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

Technical Note **LIGO-T11XXXXX–vX** 2023/06/09

LIGO SURF 2023
Interim Report I

Demonstrating Optimal Non-linear Temperature Control

Advait Mehla

California Institute of Technology
LIGO Project, MS 18-34
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/

1 Overview

Precise temperature control in the presence of noisy environments and heat loss through complex channels and a combination of conduction, convection and radiation is a challenging task. This is because the system is inherently non-linear, and subject to large disturbances in the low frequency region due to day-to-day changes in the ambient temperature on the order of several degrees. Traditional PID control based approaches struggle to attain and precisely maintain a desired set-point in the presence of such fluctuations. The optimal tuning parameters vary quite a bit with the ambient conditions and the performance thus varies with external factors.

An optimal control loop for such a non-linear system would require the knowledge of the whole state space, where the actuation is not a linear and separable function of the system parameters. We aim to achieve this kind of adaptive control by utilising neural networks to make the controller nonlinear. We plan to train the neural network via reinforcement learning, but this will be dealt with mostly in the latter half of the SURF.

After my arrival here at Caltech and discussions with Prof. Adhikari, Dr. Caracoba and Dr. Bhatt at the LIGO 40m lab, I learnt a lot more about the trajectory of the project and got short and long term targets to work towards. I will be working on this project with Andrei Diaconu, who is a Caltech undergrad also SURFing at the LIGO 40m lab.

Before graduating to the actual system, which is a large seismometer surrounded by several inches of insulating foam, we decided that it would be better to test our electronics and methods on a smaller, toy model that would be easier to simulate as well as experiment with. We found a small cylindrical aluminium mass, that I refer to as the 'puck' moving forward.

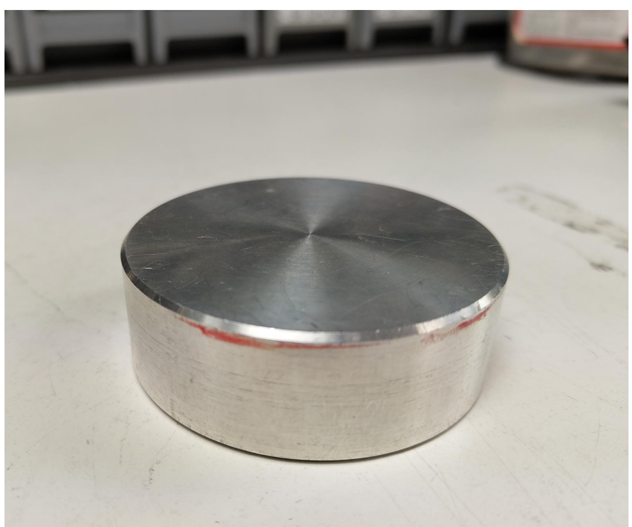

Figure 1: Aluminium puck with a diameter of 75 mm and thickness of 25 mm

It has a mass of 298 g, and this gives it a relatively small heat capacity that should allow us to do simple tests within short timescales of a few hours, with a low power heater. This mass will eventually be covered with about an inch of insulation foam when we attempt to control the system.

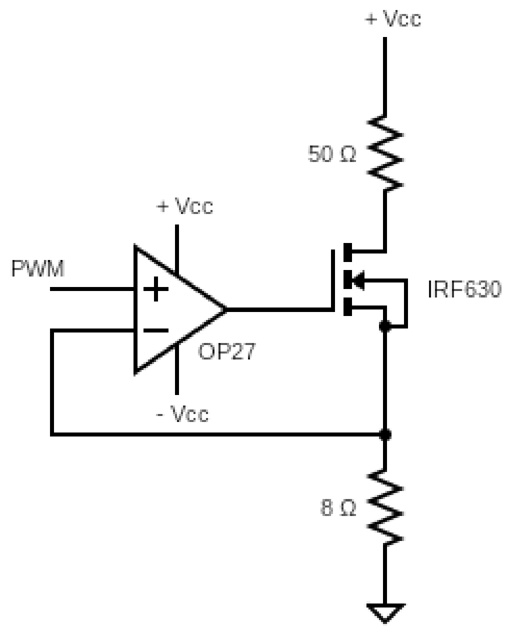

Figure 2: The heater circuit schematic

2 Heater Circuit

For the heating circuit, I have designed a power MOSFET based current source (Figure 2) which can deliver a peak power of roughly 7.5 W to the puck. The exact power level can be varied using pulse-width modulation (PWM).

