# AI-ENABLED ADAPTIVE CONTROL OF BEAM CURRENT IN AN ISOCHRONOUS CYCLOTRON USING GA-TUNED PID

C. Lopez Osses*, G. Soto, M. Backfish, I. Lopez, C. Benko,
M. Novotny, R. Sahedzada, HG. Berns, E. Prebys†
University of California, Davis, Davis, United States

### *Abstract*

A Proportional-Integral-Derivative (PID) controller has been developed and commissioned in the new digital control system at Crocker Nuclear Laboratory for feedback regulation of the extracted proton beam current from the 76-inch isochronous cyclotron. The single-input, single-output control loop uses the outer Trimming Coil (TC10) as the actuator and the extracted beam-current monitor as the feedback signal. Since the response of the extracted beam current to TC10 adjustments can vary with cyclotron operating conditions, PID gain selection was automated using a genetic algorithm. The algorithm searches predefined ranges of proportional, integral, and derivative gains and minimizes a weighted cost function that combines normalized beam-current tracking error, residual steady-state error, TC10 actuator motion, actuator saturation, and oscillatory response penalties. The resulting gains are then applied to the real-time PID controller for beam-current setpoint tracking. Initial commissioning measurements demonstrate stable closed-loop regulation of the extracted beam current and indicate that GA-assisted PID tuning is a viable approach for automated current control in cyclotron operation. This work serves a proof of principle for more complex machine learning implementation planed for the near the future.

## INTRODUCTION

The Crocker Nuclear Laboratory (CNL) [1] at the University of California, Davis is modernizing its 76-inch isochronous cyclotron by integrating a deterministic digital control layer with the existing analog infrastructure. Building on prior CNL work in automated monitoring and AI-ready control [2, 3], the upgraded system combines LabVIEW-based control, CompactRIO hardware, Python interfaces, and ZeroMQ communication. This architecture enables synchronized ADC/DAC operation, time-aligned logging, real-time visualization, and software-defined feedback control at a nominal 100 Hz update rate, providing the foundation for autonomous cyclotron tuning and regulation.

In routine operation at CNL, stabilization of the extracted beam current is performed using a beam-based tuning strategy, in which operators monitor the extracted beam current read-back and manually adjust the cyclotron magnetic field. This practice is motivated by the role of trim coils (TC's) in cyclotrons; where localized magnetic-field corrections can be used to compensate field errors, influence beam phase, and to provide isochronous fields for different beam species [4, 5]. TC10 provides a wide enough magnetic field that it can compensate for magnetic fluctuations from the main field magnet. Therefore, TC10 serves as an empirical control input for correcting slow changes in beam transmission and extraction conditions caused by power source fluctuations, magnetic-field drift, or other operating variations.

This work applies the upgraded digital control layer to Single Input Single Output (SISO) feedback regulation of the extracted beam current using TC10 as the actuator. Because the TC10-to-beam-current response can vary with source conditions, magnetic-field drift, and beam-transport nonlinearities, fixed PID gains may not remain optimal across operating points. A genetic algorithm (GA) is therefore used to tune the bounded PID gains by minimizing normalized beam-current error [6, 7], following the use of global optimization methods for accelerator parameter tuning [4]. The goal is to reproduce manual TC10 tuning in real time while reducing operator intervention and improving beam-current stability [8, 9].

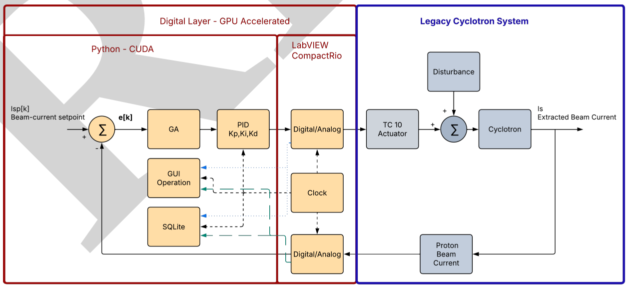

Figure 1: Closed-loop GA-tuned PID control architecture for extracted beam-current regulation.

## SYSTEM ARCHITECTURE

### *Operational Feedback Architecture*

The beam-current regulation system is implemented as a software-defined feedback layer within the upgraded CNL digital control infrastructure. The extracted beam current is measured at the Main Beam Line using SEM monitor and Keithley 617 electrometer, while TC10 is used as the feedback actuator. The Keithley analog output is sampled by the control-system ADC at a nominal rate of 100 Hz however, the effective feedback bandwidth is limited by the electrometer measurement and trigger/readout dynamics rather than by the ADC sampling rate alone. Therefore, the 100 Hz ADC stream provides time-aligned digital logging and controller input updates, while the physically meaningful beam-current response is constrained by the Keithley 617 readout time.

