ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST027

MACHINE LEARNING-BASED TUNING OF CONTROL PARAMETERS FOR LLRF SYSTEM OF SUPERCONDUCTING CAVITIES

J. Diaz Cruz*, S. Biedron[1], S. Sosa, R. Pirayesh, University of New Mexico, Albuquerque, NM, USA
[1]also at Element Aero, Chicago, IL, USA

Abstract

The multiple systems involved in the operation of particle accelerators use diverse control systems to reach the desired operating point for the machine. Each system needs to tune several control parameters to achieve the required performance. Traditional Low-Level RF (LLRF) systems are implemented as proportional-integral (PI) feedback loops, whose gains need to be optimized. In this paper, we explore Machine Learning (ML) as a tool to improve a traditional LLRF controller by tuning its gains using a Neural Network (NN). We present the data production scheme and a control parameter optimization using a NN. The NN training is performed using the THETA supercomputer.

INTRODUCTION

The LLRF system is in charge of controlling the amplitude and phase of the electromagnetic field that drives superconducting RF (SRF) cavities. For facilities like the Linac Coherent Light Source II (LCLS-II), tight requirements for amplitude and phase are defined: 0.01% and 0.01 degrees, respectively [1]. The quality of the X-rays produced by this type of facilities depends on the quality of the electron beam.

LLRF systems wit a single source single cavity configuration, like the one for LCLS-II, use traditional PI loops for amplitude and phase control, having a total of 4 parameters to be tuned. Tuning these parameters takes into account quantities like the cavity gain and cavity bandwidth, the closed-loop bandwidth and latency, and the amplitude set-point. In this paper, we propose a ML-based tuning of the LLRF controller parameters, which uses a NN to calculate the optimal proportional and integral gains to minimize amplitude and phase errors. We also present a data production scheme based on simulations and an algorithm for stochastic optimization [2].

LLRF MODEL AND DATA PRODUCTION

Traditional LLRF controllers for particle accelerators are PI controllers like the one shown in Fig. 1. It usually consists of a couple of feedback loops: one for amplitude and one for phase. Therefore, the controller has 4 parameters: proportional and integral gains for amplitude and phase. Tuning this parameters is not a trivial task and can be time consuming for control room operators, specially when the accelerator has a large amount of SRF cavities (280 SRF cavities in the case of the LCLS-II). The tuning process can be automated based on control theory and the desired behaviour of the closed-loop, taking into account quantities like the cavity gain and cavity bandwidth, the closed-loop bandwidth and latency, and the amplitude set-point. However, this automation does not guarantee optimal parameters. Furthermore, drift in the system parameters would require to perform characterization of the cavity and the closed-loop system parameters multiple times.

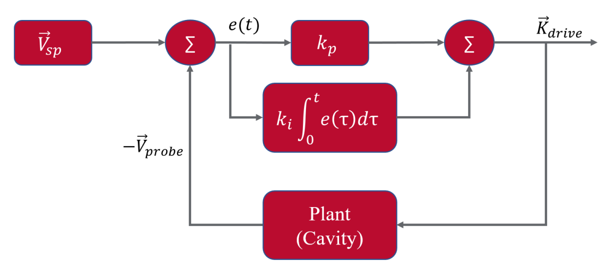

Figure 1: Simplified diagram of a traditional PI LLRF controller.

In this paper, we propose a tuning process to minimize the amplitude and phase stability errors, and a NN can be trained to learn this optimization. In the next subsections, we explain the cavity model and the simulations of the cavity field feedback loop.

Cavity Model

A multi-cell SRF cavity can me modeled as a group of RLC circuits (resonant circuits), each one corresponding to an eigenmode of the cavity. Figure 2 is the equivalent RLC circuit of each eigenmode. The differential equations that describe the electrodynamics of the systems are derived in [3] and result in the following set of equations:

$$V = Se^{j\theta}, \tag{1}$$

$$\frac{d\theta}{dt} = w_d, \tag{2}$$

$$\frac{dS}{dt} = -w_f S + w_f e^{-j\theta}(2K_g\sqrt{R_g} - R_b I), \tag{3}$$

where V is a representative measure of each mode's energy, with magnitude S and phase θ, w_d is the detuning frequency and w_f is the cavity bandwidth. K_g is the incident wave amplitude, which represents the power that drives the cavity, R_g is the coupling impedance of the beam, and I represents the beam current.

Using this model for simulations and data production, we have simulated a LLRF closed-loop under different setting points of RF power and electron beam current. Additionally, different levels of cavity frequency detuning and measurement noise can be applied to the simulations.

* jorgedc@unm.edu

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST027

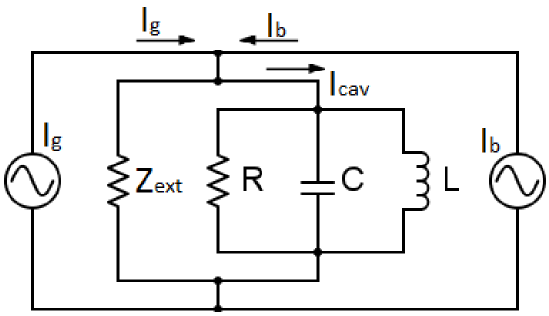

Figure 2: Cavity's circuit model of a resonant mode.

