LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T2600198–** 2026/05/15

A System Identification Framework with RL-Based Training for Technical Noise Reduction in the Caltech 40m Prototype

Michael Padilla

California Institute of Technology
LIGO Project, MC100-36
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

Contents

1 Introduction 2

2 Objective 3

3 Approach 4

4 Timeline 5

Abstract

This document proposes the development of an accurate, dynamically updatable digital twin of the Caltech 40m Prototype Input Mode Cleaner (IMC) to serve as an offline testing and training environment. We utilize a gray-box system identification approach to model the suspended optics and triangular cavity, tuned with and validated against empirical hardware data. A digital twin is compiled using the 'rtsfreerun' framework, ensuring compatibility with LIGO's real-time control architecture. Finally, a validated digital twin distribution serves as the simulated plant to train a reinforcement learning policy. Through deployment of the resulting policy, this project aims to reduce technical noise and thereby contribute to the sensitivity targets of future upgrades and next-generation detectors.

1 Introduction

Since 2015, LIGO has detected dozens of astrophysical transients, providing significant insights into objects such as black holes and neutron stars, while additionally allowing us to test prominent theories and models like strong-field general relativity [1, 2]. To further push the detection horizon of gravitational wave interferometers, upgrades to LIGO and next-generation detectors are envisioned to come online in the next decade, focusing on improvements in sensitivity and more efficient power build up [3, 4]. Present noise budget considerations are categorized as high frequency (> 300 Hz), mid frequency (50-300 Hz), or low frequency (< 50 Hz). High frequency noise is primarily dominated by quantum shot noise, which will be suppressed by the upgrade's higher power and squeezing targets. Mid frequency noise is additionally addressed through changes to test-mass coatings which would lower coating thermal noise. At low frequencies, the primary consideration of these future upgrades involves changes to the test masses and suspensions to reduce radiation pressure and seismic noise. However, these design considerations will not provide their stated benefits unless low-frequency technical noise is considerably reduced to take advantage of reduced fluctuations in the interferometer's degrees of freedom [3]. This proposed project directly addresses this contribution to future design considerations through the implementation of a modeling to reinforcement learning pipeline that can contribute to a reduction in control noise.

Present day control systems which are used to 'lock' or control the various degrees of freedom of the interferometer are required for sensitive operation and can stabilize unstable systems but may inject sensor noise due to the sensor's limited sensitivity at particular frequencies. Recent work on a new approach to solving this problem has shown that the implementation of reinforcement learning (RL) algorithms to control loops can reduce injected control noise by more than a factor of 10 [5]. In order to create an effective RL model, its training environment must accurately model the behavior of the system it will control. Creating accurate models that reflect the behavior of real-world counterparts is especially difficult in the complex optical subsystems of the LIGO interferometer, particularly because of the complex nonlinear dynamics of these systems, which may be difficult to reproduce with a ground-up physics-based simulation. For this reason, this project explores an alternative approach to modeling, system identification, which will be explored to create accurate simulated models to facilitate a successful sim-to-real transfer of an RL agent policy.

2 Objective

The focus of this project is to expand the scope of the previous RL work to additional control loops in the interferometer through a system identification (sysID) approach to model building with the ultimate goal being the demonstration of reduced technical noise. The project I propose to work on this summer with the LIGO SURF program can be broken down into three pillars, each building on the last: sysID, digital twins, and RL-based control loops. Our goal is to develop a robust sysID and digital twin scheme that allows for accurate modeling and prediction of a subsystem of the interferometer by tailoring the model from the system using input and response data from the real-world counterpart. System identification will allow for the creation of digital twin models of the subsystem which will then be used for implementation of RL in control loops that have not yet been explored. The sysID to RL scheme will therefore allow us to better explore methods to reduce various sources of control noise, contributing to the design targets of future LIGO upgrades and next-generation detectors.

The three key deliverables for this project are as follows. Each deliverable is based on the successful development of the prior deliverable.

1. **A gray-box model of the input mode cleaner (IMC):** A gray-box model entails the use of both system identification methods which utilize real-world input and response data to fit unknown parameters of the system, as well as model components built using known physics and first principles. This model should include accurate modeling of the suspension system, optomechanical dynamics and couplings, as well as seismic, thermal, and technical noise sources.

2. **A distribution of digital twin models:** These models will feature slightly tuned parameters and noise sources to reflect real-world uncertainties and time-dependence. This digital twin model will be compiled using the ‘rtsfreerun‘ framework, which utilizes the exact same signal processing architecture and software interfaces as the Advanced LIGO real-time control system. This ensures the simulation handles controls and feedback identically to the physical interferometer. This additionally guarantees that our RL policy trains in an environment matching the exact deployment conditions of the 40m prototype. In tandem with this deliverable, we will also develop a system to evaluate the resulting digital twin to determine its accuracy relative to a determined requirement.

