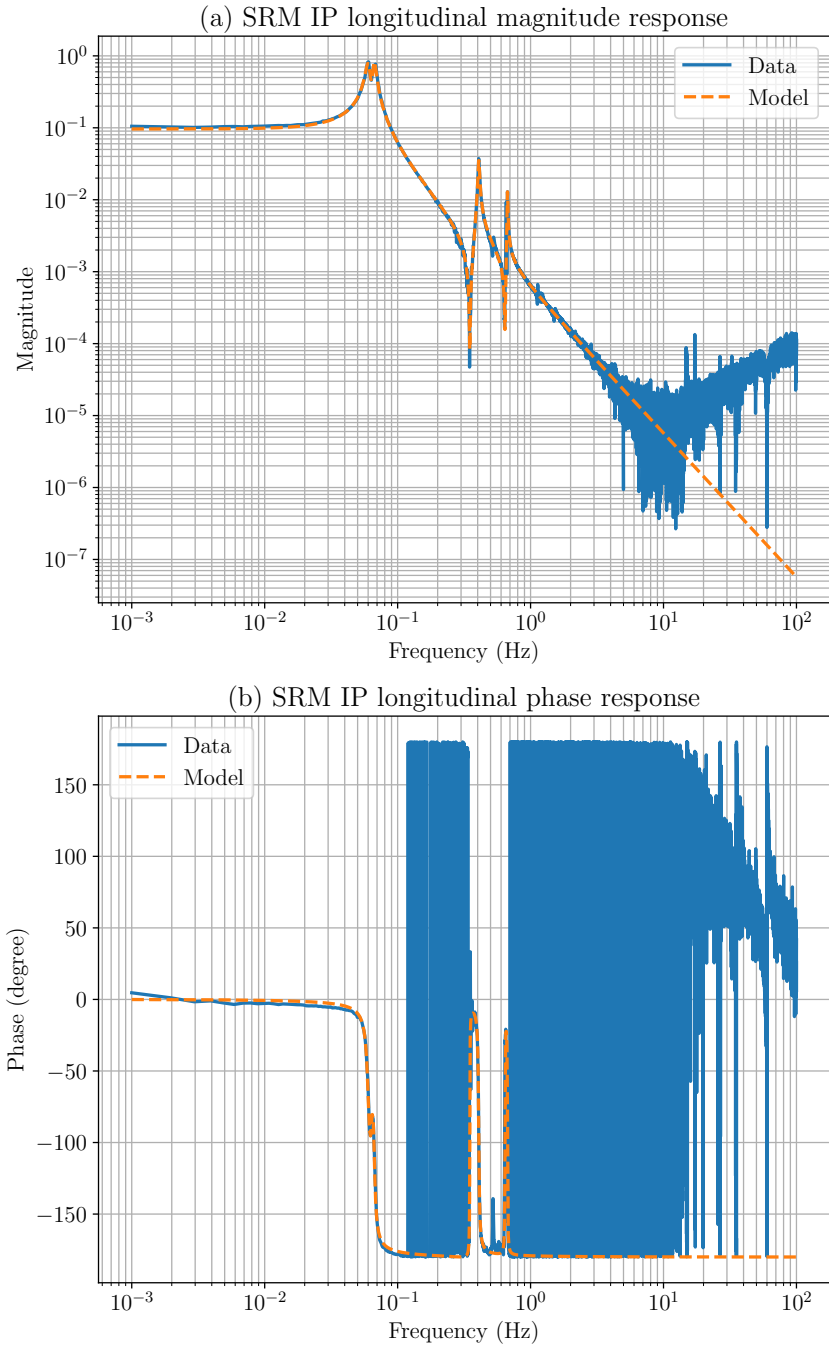


Figure 8.13: SRM longitudinal pre-isolator actuation frequency response. (a) Magnitude response. (b) Phase response. Blue solid: Measurement data. Orange dashed: Transfer function model.

Fig. 8.13 shows the SRM pre-isolator longitudinal frequency response from actuation to displacement (measured by LVDTs). The transfer function model is shown as orange dashed line in the figure and it is obtained using method discussed in Sec. 6.3.1. As can be seen, four main resonances can clearly be observed at ~ 0.06 Hz, ~ 0.067 Hz, ~ 0.4 Hz, and ~ 0.67 Hz.




Figure 8.14: SRM pre-isolator longitudinal open-loop transfer function and baseline controller. (a) and (c): Open-loop transfer function (magnitude and phase). (b) and (d): Controller (magnitude and phase).

The baseline controller for the SRM pre-isolator longitudinal degree of freedom is a PID controller with fourth-order low-pass post filtering. The controller is designed using methods in Sec. 6.4 according to the transfer function model. The open-loop transfer function and the controller is shown Fig. 8.14. The fourth-order low-pass filter has a cut-off frequency of 3.6 Hz. The gain and phase margins are 5.48 (14.78 dB) and 48.96° , respectively. Unless otherwise specified, the PID controller is used for closed-loop measurements in below sections.

Seismic and sensor noises




Figure 8.15: SRM pre-isolator LVDT and geophone noises. (a): LVDT 1 and Geophone 2. (b) LVDT 2 and geophone 2. (c) LVDT 3 and geophone 3. Blue solid: LVDT readout (pre-isolator locked). Orange dashed: Geophone readout. Green dash-dot: LVDT noise empirical model. Red dotted: Geophone noise empirical model.

Three LVDTs and three geophones are installed on the pre-isolator, each sensing a tangential displacement/velocity of the pre-isolator. The sensor noises of each sensor and their empirical models are shown in Fig. 8.15. Here, the empirical models are obtained using the frequency spectrum modeling methods discussed in Sec. 6.3.2 and the corresponding empirical model is Eqn. (6.72). For the LVDTs, the noise spectral was obtained by taking a measurement of the LVDT readouts when the pre-isolator table was locked relative to the ground. As for the geophones, the data was a measurement of the geophone readouts when the pre-isolator table was unlocked. For this reason, some reso-

nances can be observed from the spectra. To compensate for this, the weighting function in Eqn. (6.71) for modeling the noise spectrum of a geophone is set such that the weight is 0 when the coherence between any of the two geophones is higher than 0.5 and is 1 otherwise. Geophone data higher than 10 Hz are also ignored.

For each type of sensors, the three sensors are separated horizontally with 120° spacing along a common circumference on the pre-isolator table, each sensing the tangential motion of the table. To convert the individual sensor noises to the effective sensor noises in the longitudinal direction, the sensing matrices of the LVDTs and the geophones need to be invoked. The amplitude spectral density of the longitudinal sensor noise of one type of sensor can be obtained via the quadrature sum

$$N_L(f) = \left(\sum_{i=1}^3 [S_{Li}N_i(f)]^2 \right)^{\frac{1}{2}}, \quad (8.48)$$

where $N_L(f)$ is the effective amplitude spectral density of the sensor noise in the longitudinal direction, S_{Li} is the i^{th} element of the longitudinal row of the sensing matrix, and $N_i(f)$ is the amplitude spectral density of the i^{th} sensor. Note that S_{Li} for LVDTs and geophones are in general different because they are placed differently. The converted noise spectral for LVDT and geophone noises in the longitudinal direction is shown as blue solid and orange dashed curves in Fig. 8.17, respectively.




Figure 8.16: IXV seismometer noise (Y-arm direction). Blue solid: Readout. Orange dash-dot: Empirical model.

The seismometer placed at the IXV (second floor of the ITMX chamber) is considered to be a ground sensor in close proximity to the SRM suspension (both located in the corner

area of the interferometer). At low frequencies, it measures a common ground motion that is disturbing the SRM suspension and therefore can be used for sensor correction. The IXV seismometer readout measuring the seismic noise along the Y-arm direction of the interferometer (SRM is roughly aligned with the Y-arm direction.) is shown as blue solid line in Fig. 8.16. For the purpose of modeling, data at above 0.04 Hz are believed to be dominated by seismic noise and are ignored. The empirical model is then simply in the form of N_a/f^a , where N_a and a are the model parameters. The fitting empirical model is shown as orange dashed line in Fig. 8.16.




Figure 8.17: SRM pre-isolator longitudinal noises. Blue solid: LVDT noise (empirical model). Orange dashed: Geophone noise (empirical model). Green dash-dot: IXV seismometer noise (empirical model). Red dotted: KAGRA mean seismic noise.

In Fig. 8.17, a compilation of noises in the SRM pre-isolator longitudinal are shown. This includes the empirical models of the LVDT noise (blue solid), geophone noise (orange dashed dot), and the IXV seismometer noise (green dash-dot). The mean seismic noise in KAGRA taken from Fig. 7.2 is also shown. These noises are important in later sections as the transfer function models for \mathcal{H}_∞ synthesis in later sections are largely based on these four noises. Transfer function models for each noise could vary depending on the application and they shall be specified later.

KAGRA complementary filters and sensor correction filters

During the two-week joint observing run with KAGRA and GEO600 (O3GK), a gravitational-wave detector in Germany, KAGRA has suffered from high seismic activity [75, 24]. This has eventually limited the duty cycle of the detector to only 54% [24]. At the time, KAGRA's suspensions had no active isolation capability and this was due to the tight

commissioning schedule. For O3GK, the pre-isolators of the suspensions only use LVDTs, which couples perfectly to seismic noise, for damping of the resonances. Therefore, it is no surprise that the KAGRA became susceptible to seismic disturbances.

Before O3GK, it was concluded that the low-frequency noise of the geophones on the pre-isolator is too high for them to be implemented in a sensor fusion configuration for inertial damping (damping with respect to an inertial frame). To elaborate, the cross-over frequency of between the LVDT and geophone noise is close to ~ 0.2 Hz in Fig. 8.17, which is much higher than the expected ~ 0.07 Hz in Fig. 8.4. This forces a higher blending frequency of the complementary filter, which could potentially amplify the secondary microseism coupling in the LVDT readout, leading to an adverse effect. Therefore, the focus for active isolation was switched to using a sensor correction configuration. The pioneer studies of sensor correction in KAGRA is given in Ref. [31] and it was shown that the secondary microseism can be suppressed and that the sensor correction scheme can be a feasible solution to the active isolation problem. However, due to a lack of experience, the sensor correction filter used for the study is not properly design to attenuate the low-frequency noise of the seismometer. This caused low-frequency drift of the pre-isolator table when sensor correction is engaged [58] and therefore is not used during O3GK. The particular sensor correction filter is shown in subplot (a) in Fig. 8.18. As can be seen, it has a ~ 3 times amplification around 0.013 Hz and a passband above that.

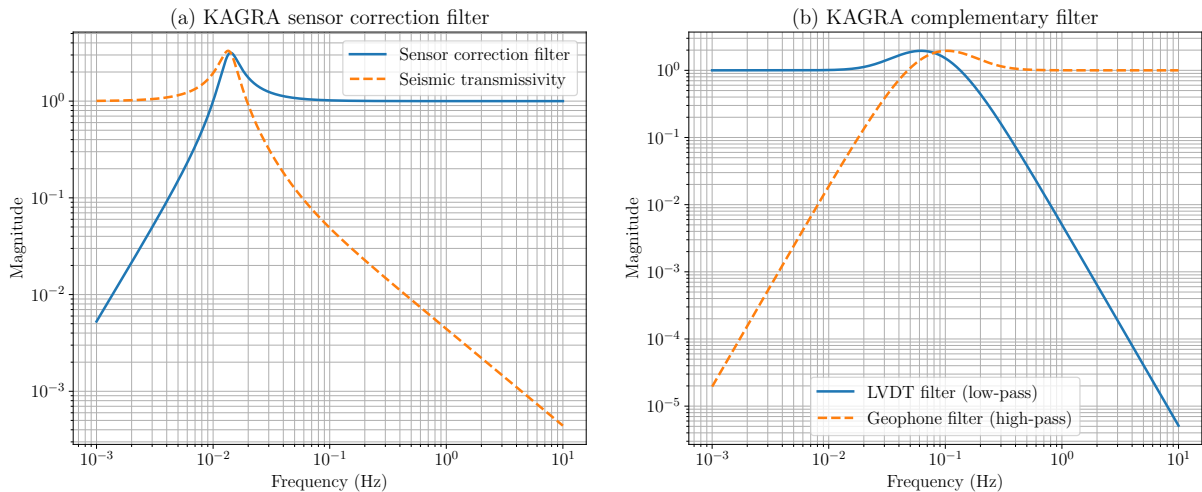


Figure 8.18: (a): KAGRA sensor correction filter. Blue solid: sensor correction filter. Orange dashed: seismic transmissivity. (b) KAGRA complementary filter. Blue solid: LVDT filter (low-pass) Orange dashed: Geophone filter (high-pass).

To remedy the issue of lacking active isolation, inertial damping was brought back during the pre-O4 commissioning [76]. As part of the suspension system unification³, a common set of complementary filters was simply set for all suspensions for the sensor fusion of LVDT and geophone. In Fig. 8.18, the KAGRA complementary filters are

³Different types of suspensions have very different control topologies back in O3.

shown in subplot (b). And, this set of filters is intended to be used during O4. The complementary filters were obtained from Ref. [28] and they are described by Eqn. (7.11) and (7.12), with $n_l = n_h = 3$. The blending frequency (for longitudinal and transverse direction) was set to 0.08 Hz.