Closed Loop Simulations

Following the example of the CMOC code [3], we developed our own Python code to simulate a feedback LLRF loop. It consists of a cavity, a power source (a Solid State Amplifier in the case of the LCLS-II) and a PI controller. The following perturbations are implemented: beam loading disturbances, cavity detuning and measurement noise. The effects of these perturbations in the stability of the cavity are explain in detail in [4, 5]. It is important to clarify the cavity detuning model: it is a sinusoidal variation of the cavity resonant frequency. Notice that the frequency of this sinusoidal variation represents the frequency of the detuning source (for example, a pump) and the amplitude represents the effective detuning of the cavity (for example, 15 Hz of detuning), which also has units of Hz.

In Fig. 3, the top plot shows a cavity field (blue curve) reaching the set point (dotted black line) due to the action of the control signal U (red curve). We can see the saturation of the power source for the first 12 ms of the simulation, and the effect of 10 Hz detuning (oscillating at 200Hz) in the control signal U. We can also see the effect that beam loading has over the control signal U, when more power is required when the beam is present (from 15 ms to 25 ms).

The bottom plot of Fig. 3 shows a detail of the cavity's voltage signal to see how it is affected by the perturbations: we see the oscillations due to microphonics, the noise related to the measurement noise, and an undershoot and overshoot related to the start and end of the beam. Notice that this plot starts at 15 ms. The upper and lower limits for amplitude stability are shown for reference. The magnitude of the effects related to the perturbations mentioned above is a function of the control parameters (proportional and integral gains). In the next section we explain how to optimized this parameters to minimize the stability error.

CONTROL PARAMETER OPTIMIZATION

The 0-dB crossing of the closed-loop depends on the proportional gain k_p. Figure 4 shows the relation between 0-dB crossing (and therefore k_p) and the RMSE of the cavity's field amplitude. There is an optimal k_p that minimizes the error. Applying an algorithm for stochastic optimization using

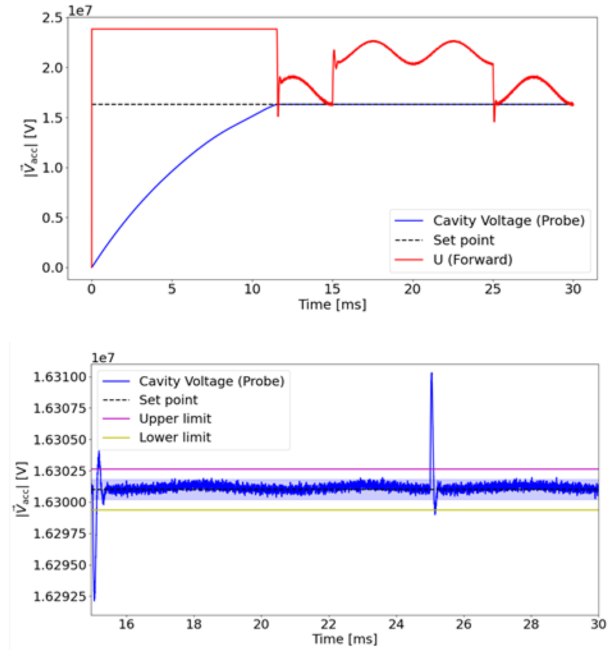

Figure 3: Top: Simulation of the LLRF closed-loop with beam loading disturbances, cavity constant detuning and measurement noise. Bottom: Detail of the cavity's field.

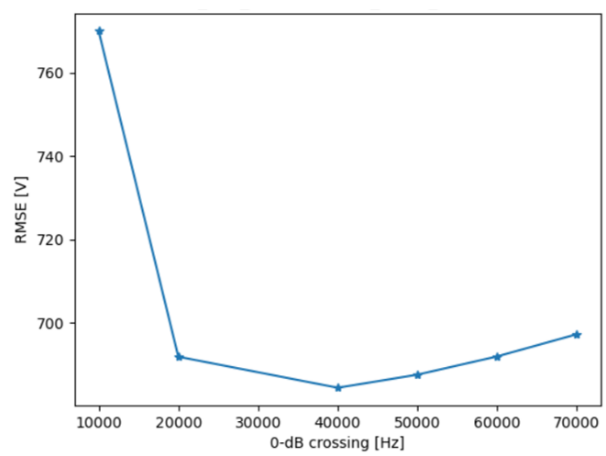

Figure 4: Relation between 0-dB crossing (or proportional gain) and RMSE of cavity's field amplitude.

the Python library Noisyopt, described in [2] and represented by the following equation:

$$\min_{k_p} f(k_p) = \min_{x} E[F(k_p, \xi)], \quad (4)$$

one can find the optimal value of the gain. In Eq. (4), $f(x)$ represents the RMSE error as a function of the gain, which cannot be directly evaluated, and $F(x, \xi)$ represents the function that we can simulate and their dependency on a noise ξ. Using the optimization algorithm and the simulations of the closed-loop LLRF system described above, the results of the optimization are shown in Fig. 5. The algorithm can find the optimal gain.