This is a makeshift circuit currently implemented on a breadboard for preliminary testing, but can be scaled up easily once we decide the requirements for the actual system and buy high power resistors and heating elements. The input signal is applied by a Raspberry Pi 4B, which has a few 3.3V PWM pins. The mass acts as an integrator and averages the power supplied to it, letting us regulate the power supplied easily. This method also minimizes the heat dissipation through the MOSFET, by operating it only in the fully-on state in which the drain-source resistance is minimized. The Pi was chosen because we eventually plan to use neural networks, and evaluating control signals at a high rate would demand a considerable amount of computational power.

In this circuit, V_S of the MOSFET is equal to the input voltage, which is about 3.1 V when the input is high, slightly lower than 3.3 V due to an indicator LED which is driven by the same pin for visual feedback. Thus, the average current through the heating element can be controlled easily via PWM, and has a peak value of $3.3/8 = 0.3875$ A. This gives us a peak power of $0.3875^2 \times 50 = 7.51$ W through the heating resistor. This is housed in brass, which lets it efficiently conduct heat to the puck by just taping it to the top.

The resistance combination was limited by availability of high power resistors in the lab, and carefully arranged to make sure none of them burn out. The 8 Ω resistance is therefore

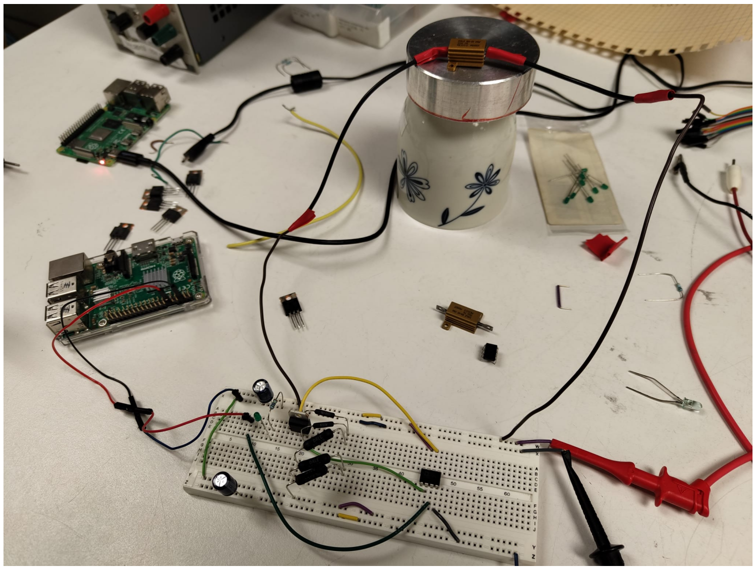

Figure 3: The full circuit, with the heating resistor kept on the puck

actually a series-parallel combination several individual resistors visible on the breadboard.

3 Thermal Modelling

In order to design a controller for the full puck-foam system, we have to first design a model to accurately represent this combined mass. The puck is fully wrapped by a layer of insulation of 1 inch thickness (d). The foam will be further wrapped by aluminium foil to lower the emissivity and insulate it even more.

In this system, the puck is at a temperature T_p, has a mass and specific heat capacity of m_p and c_p. Similarly, the surface of the foam is at T_f and has m_f and c_f. I assumed that the temperature gradient within the foam can be neglected, and the entire mass can be assumed to be at T_f for estimating its rate of heating. Additionally, the cross-sectional area of conduction is considered to be the total surface area of the cylinder enclosed by the surface at half of the foam-width (represented by A).

Now, the puck loses heat only through conduction across the foam, and gains heat from the heater at a rate H. This can be expressed as below -

$$\frac{dQ_p}{dt} = m_p c_p \frac{dT_p}{dt} = -\frac{kA}{d}\left(T_p - T_f\right) + H$$

Here, k is the foam conductivity.

Similarly, for the foam, the influx of heat is due to conduction from the puck, while it loses heat to the environment via radiation through the total exposed surface area (S).

$$\frac{dQ_f}{dt} = m_f c_f \frac{dT_f}{dt} = \frac{kA}{d}(T_p - T_f) - Se\sigma\left(T_f^4 - T_{env}^4\right)$$

Here, $e \approx 0.07$ for aluminium foil.