Miscellaneous remarks

The \mathcal{H}_∞ filters that are presented in later sections are benchmarked against the original KAGRA complementary filters and sensor correction filter shown in Fig. 8.18. The performance are compared in several ways. This includes comparing the spectra of the super sensor readouts and sensor correction readouts during free wing of the pre-isolator table. These sensors are also compared in a closed-loop condition where the controller in Fig. 8.14 is engaged. In such a case, the optical lever at the optics stage is used as an out-of-loop sensor to monitor the motion of the optics.

It should be noted that, in reality, the geophone signals need to be pre-filtered by high-pass filters to avoid the overflow of digital signals. For KAGRA geophones, the geophones signals are pre-filtered by a fourth-order high-pass filter with a cut-off frequency at 0.003 Hz. Therefore, signals involving the geophone measurements be appeared to be rolled off below that frequency. With the pre-filtering, the complementary condition of the complementary filters is no longer valid. This gives a super sensor that has a distorted frequency response around the blending frequency. Therefore, the pre-filters must be taken into account and the complementary filters need to be adjusted accordingly. Refs. [28, 61] provide one way of “normalizing” the complementary filters but the resulting filters may have excessive noise amplification around the blending frequency. The treatment of pre-filtering is to be discussed in Sec. 9.1.2. But, it should be noted that the actual \mathcal{H}_∞ complementary filters implemented in the system are treated in some way to retain the complementary condition. Hence, the \mathcal{H}_∞ super sensor noise might not behave exactly like that in the simulation, but is slightly higher noise at higher frequencies. On the other hand, the KAGRA complementary filters are not treated so the KAGRA super sensor has a distorted frequency response around the blending frequency in reality.

It is worth clarifying one performance index that is commonly used for evaluating the “size” of a signal, that is, the root mean square (RMS) value. The expected RMS value of a signal can be obtained by integrating its entire power spectral density with respect to frequency and then taking the square root of that value. This gives the an expected RMS value of the signal. For active isolation systems, and for subsystems in a gravitational-wave detector in general, signals can have very different order of magnitude at different frequencies. In most cases, the amplitude of a signal at high frequency is much lower than that at low frequency. This means that the RMS value is typically dominated by the low frequency components so the high frequency components are not well represented. To remedy this, integrated RMS $S_{\text{RMS}}(f)$ is introduced as a function of frequency, and it is

defined as

$$S_{\text{RMS}}(f) = \left[\int_f^\infty S_{\text{ASD}}(f')^2 df' \right]^{\frac{1}{2}}, \quad (8.49)$$

where $S_{\text{ASD}}(f)$ is the amplitude spectral density of a signal. The integrated RMS is meant to be viewed from high frequency to low frequency and it shows how the RMS value gets accumulated, which would be useful for identifying RMS contributions from high frequency components. The terms RMS and integrated RMS are sometimes used interchangeably since they can be easily distinguished (RMS is a value whereas integrated RMS is a function of frequency).

Here are some disclaimers regarding the results in later sections. First of all, the comparison only yields relative performance. Therefore, the results do not evaluate the optimality of the control filters. Performance indices, such as the RMS value of a sensor signal, do not necessarily translate to real performance such as lock stability of the interferometer. To test the active isolation performance, all control filters of all suspension must be changed for a sufficient long period during the observation where metrics such as number of lock-losses per number of earthquakes can be observed and compared. These tests maybe too costly to be conducted for current KAGRA in such an early stage where robustness is pursued.

8.3.2 \mathcal{H}_∞ sensor fusion

In this section, the \mathcal{H}_∞ sensor fusion using the LVDT and the geophone is demonstrated. In this case, the LVDT is measuring relative displacement of between the pre-isolator and the ground. Therefore, the seismic noise is considered as part of the sensing noise in the LVDT readout. To use the \mathcal{H}_∞ method for sensor fusion in Sec. 8.3.2, the sensing noises of the LVDT and the geophone need to be modeled as transfer functions. For the geophone noise, the modeling is straightforward, as shown in sec. 6.3.2. As for the seismic noise coupled LVDT sensing noise, it is not simple since the seismic noise need to be estimated from a seismometer readout. The seismometer readout at low frequency, especially below ~ 0.1 Hz, is dominated by the seismometer noise, as shown in Fig. 8.17. This makes seismic noise at low frequency hard to estimate. As a simple estimation, the seismic noise in Fig. 8.17 is processed as follows. At frequencies lower than the local minima around ~ 0.1 Hz, the values are replaced with the value of the minima, flattening the spectrum at lower frequencies. This adds the secondary microseismic peak to the LVDT noise. The quadrature sum of the processed seismic noise and the LVDT noise is modeled as the sensing noise with a transfer function for the purpose of \mathcal{H}_∞ synthesis.

The modeled noises are used to optimized complementary filters via \mathcal{H}_∞ synthesis. The minimum \mathcal{H}_∞ norm for this case is $\gamma = 7.055$, meaning that the super sensor noise has an amplification 7.055 times above the lower boundary. The \mathcal{H}_∞ complementary filters




Figure 8.19: \mathcal{H}_∞ complementary filters for the sensor fusion of the seismic noise coupled LVDT and the geophone of the SRM pre-isolator in the longitudinal direction. Blue solid: Seismic noise coupled LVDT filter (low-pass). Orange dashed: Geophone filter (high-pass)

are shown in Fig. 8.19. As shown in the figure, the low-pass filter has a clear difference compared to KAGRA complementary low-pass filter shown in subplot (b) of Fig. 8.18. That is, there exists a notch-like feature in the \mathcal{H}_∞ around the secondary microseism. This gives a ~ 10 times attenuation at 0.2 Hz and, in comparison, the KAGRA complementary filter only offers ~ 2.5 times attenuation around the microseism. A similar feature is also engineered in the example Virgo complementary filter shown in Ref. [27] for the purpose of providing additional suppression to the microseism. However, the Virgo complementary filters are only given as an example and it was not entirely how the filter design was approached. In contrast, the \mathcal{H}_∞ method automatically generates the necessary features that are required to properly filter the sensing noises. The notch-like feature in \mathcal{H}_∞ low-pass filter in Fig. 8.19 was naturally created as a result of optimization.

The complementary filters are used to estimated the super sensor noises. The KAGRA super sensor noise (green dash-dot) and the \mathcal{H}_∞ super sensor noise (red dotted) are shown in Fig. 8.20 together with the original seismic noise coupled LVDT and geophone sensing noise models (blue solid and orange dashed, respectively). The amplitude spectral densities of the noises are shown in subplot (a) and the integrated RMS of the noises are shown in subplot (b). Note that the RMS of the geophone noise is not shown since it has a much larger value than the others and it is not an important value.

In comparison, the KAGRA super sensor has poorer noise performance compared to the \mathcal{H}_∞ super sensor. As can be seen, at frequencies below 0.05 Hz, the amplitude spectral density of the KAGRA super sensor is around an order of magnitude higher




Figure 8.20: Sensor noise and estimated super sensor noises in the SRM pre-isolator longitudinal direction. (a) Amplitude spectral density. (b) RMS. Blue solid: Seismic noise coupled LVDT. Orange dashed: Geophone. Green dash-dot: KAGRA super sensor. Red dotted: \mathcal{H}_∞ super sensor.

than that of the \mathcal{H}_∞ super sensor. This indicates that the KAGRA complementary high-pass filter does not filter the geophone noise. The \mathcal{H}_∞ super sensor has a slightly higher noise amplification around ~ 0.1 Hz, corresponding to the amplification band of the complementary filters, and rolls off along the geophone noise at higher frequencies with a slight multiplicative offset. Even with the presence of the microseism, the \mathcal{H}_∞

super sensor noise follows the shape of the lower limit well at all frequencies. On the other hand, the microseismic peak appears in the KAGRA super sensor noise so it is not filtered well. At higher frequencies, the KAGRA super sensor noise rolls off close to the geophone noise. Although the KAGRA super sensor performs better at high frequency, it introduces a significant amount of drift at low frequency compared to the original LVDT and \mathcal{H}_∞ noise while providing little seismic noise attenuation. As shown in subplot (b) of Fig. 8.20, the RMS value of the KAGRA super sensor noise is $0.8694 \mu\text{m}$, which is the highest among those of the original LVDT noise and the \mathcal{H}_∞ super sensor noise, which are $0.2485 \mu\text{m}$ and $0.3578 \mu\text{m}$, respectively.

The \mathcal{H}_∞ complementary filters are implemented to the SRM pre-isolator and open-loop spectra were simultaneously measured for the KAGRA and \mathcal{H}_∞ super sensors. In this particular case, there was a mistake when synthesizing the \mathcal{H}_∞ -optimal complementary filters. The intrinsic noise of the LVDT was amplified by a factor of ~ 3.56 during the modeling. The actual noise model used is not that one shown in Fig. 8.20 and the actual complementary filters are slightly different from that in Fig. 8.19. The \mathcal{H}_∞ norm corresponding to the actual sensor fusion configuration is $\gamma = 8.913$, which correspond to a slightly higher noise amplification. Therefore, the measurements here are for references only.

Fig. 8.21 shows the open-loop spectra of the KAGRA and \mathcal{H}_∞ super sensors. The LVDT and geophone readouts are also shown as references. The spectra correspond to measured displacement of the pre-isolator table so they do not represent the sensing noises exactly. However, any observable differences between the spectra are explained by difference in sensing noise. Therefore, the open-loop spectra can be compared this way.

As can be seen, similar observations from Fig. 8.20 can be made from the real measurements in Fig. 8.21. At low frequencies, the KAGRA super sensor has a higher noise compared to that of the \mathcal{H}_∞ super sensor. As a result, the KAGRA super sensor readout has an RMS value of $0.55 \mu\text{m}$, which is highest among the LVDT and the \mathcal{H}_∞ super sensor. The RMS value of the LVDT readout and the \mathcal{H}_∞ super sensor readout are $0.342 \mu\text{m}$ and $0.356 \mu\text{m}$, respectively. Around 0.1 Hz , the \mathcal{H}_∞ super sensor has a noise amplification that corresponds to the peak of the complementary filters. The \mathcal{H}_∞ super sensor has a better seismic attenuation at 0.2 Hz and a similar performance at 0.3 Hz compared to the KAGRA super sensor. In this measurement, the microseismic peak shifted from the expected 0.2 Hz to $\sim 0.3 \text{ Hz}$, which means the \mathcal{H}_∞ super sensor is operating sub-optimally and this explains why the difference in seismic attenuation between the two super sensors is subtle. At above 0.7 Hz , the \mathcal{H}_∞ super sensor has a clear noise amplification relative the geophone noise. This is a consequence of the compensation of the geophone pre-filter, which is to be discussed in Sec. 9.1.2. It should be noted that the signal-to-noise ratio of both super sensors around the first resonance is very low. This means that the super sensors cannot be used to suppress the main resonance motion. This makes the super sensor




Figure 8.21: SRM pre-isolator longitudinal freeswing readout. (a) Amplitude spectral density. (b) RMS. Blue solid: LVDT (seismic noise coupled). Orange dashed: Geophone. Green dash-dot: KAGRA super sensor. Red dotted: \mathcal{H}_∞ super sensor.

configuration, using only the seismically-coupled LVDT and geophone, not very suitable for damping the resonances. This also explains why inertial damping was abandoned during pre-O3 commissioning and focus was switched to sensor correction instead.

8.3.3 \mathcal{H}_∞ sensor correction

It is tempting solve the sensor correction problem by treating it exactly as a sensor fusion problem. This is because the form of sensor correction noise in Eqn. (7.18) looks exactly like the super sensor noise in Eqn. (7.8). To show why this is sub-optimal, consider the seismic noise model (blue solid) and the seismometer noise model (orange dashed) in Fig. 8.22. Here, the seismic noise model and the seismometer noise are modeled against




Figure 8.22: SRM pre-isolator longitudinal sensor correction noise models and estimated sensor correction noises. Blue solid: Seismic noise. Orange dashed: Seismometer. Green dash-dot: LVDT. Red dotted: Trial \mathcal{H}_∞ sensor correction. Lilac dash-dot-dot: Modified seismic noise. Brown dash-dot-dot-dot: \mathcal{H}_∞ sensor correction

the KAGRA mean seismic noise and the IXV seismometer noise shown in Fig. 8.17, respectively. Plugging these two noises in to the generalized plant Fig. 8.3 and solving for a “trial” sensor correction filter gives a sensor correction noise (red dotted) shown in Fig. 8.22. Note that the trial version of sensor correction is shown as a demonstration only and this configuration is not intended to be used. The trial sensor correction filter (blue solid) and the corresponding seismic noise transmissivity (orange dashed) are plotted in subplot (a) of Fig. 8.24 as a reference.

Firstly, from Fig. 8.22, the flat region of the seismic noise at frequencies below 0.1 Hz is not an estimation of the seismic noise as it was artificially made flat to neglect the seismometer noise in the measurement. This means that the seismometer noise in the sensor correction readout might be under-suppressed and the extent depends on how the data is treated. At frequencies above ~ 0.04 Hz, the sensor correction noise follows the spectral shape of the seismometer noise, indicating an “optimal” roll off. However, ultimately, the readout of the sensor corrected LVDT in Eqn. 7.17 is used for feedback




Figure 8.23: SRM pre-isolator longitudinal trial sensor correction filter and seismic transmissivity. Blue solid: Sensor correction filter. Orange dashed: Seismic transmissivity.

control and it has a sensing noise composed of the relative sensor noise $N_{\text{rel}}(s)$ and the sensor correction noise $N_{\text{sc}}(s)$. With the relative sensor noise (green dash-dot), which is the LVDT noise, also plotted in Fig. 8.22, it is clear that the seismic noise is over-suppressed. This is indicated by the fact that the sensor correction noise goes lower than the LVDT noise at frequencies above ~ 0.13 Hz. At those frequencies, the corrected LVDT readout is dominated by the intrinsic noise of the LVDT and further suppression of the seismic noise is not useful. This falsifies the conventional method of requiring the seismic noise transmissivity to be as low as possible, i.e. the sensor correction filter does not have to be exactly 1 at the pass band.