---

* cslopez@ucdavis.edu
† eprebys@ucdavis.edu

The ADC feedback path, PID controller, TC10 command path, and supervisory GA tuning loop are summarized in Fig. 1.

### *TC10-to-Beam-Current Plant Response*

Before closed-loop controller testing, isolated TC10 current steps were applied to identify the open-loop TC10-to-beam-current response. The normalized transient response, shown in Fig. 2, was well represented by a first-order plus dead-time model. The corresponding response metrics are summarized in Table 1.

The identified response had an effective time constant of approximately 1 s and a measured 10–90% response time of about 2.17 s. This approximately first-order behavior supports the use of a conservative PI/PID feedback controller and provides a practical basis for selecting safe PID gain limits for online beam-current regulation.

Table 1: Measured TC10-to-beam-current Response Metrics from Repeated Clean TC10 moves

| **Response metric** | **Result [s]** |
|---|---|
| First visible response / 10 % point | 0.28 ± 0.03 |
| 50 % response | 0.87 ± 0.10 |
| 90 % response | 2.45 ± 0.27 |
| 10–90 % rise/fall time | 2.17 ± 0.26 |
| First-order equivalent time constant, $\tau$ | ~1.0 |

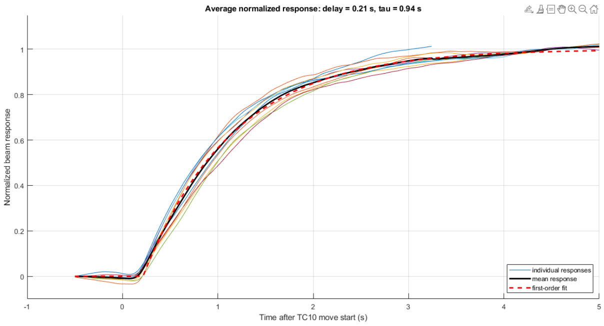

Figure 2: Normalized first-order TC10-to-beam-current response obtained from repeated clean TC10 moves.

### *PID Feedback Control for TC10 Beam-Current Regulation*

The TC10 beam-current regulation is treated as a PID gain-selection problem. as summarized in Fig. 3. For each measured beam-current sample $y_k$, the tracking error relative to the reference $r$ is

$$e_k = r - y_k.$$

The controller response is determined by the gain vector

$$\theta = (K_p, K_i, K_d),$$

which weights the proportional, integral, and derivative contributions to the TC10 correction.

The performance of a PID gain set is quantified using

$$J(\theta) = w_1 J_{\text{track}} + w_2 J_{\text{ss}} + w_3 J_{\text{move}} + w_4 J_{\text{sat}} + w_5 J_{\text{osc}}.$$

The cost function combines tracking accuracy, steady-state regulation, TC10 movement, actuator saturation, and oscillation penalties. Therefore, the desired PID gains are those that minimize beam-current error while keeping the TC10 response smooth, bounded, and stable. The implemented controller also includes a deadband, derivative filtering, anti-windup protection, TC10 limits, and ramp-rate constraints to ensure safe real-time operation.

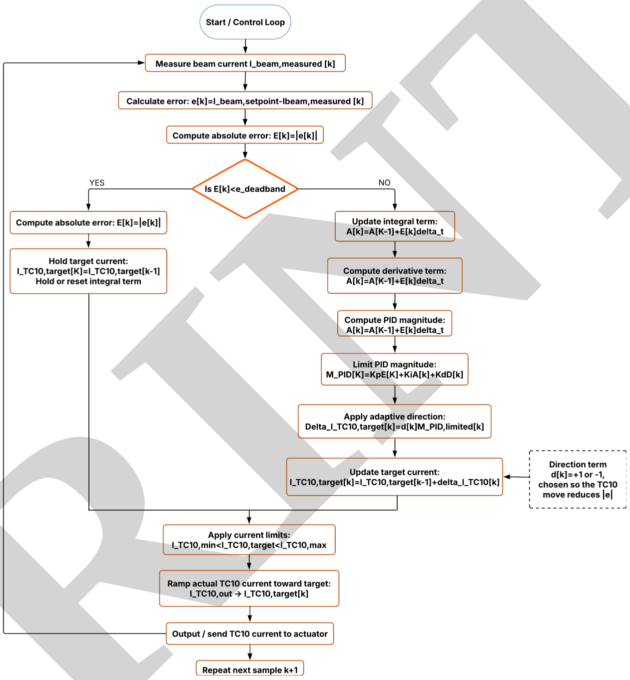

Figure 3: Adaptive-direction PID control algorithm with deadband for TC10 beam-current regulation.