3. **A trained and deployed RL model:** The digital twin distribution serves as the environment, or simulated plant, in which a neural network acts as the controller or alongside a standard linear controller. The resulting model policy will be integrated directly into the 40m prototype and validated against existing linear control loops.

Not only will this work help with the ultimate technical noise reduction goal of the project, but the resulting digital twin models have applications in future prediction and minimizing downtime in the system. Currently, testing new control configurations on the physical interferometer carries the operational risk of destabilizing the system and breaking the cavity lock, potentially costing hours of active sensing time. With the ability to create digital

twins, many instances of the model can be simulated offline, each with a different control configuration. This would allow commissioners to predict more accurately the resulting interferometer state after a specific adjustment, reducing potential downtime and increasing the chances for event detection.

Further, this framework allows the digital twin to be dynamically updated with real-time data. In the real interferometer, state parameters naturally drift overtime, such as the thermal state. A static model therefore can become inaccurate or obsolete over time. By utilizing a sysID pipeline, the interferometer can be periodically probed to determine its current parameters. This allows for the baseline of the digital twin to be constantly updated, ensuring the simulation remains an accurate, real-time reflection of the physical interferometer.

3 Approach

The project will first involve the construction of a sysID model for the input mode cleaner (IMC) of the CIT 40m Prototype. The IMC is a triangular, suspended cavity designed for spatial mode cleaning and frequency stabilization (see Figure 1). To do this, we will take a gray-box approach to model building, which uses a combination of first principles from well understood dynamics of the IMC with model fitting from input and response data that we will collect in the CIT 40m prototype. We will use Julia and the Julia tool SciML to construct physics-based models of subsystems of the IMC such as the single-stage suspensions (see Figure 2). Python will be used to create and tune the sysID models. Finally, we will use a combination of Matlab tools like Simulink as well as Python to construct and interact with the digital twin models. We are particularly interested in modeling the dynamics of optics, suspension systems, and perturbations such as seismic and sensor noise, which are most important for an accurate control loop model. Some prior analyses have been performed to better understand the behavior of the IMC and single-stage suspension dynamics [6, 7]. This prior work will be used as a conceptual foundation for our constructed sysID models.

Next, we will develop a framework to evaluate the performance of our gray-box model using relevant frequency- and time-domain characteristics, alongside corresponding error metrics. This evaluation may additionally involve residual analysis of the model's predictions.

Following the construction of our sysID model, we will incorporate a digital twin framework which uses our model in tandem with real-time data to create a digital counterpart to the real-world system to allow for realistic modeling and prediction of the present system. To do this, we will take our constructed sysID model and translate it directly into a Python module using the existing 'rtsfreerun' Gitlab repository. This allows us to instantiate our model and inject signals within a Python script. This structure enables the implementation of non-linear control loop schemes, making the transfer to the RL stage straightforward. We will create a distribution of many digital twins, each with slightly varying parameters and noises based on their determined uncertainties. The distribution of many digital twin models ensures that the resulting policy is robust to varying conditions in the interferometer.

With the development of digital twins and underlying sysID models, we can develop an RL scheme for non-linear control loops for the various degrees of freedom in the IMC. We treat the digital twin as the simulated plant in which a neural network acts as the controller or

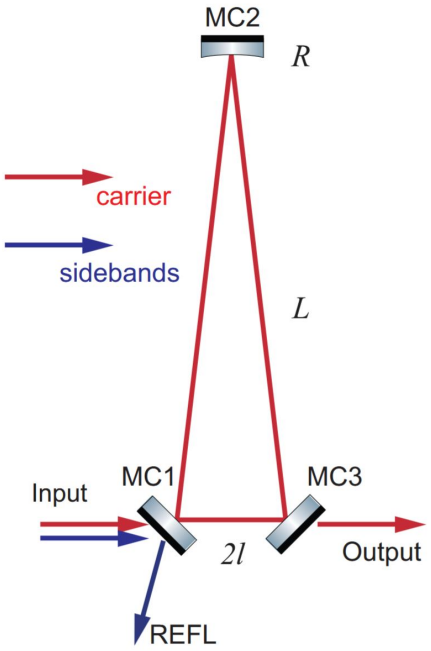


Figure 1: Simple diagram of the IMC topology. Reproduced from [6].