To resolve the issues above, the seismic noise data in Fig. 8.17 is modified as follows. For low frequency data, instead of flattening the curve below the frequency of the first minima at frequencies below the secondary microseism, the curve was extrapolated. The extrapolation for the i^{th} data point is done by taking the average between the $(i+1)^{\text{th}}$ data point of the seismic noise data and the i^{th} value of the LVDT noise. This goes with decreasing values of i , starting from the index just below the local minima next to the secondary microseism. When the value of the i^{th} seismic noise data point is higher than that of the $(i+1)^{\text{th}}$ data point, it is set to the value of the $(i+1)^{\text{th}}$ data point. This flattens the curve at low frequency, making it lower than the LVDT noise. As for high frequency data, the values of the seismic noise data points at frequencies higher than the microseismic peak are set to the value of the microseismic peak, making a flat line at high frequency. The purpose of this is to force the \mathcal{H}_∞ sensor correction filter to be a high-pass filter. Without this modification, the sensor correction filter will amplify both

seismic noise and seismometer noise at high frequency if the target noise level is specified to the level of the LVDT noise. This has one benefit: the sensor correction performance is not sensitive to a change in microseism frequency.




Figure 8.24: SRM pre-isolator longitudinal sensor correction filter and seismic transmissivity. Blue solid: Sensor correction filter. Orange dashed: Seismic transmissivity.

The corresponding transfer function model of the modified seismic noise is shown as lilac dash-dot-dot curve in Fig. 8.22. As can be seen, the LVDT noise becomes higher than the minimum of the seismometer noise and the modified seismic noise at most frequencies. Therefore, it can be used as the weights for specifying the upper bounds of both the suppressed seismic noise and the filtered seismometer noise, as discussed in Sec. 8.2.2. The sensor correction filter was optimized using the \mathcal{H}_∞ method discussed in Sec. 8.2.2. The \mathcal{H}_∞ sensor correction filter (blue solid) and the corresponding seismic transmissivity (orange dashed) are shown in subplot (b) of Fig. 8.24. As can be seen, at high frequency, the amplitudes of the seismic transmissivity and the sensor correction filter plateau at ~ 0.055 and ~ 0.945 , respectively.

The resulting \mathcal{H}_∞ norm of the optimization is $\gamma = 3.56$. This means that the resultant sensor correction noise can be estimated as 3.56 times that of the LVDT noise⁴, except at frequencies higher than the microseismic peak. The estimated sensor correction noise is shown in Fig. 8.22 as brown dash-dot-dot-dot line. As can be seen, the estimated sensor correction noise follows the shape of the LVDT noise roughly, with a multiplicative offset of ~ 3.56 . Note that the sensor correction noise is estimated using the modified seismic noise model, which is not true at frequencies higher than the microseism. At those

⁴This explains the mistake made earlier in Sec. 8.3.2.

frequencies, the sensor correction noise would be lower. So this means that the sensing noise of the corrected LVDT would converge to the intrinsic noise of the LVDT at those frequencies.

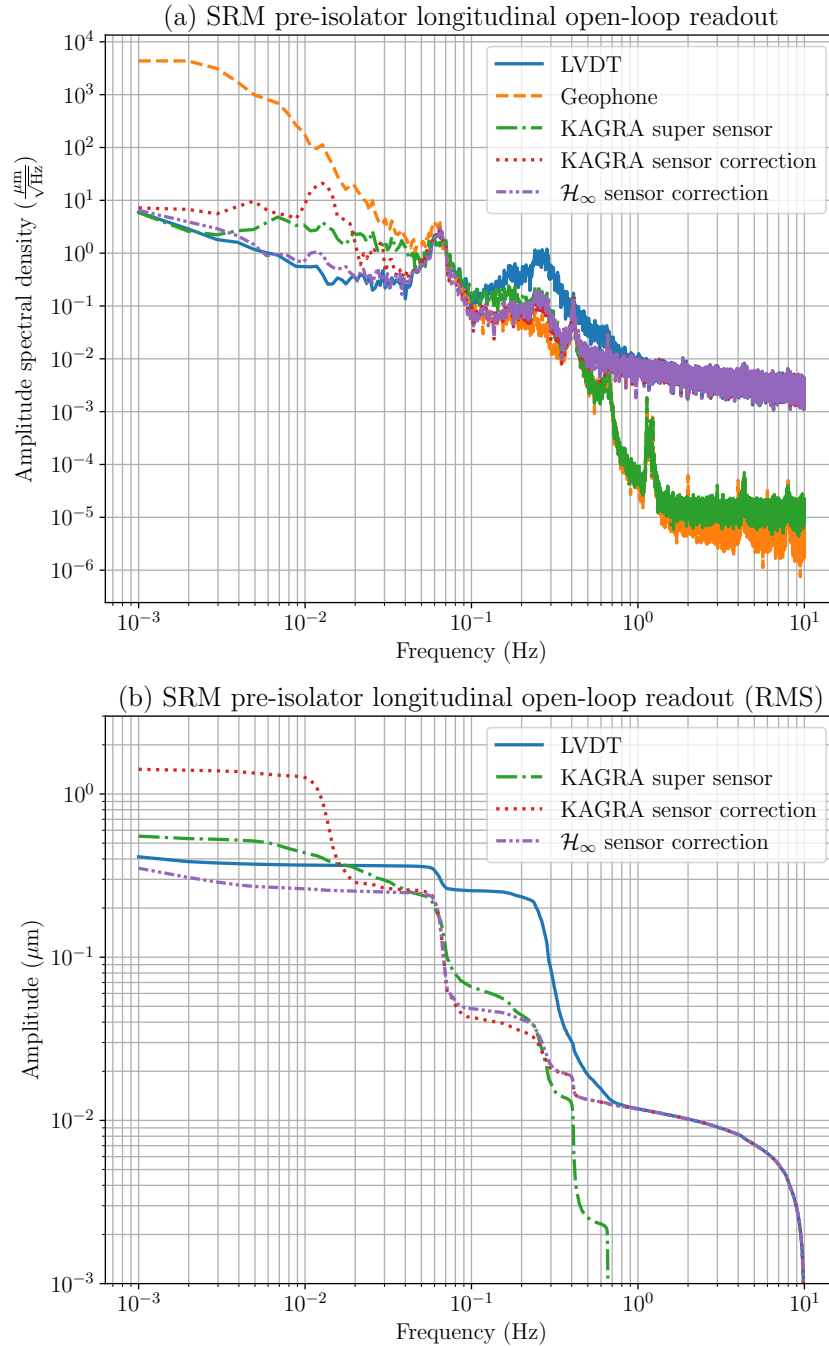


Figure 8.25: SRM pre-isolator longitudinal open-loop readout. (a) Amplitude spectral density. (b) RMS. Blue solid: LVDT. Orange dashed: Geophone. Green dash-dot: KAGRA super sensor. Red dotted: KAGRA sensor correction. Lilac dash-dot-dot: \mathcal{H}_∞ sensor correction.

The \mathcal{H}_∞ sensor correction filter was put into test in the SRM pre-isolator. Fig. 8.25 shows a comparison between the open-loop readouts of the KAGRA super sensor (green

dash-dot), KAGRA sensor correction (red dotted), and the \mathcal{H}_∞ sensor correction (lilac dash-dot-dot). The LVDT and geophone readouts are also shown as references. At low frequencies, the KAGRA sensor correction has the worst performance due to the noise amplification at 0.013 Hz resulting from the design of the sensor correction filter shown in subplot (a) in Fig. 8.18. In contrast, the noise amplification area of the \mathcal{H}_∞ falls between ~ 0.05 Hz and ~ 1 Hz, which correspond to the trough area formed by the seismometer and seismic noise shown in Fig. 8.22. The misplacement of the noise amplification in the KAGRA sensor correction filter caused an excessive low-frequency drift in the readout and the RMS value is $1.417 \mu\text{m}$ as shown in subplot (b) of Fig. 8.25. In comparison, the RMS value of the LVDT, the KAGRA super sensor, and the \mathcal{H}_∞ sensor correction readouts are $0.412 \mu\text{m}$, $0.552 \mu\text{m}$, and $0.351 \mu\text{m}$, respectively. The fact that the lower RMS of the \mathcal{H}_∞ sensor correction readout indicates that the \mathcal{H}_∞ sensor corrected LVDT is an overall better sensor, at least in the RMS sense.

At the microseism frequency, the suppression of the microseism coupling is similar for both sensor correction configuration. The microseismic peak is reduced by roughly one order of magnitude. Both of them offered a slightly better suppression compared to the KAGRA super sensor, as is indicated by the RMS plot. The RMS integrated down to 0.1 Hz reads $0.256 \mu\text{m}$, $0.0426 \mu\text{m}$ and $0.0482 \mu\text{m}$ for the LVDT readout, the KAGRA sensor correction readout, and the \mathcal{H}_∞ sensor correction readout, respectively. Although the \mathcal{H}_∞ sensor correction offers slightly lower suppression of seismic noise, the suppression is achieved with virtually no penalty whereas the KAGRA sensor correction has an excessive noise amplification in exchange. The \mathcal{H}_∞ sensor correction readout has an insignificant noise amplification compared to the LVDT noise at low frequency.

With the \mathcal{H}_∞ sensor correction filter, the corrected readout offers better suppression of the microseism while having superior noise performance at low frequency compared the KAGRA super sensor, which is intended to be used for “active seismic isolation” during O4. The corrected LVDT, being based on the LVDT readout, has poorer noise performance at high frequency compared to the KAGRA super sensor. However, this should not matter for the O4 target sensitivity, which is similar to sensitivity achieved during O3GK [24] where LVDTs were used for the control of the pre-isolators. Therefore, the \mathcal{H}_∞ sensor correction is undoubtedly already a better alternative for active seismic isolation for O4 compared to the current implementation of inertial damping at KAGRA.

The closed-loop results of the \mathcal{H}_∞ sensor correction is given in Sec. 8.3.5.

8.3.4 \mathcal{H}_∞ sensor fusion with sensor correction

With a sensor correction scheme implemented, it is worth re-introducing the sensor fusion and discuss how the geophones can improve the sensing performance. In this section, the sensor fusion between the sensor corrected LVDT in Sec. 8.3.3 and the geophone

is discussed. The \mathcal{H}_∞ sensor correction configuration in Sec. 8.3.3 has an \mathcal{H}_∞ norm of $\gamma = 3.56$, meaning that the estimated sensor correction noise is 3.56 times the noise of the LVDT. At frequencies higher than the microseism, the estimation is incorrect since the modified noise model has a higher amplitude compared to the real seismic noise. At those frequencies, the noise of the corrected LVDT is dominated by the the LVDT noise, instead of the sensor correction noise, as shown in Fig. 8.25. But, the sensor correction noise has just a slight offset of from the LVDT noise. Therefore, for simplicity, it is used as an conservative approximation of the sensing noise of the corrected LVDT.




Figure 8.26: SRM pre-isolator longitudinal noise models and estimated super sensor noise. Blue solid: Sensor corrected LVDT (model). Orange dashed: Geophone (model). Green dash-dot: \mathcal{H}_∞ super sensor with sensor correction (estimation).

The noise models of the sensor corrected LVDT and the geophone are shown as blue solid curve and orange dashed curve, respectively, in Fig. 8.26. As can be seen, this sensor fusion configuration is almost identical to the example shown in Sec. 8.2.1 as the noise model of the sensor corrected LVDT takes the shape of the LVDT noise. The procedure for \mathcal{H}_∞ optimization for sensor fusion is therefore the same as in the example and the \mathcal{H}_∞ complementary filters are shown in Fig. 8.27. The \mathcal{H}_∞ norm in this case is $\gamma = 4.65$. As shown in the figure, the complementary filters look very similar to those in Fig. 8.5 except with a higher blending frequency of ~ 0.12 Hz. The estimated super sensor noise is shown as green dash-dot line in Fig. 8.26, which is expectedly close to the lower limit at all frequencies. Again, note that the noise estimation at frequencies higher than the microseism is overestimated as the sensor corrected LVDT has a sensing noise slightly lower than that shown in Fig. 8.26.