### *Genetic-Algorithm Tuning of PID Gains*

A genetic algorithm (GA) is used as a supervisory tuning layer for the TC10 PID controller. The GA proposes candidate PID gain sets, while the real-time PID loop continues to enforce the deadband, current limits, ramp-rate limits, and other safety constraints. Each chromosome represents one set of $K_p$, $K_i$, and $K_d$ gains within predefined safe ranges. The overall tuning process is summarized in Fig. 4.

Each candidate gain set is evaluated during a short closed-loop test interval using a fitness score based on tracking error, steady-state error, TC10 actuator motion, operation near current limits, and oscillatory behavior. Poor candidates are rejected or penalized, while smoother responses with smaller beam-current error receive better scores. After each generation, selection, crossover, mutation, and elitism are used to form the next population, and the best gain set is selected for the bounded real-time PID controller.

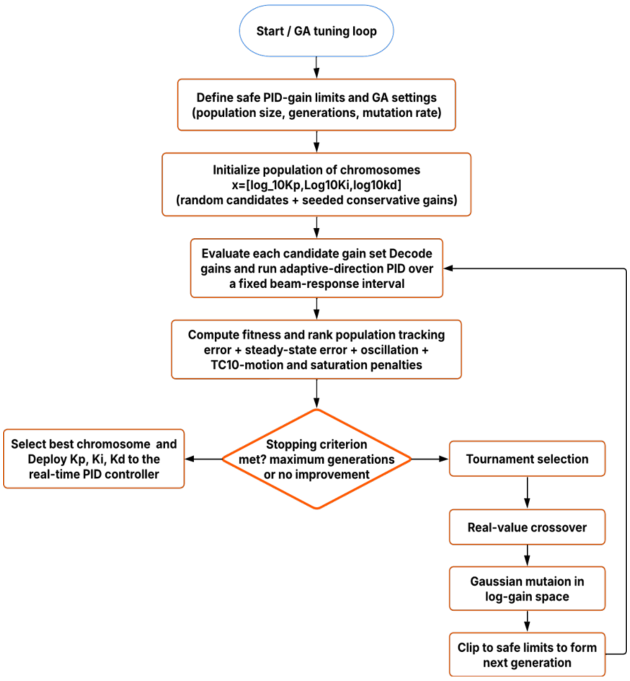

Figure 4: Genetic-algorithm tuning procedure for PID gain selection.

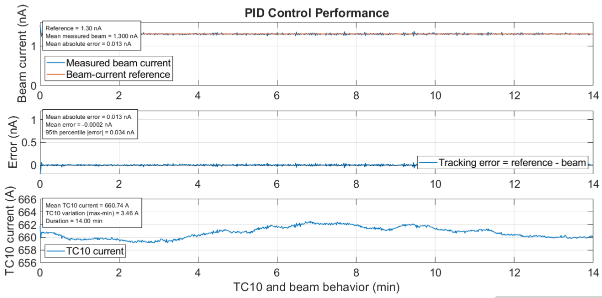

Figure 5: Fourteen-minute closed-loop PID regulation of extracted beam current using TC10 as the actuator. The upper panel shows the measured beam current and the beam-current reference, the middle panel shows the tracking error, and the lower panel shows the TC10 current readback.

## EXPERIMENTAL RESULTS

### *Closed-Loop PID Beam-Current Regulation*

Closed-loop PID beam-current regulation was evaluated using TC10 as the actuator and the extracted beam-current monitor as the feedback signal. Figure 5 shows a representative 14-minute regulation interval at a 1.30 nA beam-current reference.

The measured beam current remained close to the operator-defined setpoint during the test, while the TC10 readback shows only small corrective actuator motion. This demonstrates stable closed-loop regulation without large TC10 excursions or sustained actuator saturation.

### *Genetic-Algorithm Commissioning Results*

Seven supervisory GA tuning runs were completed, corresponding to 168 total chromosome evaluations. Each run used a population size of eight over three generations. The GA consistently converged toward proportional gains near $K_p \approx 0.02$, while the integral and derivative gains varied with the operating condition and beam-current reference. The best recorded fitness value decreased from approximately 0.120 to 0.036 in later runs, indicating improved tracking performance.