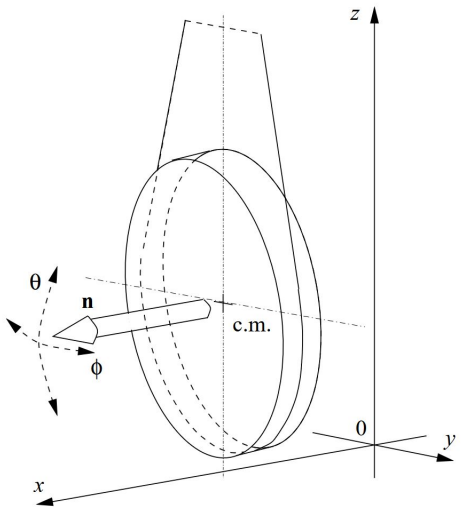


Figure 2: Schematic of the single-stage suspended mirror illustrating the relevant degrees of freedom. Reproduced from [7].

alongside a standard linear controller. We can therefore structure our training scheme to reward the network for actuation which keeps the IMC stably locked, minimizes transient response, and most importantly minimizes injected control noise. We will use the Deep Loop Shaping RL method [5] developed to reduce the injected noise of the common-hard-pitch loop. Additionally, we propose to implement a similar Maximum A Posteriori Policy Optimization architecture as was used in the Deep Loop Shaping study, as it is well suited to handle the continuous action spaces like actuator signals. To develop the architecture for our agent and training logic, we will use Python packages such as PyTorch. To evaluate the model, we will initially construct a classical linear control loop for the IMC system. This will serve as the baseline for our RL agent to be digitally compared.

Finally, after the RL agent has been sufficiently trained and validated, the resulting policy will be deployed in the real CIT 40m prototype IMC, where its performance will be analyzed and benchmarked against the existing linear controllers in the IMC.

4 Timeline

To facilitate a successful project during my 10 weeks at the California Institute of Technology, I propose a timeline of the work we will conduct over each of the ten weeks. During this period of time familiarization of the relevant tools (Matlab, Simulink, Julia, 'rtsfreerun') will be integrated with the development of associated components of the project. The bulk of the work anticipated for the project can be split up into three distinct phases, each with focus on an individual pillar of the project: System modeling, digital twin development, and RL training and deployment.

Week(s)	Milestone
1–3	**System Modeling & Data Collection:** Development and deployment of the gray box system identification scheme for the IMC. Remaining subsystems of the IMC which have not yet been modeled (ex. suspensions, individual optics, noise sources) will be constructed. Input and response data will actively be collected from sensors in the CIT prototype in collaboration with the 40m team which will be used to tune the sysID specific components of the model. A section of this data will be used to directly tune the model, while a remaining portion will be validation to evaluate its performance. These models will then be combined into a single gray box system model. Writing of the first interim report.
4–6	**Digital Twin Development:** Development and validation of the digital twin framework using empirical data from the LIGO 40m prototype. Validate the model using the held-out validation data set collected from the CIT IMC. Write the second interim report.
7–9	**RL Training & Deployment:** Implementation and training of the Deep Loop Shaping RL architecture within the digital twin environment. Deploy and validate the resulting policy directly on the physical CIT IMC.
10	**Finalization:** Finalize measurements and data analysis. Complete the project documentation and present the LIGO SURF report.

Table 1: Project timeline with key deliverables per phase.

References

[1] B. P. Abbott et al. "Observation of GWs from a Binary BH Merger". In: *Phys. Rev. Lett.* 116 (6 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. URL: https://link.aps.org/doi/10.1103/PhysRevLett.116.061102.

[2] Vitor Cardoso and Paolo Pani. "Tests for the existence of black holes through gravitational wave echoes". In: *Nature Astronomy* 1.9 (Sept. 2017), pp. 586–591. ISSN: 2397-3366. DOI: 10.1038/s41550-017-0225-y. URL: https://doi.org/10.1038/s41550-017-0225-y.

[3] P. Fritschel et al. *LSC Post-O5 Study Group Report*. Tech. rep. LIGO-T2200287. LIGO Scientific Collaboration, 2022. URL: https://dcc.ligo.org/LIGO-T2200287/public.

[4] M. Evans et al. *A Horizon Study for Cosmic Explorer: Science, Observatories, and Community*. 2021. arXiv: 2109.09882 [astro-ph.IM]. URL: https://arxiv.org/abs/2109.09882.

[5] Jonas Buchli et al. "Improving cosmological reach of a gravitational wave observatory using Deep Loop Shaping". In: *Science* 389.6764 (2025), pp. 1012–1015. ISSN: 1095-9203. DOI: 10.1126/science.adw1291. URL: http://dx.doi.org/10.1126/science.adw1291.

[6] Satoshi Tanioka et al. *Angular response of a triangular optical cavity analyzed by a linear approximation method*. 2020. arXiv: 2002.02703 [physics.optics]. URL: https://arxiv.org/abs/2002.02703.

[7] M. Rakhmanov et al. *Dynamical Properties of LIGO Single Loop Suspended Mirrors*. Tech. rep. LIGO-T000134-00-D. Laser Interferometer Gravitational-Wave Observatory (LIGO), Aug. 2000.