Fig. 8.28 shows the open-loop longitudinal displacement of the SRM pre-isolator table




Figure 8.27: \mathcal{H}_∞ complementary filters for the sensor fusion of the sensor corrected LVDT and the geophone of the SRM pre-isolator in the longitudinal direction. Blue solid: \mathcal{H}_∞ low-pass filter for the sensor corrected LVDT. Orange dashed: \mathcal{H}_∞ high-pass filter for the geophone.

using different sensing readouts. As can be seen from subplot (a), the low frequency noise performance of the \mathcal{H}_∞ super sensor with sensor correction (\mathcal{H}_∞ super sensor from hereafter) (lilac dash-dot-dot) is similar to that of the sensor corrected LVDT (red dotted). They are both close the LVDT noise (blue solid). Also, The \mathcal{H}_∞ super sensor offers a slightly better noise performance at the microseismic peak. This is also indicated by the lower integrated RMS between ~ 0.12 Hz and ~ 0.3 Hz. However, when the RMS value is integrated down to 0.1 Hz, the value becomes $0.0573 \mu\text{m}$ for the \mathcal{H}_∞ super sensor, which is slightly higher than that of the \mathcal{H}_∞ sensor correction ($0.0483 \mu\text{m}$). This is due to the noise amplification around the blending frequency of the complementary filters. At high frequencies, the \mathcal{H}_∞ super sensor noise is close to the geophone sensing noise, which might be desirable for a low-noise operation. The overall RMS value of the \mathcal{H}_∞ super sensor readout is $0.359 \mu\text{m}$, which is comparable to the RMS value of the \mathcal{H}_∞ sensor correction readout ($0.351 \mu\text{m}$). Both \mathcal{H}_∞ readouts have lower RMS values than those of the LVDT ($0.412 \mu\text{m}$) and the KAGRA super sensor ($0.552 \mu\text{m}$). In short, compared to the \mathcal{H}_∞ sensor correction, the \mathcal{H}_∞ super sensor provides a slightly better seismic noise attenuation at the microseism and has lower noise at high frequencies with a cost of slightly higher RMS value due to the noise amplification around 0.12 Hz.

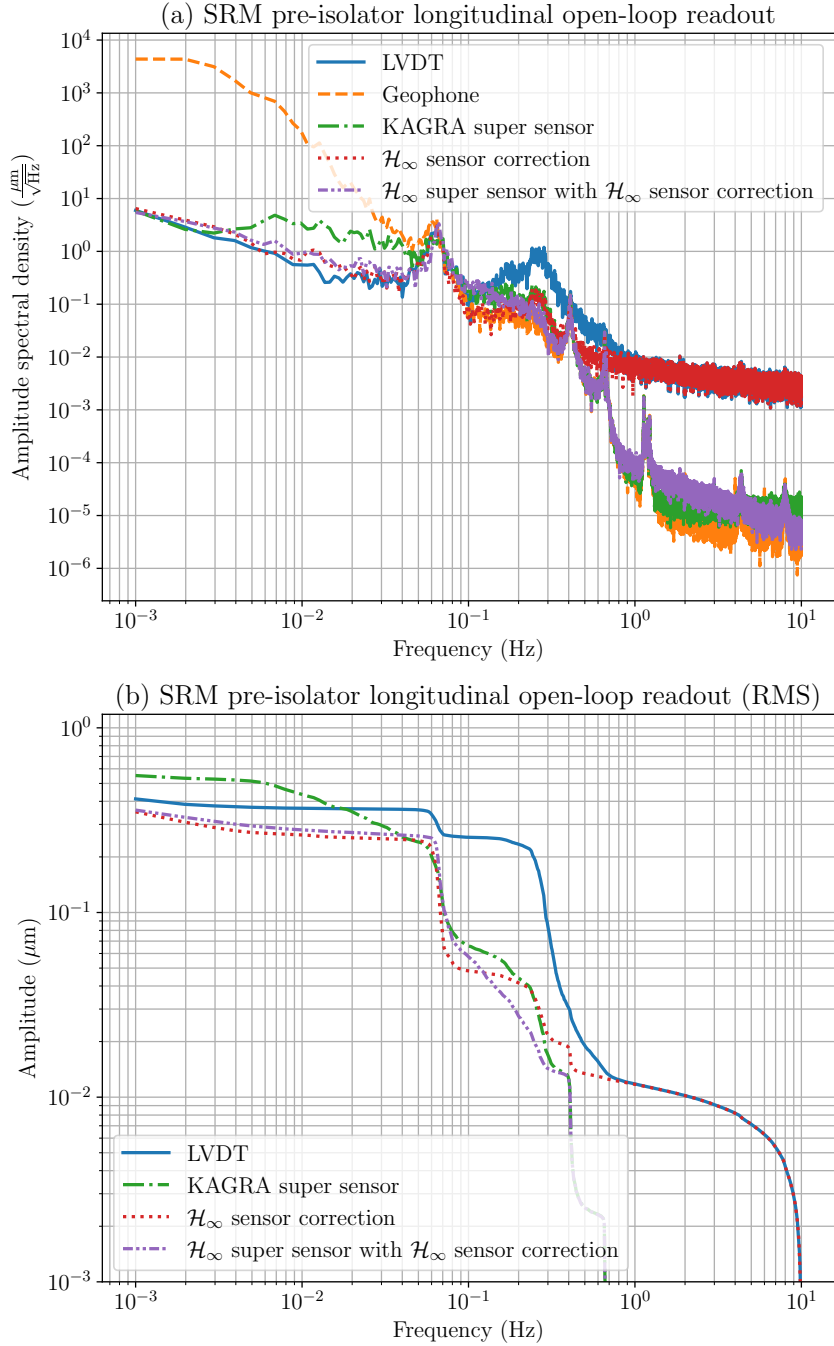


Figure 8.28: SRM pre-isolator longitudinal open-loop readout. (a) Amplitude spectral density. (b) RMS. Blue solid: LVDT. Orange dashed: Geophone. Green dash-dot: KAGRA super sensor. Red dotted: \mathcal{H}_∞ sensor correction. Lilac dash-dot-dot: \mathcal{H}_∞ sensor sensor with \mathcal{H}_∞ sensor correction.

8.3.5 \mathcal{H}_∞ sensors closed-loop performance and summary

As a verification, the \mathcal{H}_∞ sensor correction and complementary filters obtained in Sec. 8.3.3 and 8.3.4 were implemented to the SRM pre-isolator. The pre-isolator is controlled with the PID controller in Fig. 8.14 with 5 different sensors,

1. LVDT (seismic coupled),

2. KAGRA sensor corrected LVDT,
3. KAGRA super sensor,
4. \mathcal{H}_∞ sensor correction, and
5. \mathcal{H}_∞ super sensor with sensor correction.

The optical lever at the test mass stage is used as an out-of-loop sensor to monitor the longitudinal displacement of the SRM optics when the pre-isolator is controlled. Here, all measurements were done consecutively within a short period and they were all done under a similar seismic environment. This is to ensure that the closed-loop displacement levels are comparable. The seismic noise levels measured by the IXV sesimometer during the closed-loop measurements of the SRM are shown Fig. 8.29 as a reference. As can




Figure 8.29: Seismic noise levels during closed-loop measurements. Blue solid: Open loop. Orange dashed: Control with LVDT. Green dash-dot: Control with KAGRA sensor correction. Red dotted: KAGRA super sensor. Lilac dash-dot-dot: \mathcal{H}_∞ sensor correction. Brown dash-dot-dot-dot: \mathcal{H}_∞ sensor fusion with sensor correction.

be seen, the seismic noise levels between during each closed-loop measurements overlaps well.

The SRM longitudinal displacements measured by the optical lever are shown in Fig. 8.30. As shown in Fig. 8.30, the optical lever readout is unfortunately dominated by its self-noise. Nonetheless, some differences can be observed between closed-loop configuration using different sensors. Firstly, it is clear that the KAGRA sensor correction setup has introduced a significant drift to the optics at low frequency, which has increased the

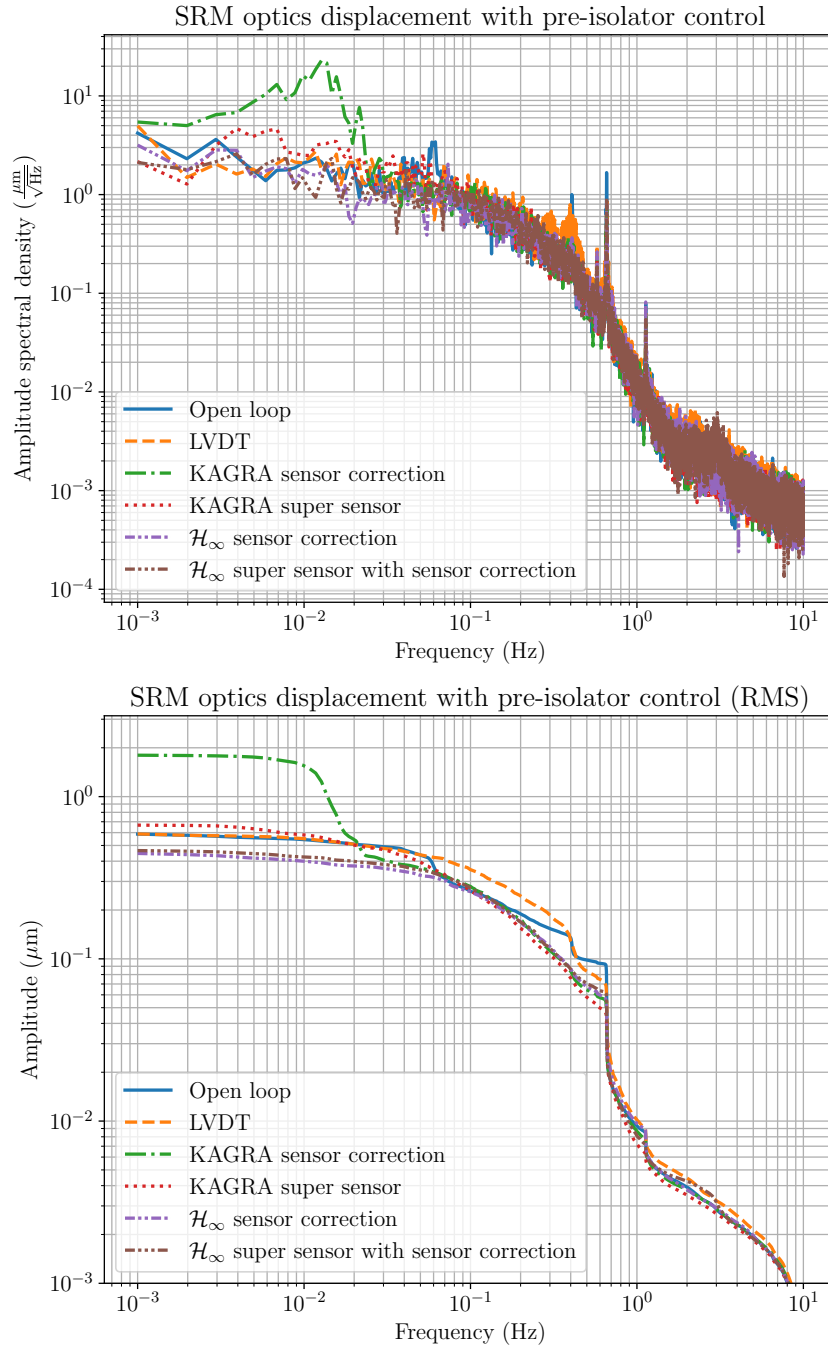


Figure 8.30: Longitudinal displacement of the SRM optics with pre-isolator control. (a) Amplitude spectral density. (b) RMS. Blue solid: Open-loop. Orange dashed: Control with LVDT. Green dash-dot: Control with KAGRA sensor correction. Red dotted: KAGRA super sensor. Lilac dash-dot-dot: \mathcal{H}_∞ sensor correction. Brown dash-dot-dot-dot: \mathcal{H}_∞ sensor fusion with sensor correction.

RMS value of the displacement to $1.8 \mu\text{m}$. This is significantly higher than those of other configurations and should not be considered as a configuration for feedback control.

At the first resonance frequency of the pre-isolator, 0.06 Hz, a peak is clearly observable from the displacement spectrum under and open-loop condition. From the RMS plot, the resonance peak contributes an RMS of around $0.1 \mu\text{m}$ (an increase from $\sim 0.2 \mu\text{m}$ to \sim

0.3 μm). And clearly, the peak is not observable in all closed-loop condition. This indicates that all sensor configurations are working properly under the closed-loop condition.

At 0.4 Hz, the second resonance peak, a difference between the LVDT control and others can be seen. The peak level is around $1 \mu\text{m}/\sqrt{\text{Hz}}$ under the open-loop condition, and is around $0.8 \mu\text{m}/\sqrt{\text{Hz}}$. Despite having an open-loop gain of ~ 10 as shown in Fig. 8.14, the suppression ratio with the closed-loop configuration using the LVDT is only $\sim 1/0.8 = 1.25$, which is way less than the open-loop gain of 10. This is due to the fact that the seismic noise level around the microseism is similar to that of the second resonance, as is also indicated in subplots (a) in Fig. 8.25 and 8.28. This means that the uncorrected LVDT measures the second resonance with a low signal-to-noise ratio and hence cannot be used to damp the peak effectively. Instead, the microseism is injected at that frequency where the open-loop gain is high, which explains the widened shape of the peak. In comparison, the peak is not observable under other closed-loop configurations. From the RMS plot, under the open-loop condition, the second peak contributes a step of RMS value of around $0.036 \mu\text{m}$ (increased from $\sim 0.111 \mu\text{m}$ to $\sim 0.147 \mu\text{m}$ from 0.42 Hz to 0.38 Hz). This step in the RMS plot is not obvious from all closed-loop configurations with sensor correction and sensor fusion. And, this indicates that the sensor correction and sensor fusion configurations can indeed be used to improve the damping performance of the suspensions.