This gives the following system of coupled differential equations -

$$\frac{dT_p}{dt} = -\frac{kA}{dm_p c_p}\left(T_p - T_f\right) + \frac{H}{m_p c_p} \tag{1}$$
$$\frac{dT_f}{dt} = \frac{kA}{dm_f c_f}(T_p - T_f) - \frac{Se\sigma}{m_f c_f}\left(T_f^4 - T_{env}^4\right) \tag{2}$$

which describe the evolution of the system of choice given the parameters.

This system can now be numerically integrated using a library like `scipy.odeint` to yield the time evolution of the puck and foam temperatures given a function describing the heating and ambient temperature. Figure 4 shows a step response plot, where the heating was turned on at a power of 1 W and then shut off to allow the system to cool freely.

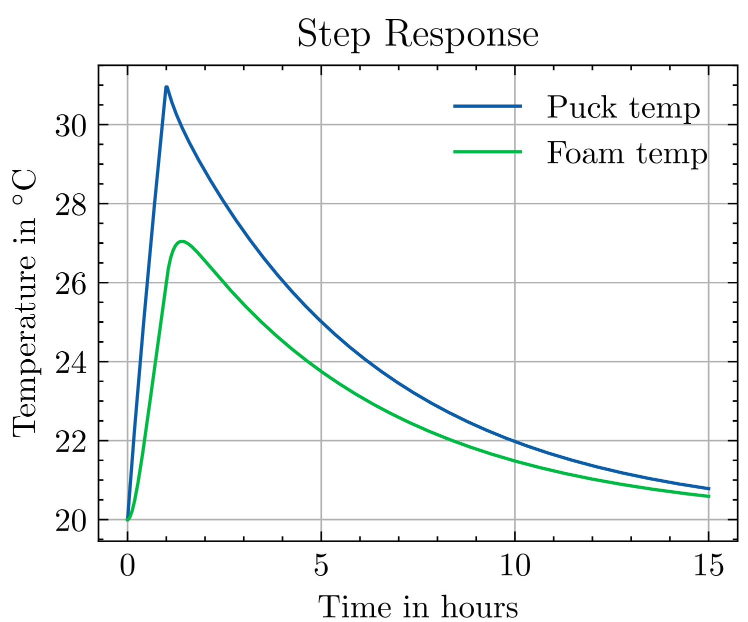

Figure 4: Heater step response

This model can be improved further by incorporating heat loss via convection from the foam to the environment. Convective heat loss can be given by

$$\frac{dQ_{conv}}{dt} = hA\left(T_f - T_\infty\right)$$

Where T_∞ is the temperature of air sufficiently far away from the foam surface, analogous to T_{env}, while h the convective heat transfer coefficient which is not a property of the fluid, but an experimentally determined quantity whose value depends on all the variables influencing convection such as the surface geometry, the nature of fluid motion, the properties of the fluid, and the bulk fluid velocity [1]. For free convection, its value can be anywhere between

2 - 25. As this is a big range, convection has not been included for now. We plan to attempt to fit this factor using experimental data eventually.

4 Future Work

The following tasks are to be undertaken by us over the coming weeks -

1. Building a temperature sensor circuit based on the AD590 temperature sensor IC, and connecting it to the puck. There is already such a circuit built for the seismometer several years ago, and we may use that as it is for the puck system. However, it still needs to be calibrated, and we plan to do this using an ice-water bath and the human body temperature as the two reference points. Once the temperature sensor is calibrated and integrated with the Raspberry Pi to read out data, we can attempt to conduct step response experiments and match the data to the simulations.

2. Fitting free or uncertain parameters of the model using the real data. We can use the Python package `lmfit` to estimate uncertain parameters in the differential equations like the convective factor, h using real data.

3. Once it is established that the un-insulated foam model works as predicted by our simulations, we will increase the complexity and include the foam in the model as well as physical system and again match the real data with the predictions of the model.

4. Finalising the heater circuit design after deciding requirements like heat-up time and maximum power for the seismometer mass, and making it on a PCB.

While I work on these action items, Andrei is currently working on designing and simulating a higher-order linear controller to control this system optimally. This is meant to be our benchmark for comparing with the eventual RL-based nonlinear controller, which is expected to outperform linear controls. He has also been reading up on reinforcement learning policies and their implementation using the Python libraries `tensorflow` and OpenAI's `gym` and `acme`.

Once the hardware is set up, we would be able to implement a linear controller designed by us and analyse its performance. This would then be followed by training and implementing the actual RL controller.

References

[1] Cengel, Y.A., *Heat Transfer: A Practical Approach. 2nd Edition, McGraw-Hill, New York.*