Table 2: Summary of GA-based PID Tuning Runs

| **Run** | **Ref. [nA]** | $K_p$ | $K_i$ | $K_d$ | **Cost** |
|---|---|---|---|---|---|
| 1 | 50 | 0.02186 | 0.000114 | 0 | 0.1201 |
| 2 | 50 | 0.02332 | 0.000123 | 0.000096 | 0.1442 |
| 3 | 50 | 0.02234 | 0.000594 | 0.000095 | 0.1359 |
| 4 | 50 | 0.02292 | 0.000149 | 0.000092 | 0.1667 |
| 5 | 50 | 0.02275 | 0.000254 | 0 | 0.1153 |
| 6 | 80 | 0.02052 | 0.000661 | 0.000097 | 0.0438 |
| 7 | 80 | 0.02469 | 0.000356 | 0 | 0.0361 |

Table 2 summarizes the completed online tuning runs. These preliminary results demonstrate that the GA infrastructure was able to apply, evaluate, and reject candidate PID gains during live beam-current regulation. However, the number of completed evaluations was not sufficient to claim full convergence to a globally optimized gain set.

## CONCLUSION

These results demonstrate the initial commissioning of a digital closed-loop beam-current regulation system for the CNL cyclotron. Using TC10 as the actuator, the PID controller maintained the extracted beam current close to the requested reference with small tracking error and limited trim-coil motion.

The GA commissioning tests showed that supervisory optimization can be integrated with the real-time PID layer by applying, evaluating, and rejecting candidate gain sets during online operation. Although the number of GA evaluations was limited, the combined PID and GA results establish a practical foundation for adaptive beam-current regulation and future AI-assisted cyclotron control.

## REFERENCES

[1] “Crocker Nuclear Laboratory”. `http://crocker.ucdavis.edu`

[2] C. L. Osses *et al.*, “Automated control and monitoring system for the Crocker Nuclear Laboratory cyclotron”, in *Proc. IPAC’25*, Taipei, Taiwan, May 2025, pp. 3197–3199. `doi:10.18429/JACoW-IPAC2025-THPS113`

[3] C. L. Osses *et al.*, “AI-ready control infrastructure for cyclotron systems using GPU-accelerated Python GUIs and LabVIEW over ZeroMQ”, in *Proc. NAPAC’25*, Sacramento, CA, USA, Aug. 2025, pp. 42–44. `doi:10.18429/JACoW-NAPAC2025-MOP003`

[4] M. Frey, J. Snuverink, C. Baumgarten, and A. Adelmann, “Matching of turn pattern measurements for cyclotrons using multiobjective optimization”, *Phys. Rev. Accel. Beams*, vol. 22, no. 6, p. 064602, Jun. 2019. doi:10.1103/PhysRevAccelBeams.22.064602

[5] E. D. Hudson, J. A. Martin, M. L. Mallory, F. E. McDaniel, and F. Irwin, “Magnetic field trimming studies for a separated-sector cyclotron”, in *Proc. Cyclotrons’75*, vol. 24, Zürich, Switzerland, Aug. 1975, pp. 201–204, 1975. doi:10.1007/978-3-0348-5520-4_41

[6] T. Duriez, S. L. Brunton, and B. R. Noack, *Machine learning control: taming nonlinear dynamics and turbulence*. Springer, 2017. doi:10.1007/978-3-319-40624-4

[7] J. M. Herrero, X. Blasco, M. Martínez, and C. Ramos, “Optimal pid tuning with genetic algorithms for nonlinear process models”, *IFAC Proceedings Volumes*, vol. 35, no. 1, pp. 101–106, 2002. doi:10.3182/20020721-6-ES-1901.00658

[8] A. Edelen and X. Huang, “Machine Learning for Design and Control of Particle Accelerators: A Look Backward and Forward”, *Annu. Rev. Nucl. Part. Sci.*, vol. 74, no. 1, pp. 557–581, Sep. 2024. doi:10.1146/annurev-nucl-121423-100719

[9] R. Roussel *et al.*, “Bayesian optimization algorithms for accelerator physics”, *Phys. Rev. Accel. Beams*, vol. 27, no. 8, p. 084801, Aug. 2024. doi:10.1103/PhysRevAccelBeams.27.084801