As for the resonance peak at 0.67 Hz, the suppression ratio is low for all closed-loop configuration. This may be due to the fact that the control has a low open-loop gain and a low signal-to-noise ratio for all configurations. Also, there is a degenerating pendulum mode, that involves the longitudinal and pitch motion of the payload but not the pre-isolator, at this frequency. From the open-loop readouts in Fig. 8.25 and 8.28, the peak at 0.67 Hz appears to be lower than the peak at 0.4 Hz. Whereas in the optics open-loop readout in Fig. 8.30, this is reversed. Therefore, the observed peak at 0.67 Hz may not be entirely due to the motion of the pre-isolator, but rather, the motion of the aforementioned pendulum.

Pre-isolator control configuration	Optics displacement readout RMS (μm)
Open loop	0.587
LVDT	0.589
KAGRA sensor correction	1.8
KAGRA super sensor	0.667
\mathcal{H}_∞ sensor correction	0.446
\mathcal{H}_∞ super sensor with sensor correction	0.465

Table 8.1: RMS value of the SRM optics displacement readout under different pre-isolator control configuration.

The RMS values of the SRM optics displacement measured by the optical lever under different pre-isolator control configurations are summarized in Table 8.1. The values for

the open-loop, LVDT, KAGRA sensor correction, KAGRA super sensor, \mathcal{H}_∞ sensor correction, and \mathcal{H}_∞ super sensor configurations are $0.587\ \mu\text{m}$, $0.589\ \mu\text{m}$, $0.667\ \mu\text{m}$, $0.446\ \mu\text{m}$, and $0.465\ \mu\text{m}$, respectively. Again, note that the spectrum of the displacement readout of the optics is dominated by the self-noise of the optical lever. Therefore, the values in Table 8.1 do not represent the actual RMS values of the optics displacement. But, the differences can be used to evaluate the relative performance between different configurations. And, as can be seen, both \mathcal{H}_∞ configurations offer better closed-loop performance compared to the LVDT, and the original KAGRA sensor correction and sensor fusion configurations. Between the two \mathcal{H}_∞ configurations, the \mathcal{H}_∞ sensor correction configuration has a slightly lower closed-loop RMS value. This is because the use of geophones in the sensor fusion configuration introduces a slight addition of noise at low frequency, which eventually contributes to the subtle increase in RMS value. Nevertheless, the \mathcal{H}_∞ super sensor configuration offers ~ 3 orders of magnitude reduction in noise level at 10 Hz compared to the sensor correction configuration, as shown in Fig. 8.28 in Sec. 8.3.4.

Pre-isolator sensors	Sensor readout RMS (μm)
LVDT	0.412
KAGRA sensor correction	1.417
KAGRA super sensor	0.552
\mathcal{H}_∞ sensor correction	0.351
\mathcal{H}_∞ super sensor with sensor correction	0.359

Table 8.2: RMS values of the SRM pre-isolator displacement readouts under different sensing configuration (pre-isolator open loop).

Table 8.2 shows a summary of the RMS values of different sensing configurations of the pre-isolator under the open-loop condition. These values are reiterated from Sec. 8.3.3 and 8.3.4. As can be seen, the open-loop results agree with those measured in closed-loop condition in Table 8.1, with the original KAGRA configurations being worse than the LVDT and the \mathcal{H}_∞ sensors being superior. In summary, the \mathcal{H}_∞ sensor correction and sensor fusion configurations show improvements in noise performance compared to the LVDT, the original KAGRA sensor correction, and the original sensor fusion configurations. The \mathcal{H}_∞ sensor correction and the \mathcal{H}_∞ super sensor (with sensor correction) have similar noise performance in terms of RMS value, with the sensor fusion being slightly worse. However, the \mathcal{H}_∞ sensor fusion offers significantly lower noise at high frequency and should be considered when the pre-isolator control noise becomes significant to the detector sensitivity.

8.3.6 \mathcal{H}_∞ control

In this section, the \mathcal{H}_∞ control of the pre-isolator is discussed. Two types of controllers are presented here. The first problem is the active isolation control that seeks optimal

trade-off between seismic isolation and noise attenuation. And, the second problem is damping control similar to the example in Sec. 8.2.3. The former control system was unfortunately not realizable due to the impracticality and complexity of \mathcal{H}_∞ controllers. To compensate for this, the latter control serves as a simplified example of \mathcal{H}_∞ control but is not a realistic control problem in active isolation systems.

Active isolation control

For active isolation control at the pre-isolator stage, the goal is to attenuate the seismically induced motion without injecting too much control noise (Including actuation noise and sensor noise, although only sensor noise have been discussed so far.). The seismically induced motion is the corresponding disturbance under the feedback control framework shown in Fig. 5.5. Without control, the pre-isolator table displacement can be written as $X(s) = D(s) = P_g(s)X_g(s)$, where $D(s)$ is the disturbance, i.e. the seismically induced motion, $X_g(s)$ is the ground displacement, and $P_g(s)$ is the transfer function from ground displacement to the pre-isolator table displacement, i.e. the passive attenuation. Both transfer functions, from ground displacement or from actuation, refer to a similar same dynamics of the inverted pendulum. Therefore, assuming that the pre-isolator follows the ground motion perfectly below the main resonance frequency, the ground transfer function $P_g(s)$ can be approximated as the normalized actuation transfer function $P_g(s) = P(s)/|P(0)|$, where $P(s)$ is the transfer function from actuation to the pre-isolator displacement. Hence, the disturbance is $D(s) = P(s)X_g(s)/|P(0)|$, which is already obtainable based on transfer function models in previous sections ($P(s)$ from Sec. 8.3.1 and $X_g(s)$ from Sec. 8.3.3).

Suppose the \mathcal{H}_∞ sensor correction configuration is used (so the geophone can be used as an out-of-loop sensor to monitor the control performance), the corresponding sensing noise of the feedback control problem, like how it is done in Sec. 8.3.4, can then be approximated as the LVDT noise multiplied by the \mathcal{H}_∞ norm of the sensor correction configuration, which is $\gamma = 3.56$. In principle, the sensing noise should be re-measured with the sensor correction scheme implemented but the estimated noise is used here for simplicity. For the purpose of position control of the pre-isolator table, the disturbance is multiplied by a factor of $[(s + \omega_0)/(s + \epsilon)]^2$ with $\omega_0 = 2\pi \times 0.01$ rad/sec and $\epsilon = 2\pi \times 0.001$ rad/sec, which are arbitrarily chosen for illustration purpose. As opposed to first-order one described in Sec. 8.2.3, the second-order transfer function is chosen such that the relative order between the disturbance and the noise is at least one, which is required for \mathcal{H}_∞ synthesis to generate an integrator-like structure at low frequency. The disturbance (blue solid) and sensing noise model (orange dashed) are shown in Fig. 8.31.

With the disturbance, sensing noise, and actuation plant obtained, the weights and the transfer function models in the generalized plant for feedback control in Fig. 8.8 are all defined assuming the optimal weightings $W_D(s) = 1/\hat{N}(s)$ and $W_N(s) = 1/\hat{D}(s)$. The




Figure 8.31: SRM pre-isolator longitudinal noise models and displacement. Blue solid: Disturbance (model), i.e. The seismically induced displacement. Orange dashed: Noise (model), i.e. the sensing noise of the \mathcal{H}_∞ sensor corrected LVDT. Green dash-dot: The closed-loop displacement with \mathcal{H}_∞ controller.

feedback controller is optimized using \mathcal{H}_∞ synthesis according to Sec. 8.2.3. In Fig. 8.31, the estimated closed-loop displacement (green dash-dot) plotted. As can be seen, the main resonance peaks at 0.06 Hz and 0.067 Hz are suppressed to a level close to the sensing noise. The peaks at higher frequencies are not suppressed since they are below the noise level of the sensor, and in fact, they are amplified roughly by a factor of the \mathcal{H}_∞ norm, which is $\gamma = 1.9$ in this case. At high frequencies, the pre-isolator displacement rolls off towards the natural attenuation of the suspension. Again, this shows how \mathcal{H}_∞ control can suppress control noise such that the feedback control does not ruin the passive isolation performance at high frequency.

The frequency responses of the open-loop transfer function ((a) and (c)) and the \mathcal{H}_∞ controller ((b) and (d)) for SRM pre-isolator longitudinal direction is shown as blue solid lines in Fig. 8.32. The \mathcal{H}_∞ controller for the SRM pre-isolator was implemented to the KAGRA digital system using the Foton utility. To achieve this, the 46th-order \mathcal{H}_∞ controller was split into 3 transfer functions, which is equivalent to the \mathcal{H}_∞ controller when connected in series. This is necessary in practice since the Foton utility only accepts filters up to 20th order. The controller, as interpreted by Foton, is shown as orange dashed line in Fig. 8.32. As can be seen, the Foton utility unfortunately did not interpret the controller correctly. This resulted in an unexpected controller, which has unwanted peaks around the resonance frequencies.

Fig. 8.33 shows the Nyquist plot of the expected \mathcal{H}_∞ control (subplot (a)) and that of




Figure 8.32: SRM pre-isolator open-loop transfer function and \mathcal{H}_∞ controller. (a) and (c) Open-loop transfer function (magnitude and phase). (b) and (d) \mathcal{H}_∞ controller. Blue solid: Expected \mathcal{H}_∞ controller. Orange dash-dot: Foton interpreted \mathcal{H}_∞ controller.

the Foton interpreted control (subplot (b)) for the SRM pre-isolator. While it is unclear from the Bode plot in Fig. 8.32 that the \mathcal{H}_∞ control is stable, the Nyquist plot shown in subplot (a) in Fig. 8.33 clearly indicates the \mathcal{H}_∞ control of the SRM pre-isolator is stable as there is no encirclement around the -1 point. However, the same cannot be said for the Foton interpreted \mathcal{H}_∞ controller. In subplot (b), it is clear that the Foton interpreted \mathcal{H}_∞ controller resulted in an unstable system since there are multiple encirclement around the -1 point while the plant is open-loop stable. This means that the implementation of \mathcal{H}_∞ is unsuccessful and the system was tested to be closed-loop unstable in reality.

Nevertheless, there are certain implications that can be drawn from the the \mathcal{H}_∞ controller. First of all, the \mathcal{H}_∞ controller is a broadband controller, if not like a PID controller in Fig. 8.14. The V-shape magnitude response at low frequency is very similar to that created by the integral and derivative controller. This means that the PID controller is actually not that far from optimal. However, the \mathcal{H}_∞ controller clearly has lower loop-




Figure 8.33: Nyquist plot of the SRM pre-isolator longitudinal control. (a) Expected \mathcal{H}_∞ control. (b) Foton interpreted \mathcal{H}_∞ control.

gain at the main resonances, which indicates that the critical damping gain is, in fact, too high for this configuration. At the microseism, the open-loop gain is lower than 1, meaning that there is effectively no feedback action that actively attenuates the seismic noise. Therefore, active vibration isolation, at least for the SRM suspension, is more like “not injecting seismic noise” rather than suppressing seismic noise at the microseism. At high frequencies, the roll-off is 4th to 5th-order and the corner frequency, judging from the controller response, is around 0.3 Hz, which is roughly one decade lower than that in the PID controller. While the \mathcal{H}_∞ controller may not be implementable, a PID controller can be re-designed according to these features, either via manual shaping or \mathcal{H}_∞ loop-shaping with fixed controller structure, in the future if low-noise operation is desirable.

Example damping control

To exemplify the \mathcal{H}_∞ control, the yaw degree of freedom of the SRM pre-isolator is considered instead. The pre-isolator yaw only has one resonance and hence the plant is much simpler than that of the longitudinal direction. This hopefully would result in an \mathcal{H}_∞ controller that can be implemented. For the yaw degree of freedom, the disturbance in this direction cannot be easily characterized. Firstly, there is no natural yaw disturbance, i.e. shear on the ground. Secondly, the disturbance of the yaw degree of freedom comes from cross-coupling from other degrees of freedom, which cannot be easily estimated. Therefore, for simplicity, the disturbance here is assumed to be the actuation plant $P(s)$. And, during the actual test, the disturbance is injected via the actuators. As for the sensing noise, there is no sensor correction and the geophone is intended to be used as an out-of-loop sensor to measure the control performance. Therefore, the sensing noise is simply the LVDT noise.




Figure 8.34: SRM pre-isolator yaw noise models and displacement. Blue solid: Disturbance (model). Orange dashed: Noise (model). Green dash-dot: Closed-loop displacement with derivative control. Red dotted: Closed-loop displacement with \mathcal{H}_∞ control.

The disturbance (blue solid) and the noise (orange dashed) models in the SRM pre-isolator yaw degree of freedom are shown in Fig. 8.34. The transfer function models are used in \mathcal{H}_∞ synthesis to optimize the feedback controller and the \mathcal{H}_∞ norm is $\gamma = 1.19$. The \mathcal{H}_∞ controller and the corresponding open-loop transfer function are shown as orange dashed lines in Fig. 8.35. Note that the \mathcal{H}_∞ controller has a 4th-order low-pass post-filter at 10 Hz. As a comparison, the derivative controller that is optimized for critical damping is shown as blue solid lines. The derivative controller has a 4th-order low-pass filter optimized for a 60° phase margin.