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST027

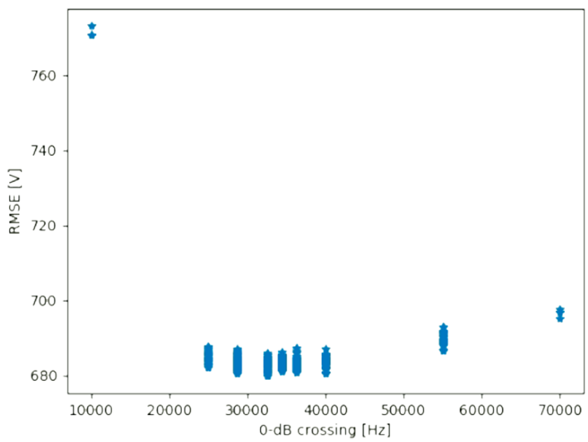

Figure 5: Optimization results.

ML ARCHITECTURE

A diagram of the ML-based LLRF control system proposed in this paper is shown in Fig. 6. For the optimization and data production phase, x_0 represents the inputs to the optimization algorithm and to the training of the ML. x_0 is the cavity detuning, measurement noise, beam current and amplitude and phase set-points of the cavity's field. y represents the optimal gains calculated by Noisyopt. x_0 and y together built the training dataset for the ML. Once the ML is trained, it will be able to calculate the optimal gains, $\tilde{y}$, for conditions not seen before, $\tilde{x_0}$, and the LLRF controller will use this optimized gains online.

For the optimization and data production phase, and learning phase, we use the resources of the THETA supercomputer at the Argonne Leadership Computer Facility [6]. The experimental phase should be implemented along with the LLRF controller in an FPGA, or in an upper level of software.

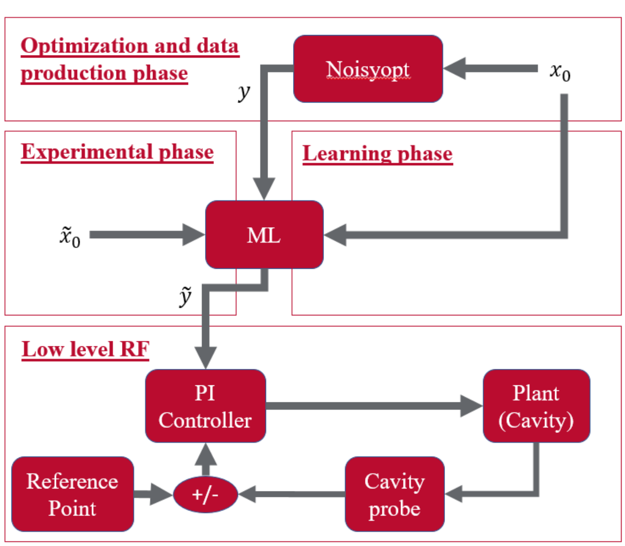

Figure 6: Diagram of the ML-based LLRf control system.

SUMMARY AND FUTURE WORK

Simulations of a closed-loop LLRF system have been implemented based on the CMOC software engine. An algorithm for stochastic optimization called Noisyopt has also been implemented to calculate the optimal proportional gain that minimizes the RMSE of the cavity's field amplitude. An ML-based LLRF controller has been proposed. For the training of the ML, data is produced with the implemented software for simulations and with the Noisyopt algorithm. We plan to develop and deploy the proposed ML-based controller to test it with a cavity emulator.

ACKNOWLEDGEMENTS

This work is being possible thanks to the guidance of the members of the LCLS-II LLRF team, specially Larry Doolittle, Carlos Serrano and Gang Huang at LBNL, and Andy Benwell and Alex Ratti at SLAC.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under award number DE-SC0019468

This work was performed under Contract No. DE-AC02-76SF00515 supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences

This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility. This research also used computing resources of Element Aero

REFERENCES

[1] C. Hovater, "Performance and Functional requirements for the LCLS-II Low Level RF System", SLAC, CA, USA, LCLSII-2.7-FR-0371-R2, Sep. 2017.

[2] Noisyopt: A Python library for optimizing noisy functions, `https://noisyopt.readthedocs.io/en/latest/index.html/`.

[3] C. Serrano *et al.*, "LCLS-II System Simulations: Physics", LBNL, Berkeley, CA, USA, Oct. 2015. `https://beamdocs.fnal.gov/AD/DocDB/0049/004978/001/physics.pdf`

[4] J. Diaz Cruz *et al.*, "Studies in Applying Machine Learning to LLRF and Resonance Control in Superconducting RF Cavities", Oct. 2019. `doi:10.48550/arXiv.1910.07648`

[5] C. Serrano, L. Doolittle, and C. Rivetta, "Modeling and Simulations", SLAC, CA, USA, LCLS-II LLRF Review, Mar. 2016.

[6] Argonne Leadership Computer Facility: Theta, `https://www.alcf.anl.gov/theta/`.