The estimated closed-loop displacements for the derivative control and \mathcal{H}_∞ control are shown as green dash-dot and red dotted lines, respectively, in Fig. 8.34. As can be seen, both derivative control and \mathcal{H}_∞ provide adequate suppression at the resonance frequency, with the derivative control providing a slightly higher attenuation. However, the derivative control injects excessive noise at higher frequency compared to the \mathcal{H}_∞ control. On the other hand, the \mathcal{H}_∞ closed-loop displacement follows the lower limit at all frequencies as expected. This shows that the \mathcal{H}_∞ method can be used to design controllers to suppress resonance peaks with low-noise injection, assuming that the disturbance can be modeled.

The \mathcal{H}_∞ controller was implemented to the actual SRM suspension and it is tested by injected the corresponding disturbance using the actuators. The derivative controller was also tested as a comparison. Fig. 8.36 shows the open-loop and closed-loop displacements of the SRM pre-isolator with the disturbance injection. Note that the closed-loop displacements were measured by the geophone and the feedback controller uses readout

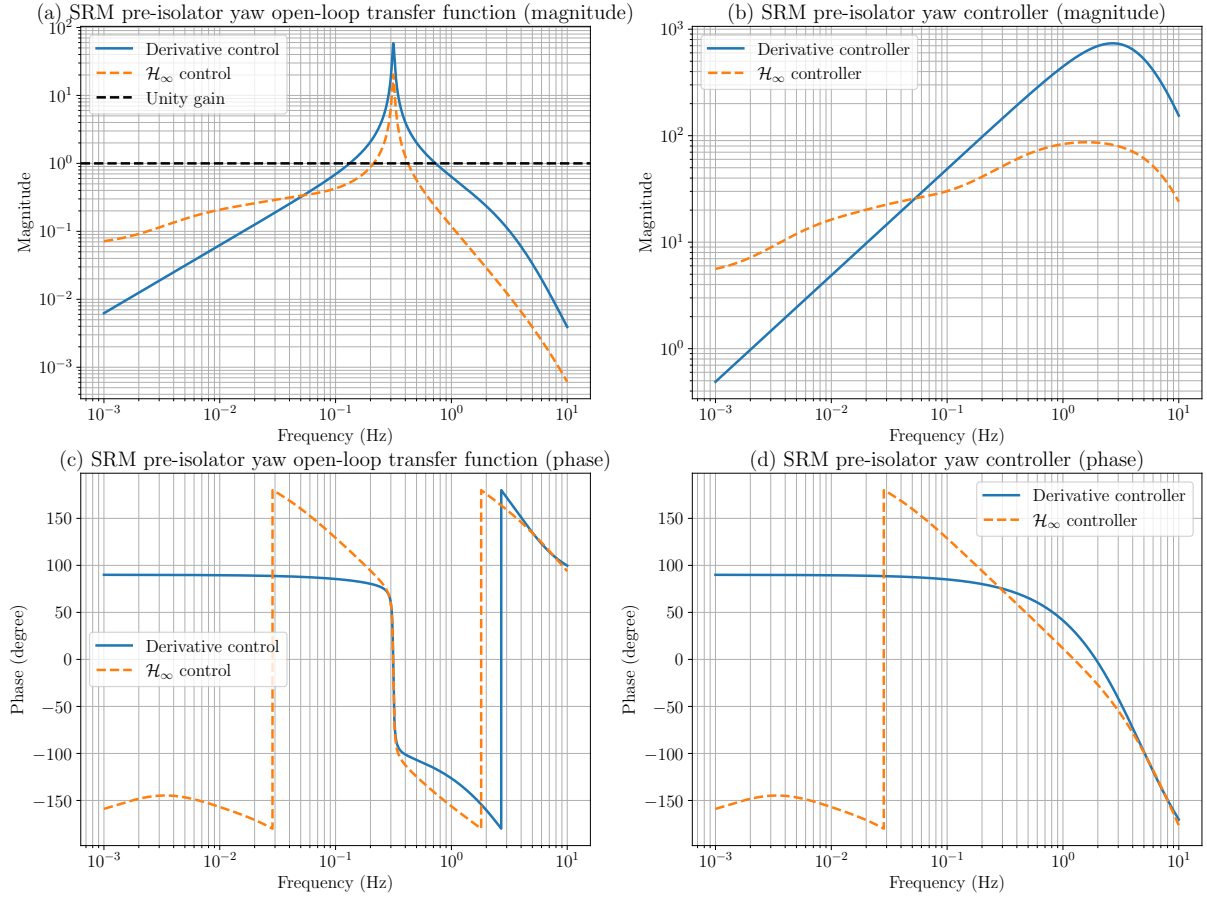


Figure 8.35: SRM pre-isolator yaw open-loop transfer function and controller. (a) and (c): Open-loop transfer function (magnitude and phase). (b) and (d): Controller (magnitude and phase). Blue solid: Derivative control. Orange dashed: \mathcal{H}_∞ control.

from the LVDT. Also, the geophone readout is not useful at low frequency, especially below ~ 0.1 Hz, due to the geophone noise. As shown in the figure, the results agrees with the conclusion in the simulated cases in Fig. 8.34. Both controllers suppress the yaw resonances well but the derivative controller resulted in a noise pump at higher frequencies. Again, this shows how \mathcal{H}_∞ method can be used to optimize controllers that attenuates control noise properly.




Figure 8.36: SRM pre-isolator yaw displacement. Blue solid: Open-loop displacement measured by LVDT. Orange dashed: Open-loop displacement measured by geophone. Green dashed: PID closed-loop displacement measured by geophone. Red dotted: \mathcal{H}_∞ closed-loop displacement measured by geophone.

Chapter 9

Concluding Remarks

9.1 Discussions

9.1.1 On \mathcal{H}_∞ methods

Undoubtedly, the \mathcal{H}_∞ method has provided a new way of optimizing control filters for active isolation systems in gravitational-wave detectors. And, the method has been verified experimentally with the KAGRA SRM suspension as shown in Sec. 8.3.3, 8.3.4, and 8.3.6. While the control schemes, such as sensor fusion, sensor correction, and feedback control, are not exclusive to KAGRA detector, the \mathcal{H}_∞ method can eventually be useful for other gravitational-wave detectors including LIGO, Virgo, and in the future, LIGO-India and the Einstein Telescope. However, in reality, there are some practical problems regarding \mathcal{H}_∞ methods that need to be addressed.

As mentioned in Sec. 8.1.2, the central controller of the \mathcal{H}_∞ problem has the same states of the generalized plant. With the weighting functions, often times, the \mathcal{H}_∞ synthesis gives high-order controllers, especially in the case of the feedback control setup discussed in Sec. 8.2.3. High-order controllers may be fine in simulation but could face problems related to numerical accuracy in reality. This is because the coefficients of the polynomial could span a huge dynamic range. For example, the lowest polynomial coefficient in the \mathcal{H}_∞ controller for active isolation in Sec. 8.3.6 is 1.549×10^{-10} whereas the highest coefficient is 3.018×10^{37} . Also, the value 1.549×10^{-10} is possibly an outlier/residual as the second lowest coefficient in the same polynomial is 1.104×10^8 , which is 18 order of magnitude higher. In fact, the simulation of time response with \mathcal{H}_∞ position control in Sec. 8.2.3 was unsuccessful and resulted in a “Badly conditioned filter coefficients (numerator)” error using `control.forced_response()` from the Python Control package. The solution is to remove such an outlier coefficient. And, it is suspected that the implementation failure of the \mathcal{H}_∞ controller in Sec. 8.3.6 might have be due to this issue. To resolve this problem in the future, one might need to consider model reduction methods for simplifying the high-order controllers or weights using lower-order

approximations [63].

In Sec. 8.2, the complementary filter shaping problem, sensor correction problem, and feedback control problem are treated as separate problems. And, indeed, the \mathcal{H}_∞ method solves the problems well individually. But in reality, they are really small problems of the active isolation problem, which ultimately seeks the minimization of the displacement of the suspended optics. The \mathcal{H}_∞ method solves each subproblem and the solution is assigned with an \mathcal{H}_∞ norm. And, by the way the problem is set up, the \mathcal{H}_∞ represents how much the error signal is offsetted from a frequency-dependent lower limit. By treating the subproblems as individual problems, the offsets become accumulative, and the sensor correction noise in the sensor fusion configuration in Sec. 8.3.4 is a good example of this phenomenon. The \mathcal{H}_∞ method is generally capable of synthesizing regulators for multiple-input and multiple-output systems. This means that all control filters in an active isolation system could be synthesized in one design process in principle, seeking the minimization of the motion suspended optics and other degrees of freedom of the suspension. The hope is to look for an overall lower \mathcal{H}_∞ norm compared to the accumulative one. However, this calls for a much sophisticated modeling of the suspension, including the characterization of transfer functions between different degrees of freedom, and this remains to be a future work.

9.1.2 On inertial sensor pre-filtering

The inertial sensors are intrinsically AC coupled and they have poor noise performance at low frequencies. In practice, the inertial sensors need to be pre-filtered with high-pass filters or else the low-frequency error in the inertial signals could accumulate and eventually saturate the process variables in the digital system, making them unusable. However, the high-pass pre-filters add non-unity response to the inertial sensors and this distorts the displacement measurements. As a result, this causes inter-calibration mismatch between the relative sensors and the inertial sensors, in the context of sensor fusion. Even if the complementary filters meet the complementary condition, the super sensor would have a non-unity response around the blending frequency, which could eventually deteriorate the control performance or even cause problems such as instability. This is especially important when the cutoff frequency of the high-pass filter is close to the blending frequency. Hence, one simple way of solving the problem is to lower the cutoff frequency of the pre-filter. However, this could be challenging due to the poor noise performance of the geophones.

In the case of KAGRA, the effect of pre-filtering is not negligible. Fig. 9.1 subplot (a) shows the ratio between the \mathcal{H}_∞ super sensor (the one in Sec. 8.3.4) readout and the LVDT readout during a transfer function measurement of the SRM pre-isolator longitudinal degree of freedom. Two \mathcal{H}_∞ super sensors are shown here. Both \mathcal{H}_∞ super sensors use




Figure 9.1: (a): Ratio between the \mathcal{H}_∞ super sensor readout and the LVDT readout. Blue solid: Pre-filter uncompensated. Orange dashed: Pre-filter compensated. (b): Super sensor response with prefilter uncompensated.

a pre-filtered geophone signal. The first one uses the \mathcal{H}_∞ complementary filters directly without compensating the effect from pre-filtering (blue solid). And the other uses post-processed complementary filters that takes into account the effect of pre-filtering (Orange dashed).

In principle, the LVDT readout and the \mathcal{H}_∞ super sensor readouts should measure the same signal and hence their ratio should be unity. However, as shown in subplot (a) of Fig. 9.1, this is not the case. First of all, the difference at high frequency above 10 Hz is due to the difference between the noise level of the LVDT and the super sensor, which is expected. At ~ 0.35 Hz, there is antiresonance in the transfer function of the pre-isolator so the signals are noise dominated, which explains the huge difference again. The same can be said for the peaks and notches around 0.06 Hz and 0.6 Hz. However, with the geophone pre-filter not compensated, there is a notch around 0.1 Hz that is not accounted by any physical response. Subplot (b) in Fig. 9.1 shows the sum of the \mathcal{H}_∞ LVDT filter and the \mathcal{H}_∞ geophone filter with the pre-filter, which corresponds to the super sensor response. As can be seen, the super sensor response is not unity as the complementary condition is not met. Also, there exists a notch around 0.1 Hz and this well explains the differences between the \mathcal{H}_∞ super sensor and the LVDT shown as blue solid curve in subplot (a).

The non-unity response of the super sensor is an uncertainty in the control system and should be eliminated if possible. To achieve this, three post-processing methods for the complementary filters are proposed here to compensate for the pre-filter. Two of them are not exclusive to the \mathcal{H}_∞ method and can be used for any complementary filters. Each method have their own pros and cons and hence they are simply suggested here without a definitive choice. It should be note that the deteriorating effect from the pre-filter is

not an \mathcal{H}_∞ problem but is a general problem for any design of complementary filters. In fact, the KAGRA complementary filters currently implemented in the suspensions do not take into account the effect of the pre-filter.




Figure 9.2: \mathcal{H}_∞ complementary filters in Sec. 8.3.4. (a): Blue solid: \mathcal{H}_∞ low-pass filter. Orange dashed: \mathcal{H}_∞ high-pass filter. Green dash-dot: \mathcal{H}_∞ low-pass filter redefined from the complementary condition. Red dotted: \mathcal{H}_∞ with high-pass pre-filter. (b): Blue solid: \mathcal{H}_∞ low-pass filter. Orange dashed: \mathcal{H}_∞ high-pass filter. Green dash-dot: \mathcal{H}_∞ low-pass filter normalized using super sensor response. Red dotted: \mathcal{H}_∞ high-pass filter normalized using super sensor response.

Consider a pair of complementary filters $H_1(s)$ and $H_2(s) \equiv 1 - H_1(s)$ for a sensor fusion of two sensors, sensor 1 and sensor 2, respectively. And, suppose a pre-filter $H_{\text{pre}}(s)$ is required for attenuating the excessive noise from the sensor 2, this redefines the complementary filter $H_2^*(s) \equiv H_2(s)H_{\text{pre}}(s)$ for sensor 2. The super sensor then has a frequency-dependent response of

$$R_{\text{super}}(s) = H_1(s) + H_2^*(s) = H_1(s) + H_2(s)H_{\text{pre}}(s) \neq 1. \quad (9.1)$$

One simple way to restore the unity response of the super sensor is to simply redefine the complementary filter for the sensor 1 using the complementary condition, i.e.

$$H_1^*(s) \equiv 1 - H_2^*(s) = 1 - H_2(s)H_{\text{pre}}(s). \quad (9.2)$$

In this case, the complementary filters actually implemented to the control system are $H_1^*(s)$ and $H_2(s)$. But effectively, the complementary filters are $H_1^*(s)$ and $H_2^*(s)$, which are complementary, when the pre-filter $H_{\text{pre}}(s)$ is also implemented. This method is used in the implementation of the \mathcal{H}_∞ complementary filters in Sec. 8.3.2 and 8.3.4 due to its simplicity and it does not ruin the noise performance around the blending frequency, which the next method would. Subplot (a) Fig. 9.2 compares the original \mathcal{H}_∞ complementary

filters (blue solid and orange dashed) and the redefined ones $H_1^*(s)$ and $H_2^*(s)$ (green dash-dot and red dotted). As can be seen, the magnitude responses of the filters are similar except at low and high frequency, which is due to the pre-filter and redefinition. In Fig. 9.1, the ratio between \mathcal{H}_∞ super sensor (with the low-pass filter redefined using Eqn. (9.2)) and the LVDT. The notch around 0.1 Hz corresponding to a non-unity super sensor response is clearly eliminated.

The second method to restore a unity super sensor response is described in Refs. [28, 61]. But, the significance is not well discussed in the references and the method is simply provided as a procedure. Here, the method is given as a re-elaboration and the significance is also given. The introduction of the pre-filter $\mathcal{H}_{\text{pre}}(s)$ resulted in a non-unity super sensor response as described in Eqn. (9.1). It is obvious that when the super sensor signal is filtered by $1/R_{\text{super}}(s)$, i.e. the inverse dynamics of the original super sensor signal with a pre-filtered signal, the unity response is restored. It is equivalent to modifying the originally designed complementary filters as

$$H_1^*(s) = \frac{H_1(s)}{R_{\text{super}}(s)} = \frac{H_1(s)}{H_1(s) + H_2(s)H_{\text{pre}}(s)}, \quad (9.3)$$

and

$$H_2^*(s) = \frac{H_2(s)}{R_{\text{super}}(s)} = \frac{H_2(s)}{H_1(s) + H_2(s)H_{\text{pre}}(s)}, \quad (9.4)$$

where $H_1(s)$ and $H_2(s)$ are the originally designed complementary filters, $R_{\text{super}}(s)$ is the super sensor response in Eqn. (9.1), and $H_1^*(s)$ and $H_2^*(s)$ are the normalized complementary filters to be implemented. It is obvious that

$$H_1^*(s) + H_2^*(s)H_{\text{pre}}(s) = 1, \quad (9.5)$$

is complementary, i.e. the super sensor response is unity when sensor 2 is pre-filtered with $H_{\text{pre}}(s)$ and the complementary filters installed are $H_1^*(s)$ and $H_2^*(s)$. The new complementary filters (with the pre filter) normalized using this method (green dashed and red dotted) are shown in subplot (b) in Fig. 9.2. As can be seen, the magnitude responses of the the low-pass filter retains as the original one, unlike the first method. However, as a trade-off, the magnitudes of both filters around the blending frequency increase, which can be seen as anti-notching the original super sensor response. This could result in excess noise amplification around the blending frequency, which could ruin the RMS performance of the super sensor. In such case, the other two alternatives should be used.

The third method is only applicable to the \mathcal{H}_∞ method and is perhaps a more elegant way of treating the pre-filter effect. However, this does not always work as the \mathcal{H}_∞ synthesis may yield no solution and hence is only given here as an alternative to the first two methods. The idea is to include the pre-filter information in the generalized plant in

Fig. 8.3 by modifying the noise model. The noise model $\hat{N}_2(s)$ of the sensor that needs to be pre-filtered is replaced by $\hat{N}_2(s)/\hat{H}_{\text{pre}}(s)$, where $\hat{H}_{\text{pre}}(s)$ is a stable transfer function that has a frequency response representing that of the pre-filter $H_{\text{pre}}(s)$, i.e. having components like $1/s$ replaced by $1/(s + \epsilon)$ as discussed in Sec. 8.2.3. This way, the noise model is amplified by the inverse of the pre-filter and the \mathcal{H}_∞ optimization will naturally generate a complementary filter containing the pre-filter to counteract that amplification. For the complementary filter $H_2(s)$, it needs to be multiplied by the inverse of the pre-filter before implementation, so $H_2^*(s) = H_2(s)/\hat{H}_{\text{pre}}(s)$ is the one that gets implemented. In such a way, the effective complementary filter, $H_2^*(s)H_{\text{pre}}(s)$, becomes the \mathcal{H}_∞ filter, as the pre-filter is applied. For the other complementary filter $H_1(s)$, i.e. the one for the sensor without pre-filter, it can be implemented as is since it is complementary to the effective complementary filter $H_2^*(s)H_{\text{pre}}(s)$. Unfortunately, a solution for the SRM pre-isolator sensor fusion does not exist and hence the method is not demonstrated here.

9.1.3 On the implementation of inertial sensors

The implementation of inertial sensors for sensor fusion is a daunting task yet not so rewarding in terms of active vibration isolation. There are many things that need to be done properly in order for the inertial sensors to work. This includes obtaining the calibration filters for the geophones, aligning the geophone sensing basis to the LVDT sensing basis, inter-calibrating the geophones and the LVDTs, and hoping that the geophones do not degrade so the parameters in the calibration filters remain unchanged (The L-4C geophones are known to degrade [48]). Error in any of the tasks could accumulate and lead to a super sensor with non-unity response, which is an uncertainty in the control systems that could lead to control limitations or instability that are hard to trace. In exchange, as shown in Sec. 8.3.4, only high-frequency noise performance is improved significantly but with the cost of a slight increase in RMS, compared the sensor correction method. With the current sensitivity target for KAGRA (similar to that achieved in O3GK), the high-frequency performance at the pre-isolator may not be required because the sensitivity is not limited by the control noise at the pre-isolator. While using the inertial sensors do not lead to massive improvements in active isolation but could lead to potential problems, it is not recommended to use the inertial sensors for inertial damping under the tight schedule of KAGRA.

In comparison, sensor correction is inarguably a simpler and effective way of attenuating the microseism in the relative readout, if implemented correctly as in Sec. 8.3.3. The \mathcal{H}_∞ sensor correction readout has the lowest noise RMS value compared to all sensors presented and the sensor fusion with sensor correction scheme resulted in a slightly higher RMS. Since many suspensions can share the same seismometer for sensor correction, the same sensor correction filter can be replicated for all suspensions at close proximity, as-

suming that the noise levels of the LVDTs in the suspensions are similar. The only task that needs to be done individually for each suspension is the inter-calibration between the LVDT readout and the seismometer readout, which is easy to achieve. Even if the task is not done, the sole implementation of the sensor correction filter could lead to sub-optimal seismic noise suppression, which is also undoubtedly useful. Therefore, the sensor correction scheme is a recommended way of achieving inertial damping, rather than using sensor fusion.

9.1.4 On the results

The results of the \mathcal{H}_∞ method shown in Sec. 8.3 has undoubtedly show improvements over the original KAGRA control filters. But, there is a caveat. The measurements in the results are not direct confirmations of the improvement of the interferometer stability, which is the quality that this research ultimately aims to achieve. This is because the measurements are local readouts of only one suspension whereas the interferometer stability is determined by the relative displacement between the suspended optics. In principle, the improvements achieved by this research can be confirmed by measuring the control signals of the interferometer, such as the differential arm length signal (DARM), when the interferometer is locked. However, the interferometer of KAGRA, including the signal-recycling cavity that involves the SRM suspension, was not available during the the period that this research was conducted. Therefore, the results in this thesis only serve as an indirect confirmation of stability improvements of the interferometer. The results must be further confirmed in the future.

9.2 Future prospects

Besides some possible improvements regarding the \mathcal{H}_∞ method that are already discussed in Sec. 9.1, there are more future works that can be done for active isolation systems in gravitational-wave detectors in general. Here, a few possible future works are given.

The RMS value of a signal is often used as a figure of merit for quantifying active isolation performance. However, the RMS value is typically dominated by low frequency signals. Therefore, high frequency signals that have much lower order of magnitude are typically not well informed. For this reason, instead of the \mathcal{H}_2 method that minimizes the RMS value, the \mathcal{H}_∞ method is chosen to be used due to its capability of specifying frequency-dependent upper bounds requirements. With that said, \mathcal{H}_2 and \mathcal{H}_∞ methods are not mutually exclusive and there exists methods can be used to solve a mixed- $\mathcal{H}_2/\mathcal{H}_\infty$ problem [77]. While the RMS value of signals, such as the test mass displacement, is often capable of indicating the stability of the interferometer, it is worth developing a mixed- $\mathcal{H}_2/\mathcal{H}_\infty$ method for active isolation systems in gravitational-wave detectors. The

aim is to achieve a minimization of the RMS value while achieving good noise attenuation performance at high frequency.

As shown in Fig. 2.5, the seismic noise in at KAGRA can vary significantly within a year. This would mean that the ideal controllers and control filters for active isolation may vary accordingly. The \mathcal{H}_∞ methods provided in Sec. 8.1 can only lead to static controllers, which are only optimal in one particular situation. At other times, the controllers are rendered sub-optimal. This calls for an adaptive control approach in seismic isolation.

The LIGO detectors suffer from lock-loss due to distant earthquakes, which generate seismic waves between 50 mHz - 60 mHz [78]. When the seismic waves hits the detectors, the seismic waves are measured by the local seismometers, which are used for the sensor correction of the relative sensors at LIGO. The sensor correction filters, like the ones in Fig. 7.7 and 8.24, have noise simplification around the frequency of the seismic wave. This means that the motions of the seismic waves are amplified and transmitted into the motion of the suspensions via the sensor correction scheme. As a result, when this happens, it results in a lock-loss of the interferometer, which ultimately leads to a reduced in duty cycle. As a solution, a control scheme called “earthquake mode” is developed to avoid such situation. In such scheme, two sensor correction configurations are prepared, with one optimized for nominal condition focusing on the suppression of the secondary microseism, and the other optimized for the earthquake mode which has a pass band at lower frequency allowing the sensor correction for the increase seismic noise at 50 mHz. The latter is engaged when a seismic disruption is foreseen that could cause the seismic noise to increase to a certain threshold. While this control scheme has successfully increase the chances of keeping the interferometer lock during earthquakes, there is a drawback. That is, the amplification of the seismometer noise at even lower frequencies and the amount of amplification is the regardless of the magnitude of the earthquakes. This is not ideal for smaller earthquakes as the seismometer noise amplification could theoretically be reduced. With the \mathcal{H}_∞ method, it is possible to optimize the sensor correction according to the detected seismic wave in real-time. In such case, the sensor correction filter is optimal at all times. Therefore, the \mathcal{H}_∞ method could complement the earthquake mode approach and combining the two would be an interesting research topic .

The static architecture of the control system has undoubtedly limited the control performance due to the lack of adaptivity. Also, the classical filter-based control system restricts the conversion from sensor signals to actuation signals to a certain way, which ultimately limits the control performance of the active isolation. All of this is due to the under-developed real-time control software used in current gravitational-wave detectors, which limits control systems to be implemented in a certain way. The lack of software development may be attributed to the fact that they have been working stably for gravitational-wave detectors. However, this is inevitably a burden for developing advanced control methods for active isolation and interferometer control, which could

potentially improve the interferometer stability. This may not be currently important for current gravitational-wave detectors, but is important for future detectors such as the Einstein Telescope and LIGO-india, which are detectors with higher sensitivities and hence much tighter noise requirements. The earthquake mode may have provided an improvement in interferometer stability but it is not perfect and there is always a risk of lock-loss during the switching of operation modes due to the complicated procedures. In contrast, modern control methods may not have this problem. For example, in a model predictive control scheme [79], the control signals are optimized on-the-fly from the sensor signals for a receding time horizon. If a seismic wave is foreseen, the model predictive controller would autonomously take the seismic wave into account and optimize the control signals according to the objective function. This provides the adaptivity that is required for active isolation in gravitational-wave detectors and could, in principle, completely replace the earthquake mode control scheme. And, because the model predictive control scheme optimizes the control signals directly without subjected to limitations of control filters, it could potentially achieve better control performance compared to a filter-based approach. However, such advanced control methods cannot be implemented due to the limitation of the software used in current gravitational-wave detectors. Therefore, the control software used in gravitational-wave detectors need to be continually developed to accommodate advanced control strategies in the future.

Bibliography

- [1] Albert Einstein. Zur allgemeinen relativitätstheorie. *Preussische Akademie der Wissenschaften, Sitzungsberichte*, 1915 (part 2):778–786, 799–801, 1915.
- [2] John Wheeler and Kenneth Ford. Geons, black holes and quantum foam: A life in physics. *American Journal of Physics*, 68, 06 2000.
- [3] Albert Einstein. Näherungsweise integration der feldgleichungen der gravitation. *Preussische Akademie der Wissenschaften, Sitzungsberichte*, 1916 (part 1):688–696, 1916.
- [4] Albert Einstein. Gravitationswellen. *Preussische Akademie der Wissenschaften, Sitzungsberichte*, 1918 (part 1):154–167, 1918.
- [5] Peter R Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. WORLD SCIENTIFIC, 2nd edition, 2017.
- [6] T.G.F. Li. *Extracting Physics from Gravitational Waves: Testing the Strong-field Dynamics of General Relativity and Inferring the Large-scale Structure of the Universe*. Springer Theses. Springer International Publishing, 2015.
- [7] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press, 10 2007.
- [8] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal Letters*, 195:L51–L53, January 1975.
- [9] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.
- [10] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.
- [11] B. P. Abbott et al. Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *Astrophys. J. Lett.*, 848(2):L13, 2017.

- [12] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. 11 2021.
- [13] T. Akutsu et al. Overview of KAGRA : KAGRA science. 8 2020.
- [14] T Akutsu et al. Overview of KAGRA : Detector design and construction history. *Progress of Theoretical and Experimental Physics*, 2021, 08 2020.
- [15] M Punturo et al. The Einstein Telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, sep 2010.
- [16] K L Dooley et al. GEO 600 and the GEO-HF upgrade program: successes and challenges. *Classical and Quantum Gravity*, 33(7):075009, mar 2016.
- [17] H. Abe et al. Performance of the KAGRA detector during the first joint observation with GEO 600 (O3GK). 3 2022.
- [18] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, s3-34(203):333–345, 1887.
- [19] Eric D. Black and Ryan N. Gutenkunst. An introduction to signal extraction in interferometric gravitational wave detectors. *American Journal of Physics*, 71(4):365–378, 2003.
- [20] LIGO. LIGO’s Laser. <https://www.ligo.caltech.edu/page/laser>.
- [21] William J. Startin, Mark A. Beilby, and Peter R. Saulson. Mechanical quality factors of fused silica resonators. *Review of Scientific Instruments*, 69(10):3681–3689, 1998.
- [22] T. Akutsu et al. First cryogenic test operation of underground km-scale gravitational-wave observatory KAGRA. *Class. Quant. Grav.*, 36(16):165008, 2019.
- [23] Geological Survey (U.S.). *Observations and Modeling of Seismic Background Noise*, Peterson, Jon.
- [24] Shinji Miyoki. Current status of KAGRA. In Heather K. Marshall, Jason Spyromilio, and Tomonori Usuda, editors, *Ground-based and Airborne Telescopes VIII*, volume 11445, pages 192 – 204. International Society for Optics and Photonics, SPIE, 2020.
- [25] G. Cella, V. Sannibale, R. DeSalvo, S. Márka, and A. Takamori. Monolithic geometric anti-spring blades. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 540(2):502–519, 2005.

- [26] F Matichard et al. Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, aug 2015.
- [27] Lucia Trozzo. *Low Frequency Optimization and Performance of Advanced Virgo Seismic Isolation System*. PhD thesis, Siena U., 2018.
- [28] J.V. van Heijningen. *Turn up the bass!: Low-frequency performance improvement of seismic attenuation systems and vibration sensors for next generation gravitational wave detectors*. PhD thesis, Vrije Universiteit Amsterdam, 2018.
- [29] T. Sekiguchi. *Study of Low Frequency Vibration Isolation System for Large Scale Gravitational Wave Detectors*. PhD thesis, Tokyo U., 2016.
- [30] Koki Okutomi. Development of 13.5-meter-tall vibration isolation system for the main mirrors in kagra. 2019.
- [31] Yoshinori Fujii. *Fast localization of coalescing binaries with gravitational wave detectors and low frequency vibration isolation for KAGRA*. PhD thesis, The University of Tokyo, 2019.
- [32] T Akutsu et al. Vibration isolation systems for the beam splitter and signal recycling mirrors of the KAGRA gravitational wave detector. *Classical and Quantum Gravity*, 38(6):065011, mar 2021.
- [33] Y Akiyama et al. Vibration isolation system with a compact damping system for power recycling mirrors of KAGRA. *Classical and Quantum Gravity*, 36, 05 2019.
- [34] G. Losurdo et al. An inverted pendulum preisolator stage for the VIRGO suspension system. *Rev. Sci. Instrum.*, 70(5):2507–2515, 1999.
- [35] Tomotada Akutsu et al. Compact integrated optical sensors and electromagnetic actuators for vibration isolation systems in the gravitational-wave detector KAGRA. *Rev. Sci. Instrum.*, 91(11):115001, 2020.
- [36] Joris van Heijningen, A. Bertolini, and Johannes van den Brand. A novel interferometrically read out inertial sensor for future gravitational wave detectors. pages 1–5, 03 2018.
- [37] L. Carbone et al. Sensors and Actuators for the Advanced LIGO Mirror Suspensions. *Class. Quant. Grav.*, 29:115005, 2012.
- [38] Takafumi Ushiba et al. Cryogenic suspension design for a kilometer-scale gravitational-wave detector. *Classical and Quantum Gravity*, 38(8):085013, mar 2021.

- [39] Alexander Franzen. ComponentLibrary. <http://www.gwoptics.org/ComponentLibrary/>.
- [40] A. (Anthony) Gerrard. *Introduction to matrix methods in optics*. Wiley series in pure and applied optics. Wiley, London ; New York, 1975.
- [41] R. Bork, M. Aronsson, and A. Ivanov. aligo cds real-time code generator application developer’s guide. LIGO document LIGO-T080135: <https://dcc.ligo.org/public/0001/T080135/005/LIGO-T080135-v5.pdf>, April 2013.
- [42] Simulation and model-based design. <https://www.mathworks.com/products/simulink.html>, 2020.
- [43] Mark Barton. Bs suspension diagonalization matrices. KAGRA document JGW-T1707205: <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/DocDB/ShowDocument?docid=7205>, 2018.
- [44] Mark Barton. Sr suspension diagonalization matrices. KAGRA document JGW-T1807663: <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/DocDB/ShowDocument?docid=7663>, 2018.
- [45] Enzo Tapia. IP LVDT calibration. KAGRA Logbook: <https://klog.icrr.u-tokyo.ac.jp/osl/?r=5551>, 2018.
- [46] Sercel. <https://www.sercel.com/Pages/default.aspx>.
- [47] Sercel. Geophones Specifications. https://www.sercel.com/products/Lists/ProductSpecification/Geophones_specifications_Sercel_EN.pdf.
- [48] Kirk Bertsche, Achim Weidemann, and Walter Wittmer. Calibration of Sercel L4C Calibration of Sercel L4C Geophones. <https://agenda.linearcollider.org/event/4838/contributions/20103/attachments/16311/26600/GeophoneCalibration.pdf>.
- [49] Terrence Tak Lun Tsang. Sensing Matrices for Optical Levers (OpLev) of the KAGRA Main Optics. KAGRA internal document JGW-T2112874: <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=12874>. GitHub repository: <https://www.github.com/terrencetec/kagra-optical-lever>.
- [50] Jer-Nan Juang and Richard Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance Control and Dynamics*, 8, 11 1985.
- [51] Yuji Nakatsukasa, Olivier Sète, and Lloyd N. Trefethen. The AAA algorithm for rational approximation. *arXiv e-prints*, page arXiv:1612.00337, December 2016.

- [52] Alvaro Valera-Rivera and Arif Ege Engin. AAA Algorithm for Rational Transfer Function Approximation With Stable Poles. *IEEE Letters on Electromagnetic Compatibility Practice and Applications*, 3(3):92–95, 2021.
- [53] Kouseki Miyo. Seismic noise of KAGRA mine. KAGRA internal document JGW-T1910436: <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=10436>.
- [54] Alessandro Bertolini, Riccardo Desalvo, Francesco Fidecaro, and Akiteru Takamori. Monolithic folded pendulum accelerometers for seismic monitoring and active isolation systems. *Geoscience and Remote Sensing, IEEE Transactions on*, 7:273 – 276, 03 2006.
- [55] Robin Kirchhoff, C Mow-Lowry, Gerald Bergmann, Manuela Hanke, P Koch, Sina Köhlenbeck, S Leavey, J Lehmann, Patrick Oppermann, J Wöhler, D Wu, Harald Lück, and K Strain. Local active isolation of the AEI-SAS for the AEI 10 m prototype facility. *Classical and Quantum Gravity*, 37, 06 2020.
- [56] Wensheng Hua. *Low Frequency Vibration Isolation and Alignment System for Advanced LIGO*. PhD thesis, Stanford University, 2005.
- [57] Dehaeze Thomas, Mohit Verma, and Collette Christophe. Complementary Filters Shaping Using \mathcal{H}_∞ Synthesis. pages 459–464, 11 2019.
- [58] Masayuki Nakano. Sensor correction investigation. KAGRA Logbook: <https://klog.icrr.u-tokyo.ac.jp/osl/?r=15119>.
- [59] Jonathan J Carter, Samuel J Cooper, Edward Thrift, Joseph Briggs, Jim Warner, Michael P Ross, and Conor M Mow-Lowry. Particle swarming of sensor correction filters. *Classical and Quantum Gravity*, 37(20):205009, sep 2020.
- [60] Michael Ross et al. Towards windproofing LIGO: reducing the effect of wind-driven floor tilt by using rotation sensors in active seismic isolation. *Classical and Quantum Gravity*, 37, 09 2020.
- [61] Robin Kirchhoff. *Implementation of an active seismic isolation system for the AEI 10 m prototype*. PhD thesis, Gottfried Wilhelm Leibniz Universität, 2021.
- [62] G. Zames. Feedback and optimal sensitivity: Model reference transformations, multiplicative seminorms, and approximate inverses. *IEEE Transactions on Automatic Control*, 26(2):301–320, 1981.
- [63] Sigurd Skogestad and I Postlethwaite. *Multivariable Feedback Control: Analysis and Design*, volume 2. 01 2005.

- [64] Vladimir Kučera. A Method to Teach the Parameterization of All Stabilizing Controllers. *IFAC Proceedings Volumes*, 44(1):6355–6360, 2011. 18th IFAC World Congress.
- [65] John Doyle, Francis A, and Allen Tannenbaum. *Feedback Control Theory*. 01 2009.
- [66] J.C. Doyle, K. Glover, P.P. Khargonekar, and B.A. Francis. State-space solutions to standard \mathcal{H}_2 and \mathcal{H}_∞ control problems. *IEEE Transactions on Automatic Control*, 34(8):831–847, 1989.
- [67] Keith Glover and John C. Doyle. State-space formulae for all stabilizing controllers that satisfy and H_∞ norm bound and relations to risk sensitivity. *Systems & Control Letters*, 11:167–172, 1988.
- [68] Pascal Gahinet and Pierre Apkarian. A linear matrix inequality approach to h_∞ control. *Int. J. Robust & Nonlinear Control*, 4(4):421–448, 07 1994.
- [69] Control Systems Library for Python. <http://github.com/python-control/python-control>, 2017.
- [70] Peter Benner, Volker Mehrmann, Vasile Sima, Sabine Van Huffel, and Andras Varga. *SLICOT—A Subroutine Library in Systems and Control Theory*, pages 499–539. Birkhäuser Boston, Boston, MA, 1999.
- [71] P.Hr. Petkov, D.W. Gu, and M.M. Konstantinow. Fortran 77 routines for Hinf and H2 design of continuous-time linear control systems. 1998.
- [72] João P. Hespanha. *Linear Systems Theory : Second Edition*. Princeton University Press,, Princeton, NJ, second edition, 2018 - 2018.
- [73] G.J. Balas, J.C. Doyle, K. Glover, A. Packard, and R. Smith. μ -Analysis and Synthesis Toolbox. The MathWorks Inc., Natick, MA, 1995.
- [74] T. T. L. Tsang, T. G. F. Li, T. Dehaeze, and C. Collette. Optimal Sensor Fusion Method for Active Vibration Isolation Systems in Ground-Based Gravitational-Wave Detectors. 11 2021.
- [75] H. Abe et al. Performance of the KAGRA detector during the first joint observation with GEO 600 (O3GK). 3 2022.
- [76] Takafumi Ushiba. External Review -Suspension commissioning for test mass suspension-. KAGRA internal document JGW-G2113194: <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=13194>.

- [77] Carsten Scherer. Mixed H_2/H_∞ Control. In Alberto Isidori, editor, *Trends in Control*, pages 173–216, London, 1995. Springer London.
- [78] E Schwartz et al. Improving the robustness of the advanced LIGO detectors to earthquakes. *Classical and Quantum Gravity*, 37(23):235007, nov 2020.
- [79] Michael Nikolaou. Model predictive controllers: A critical synthesis of theory and industrial needs. *Advances in Chemical Engineering*, 26:131–204, 2001.
- [80] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python. <http://www.scipy.org/>, 2001–.
- [81] T.T.L. Tsang. Kontrol. <https://kontrol.readthedocs.io/>, 2020–.
- [82] Fuchang Gao and Lixing Han. Implementing the nelder-mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications*, 51:259–277, 05 2012.
- [83] Rainer Storn and Kenneth Price. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11:341–359, 01 1997.
- [84] Reinoud Sleeman. Three-channel correlation analysis: A new technique to measure instrumental noise of digitizers and seismic sensors. *The Bulletin of the Seismological Society of America*, 96:258–271, 02 2006.
- [85] *Advanced gravitational wave detectors*. Cambridge University Press, Cambridge, 2012.
- [86] R. Abbott et al. GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. 8 2021.