

MODEL-BASED OPTIMISATION FOR AUTOMATED MULTI-TURN EXTRACTION TUNING AT THE PROTON SYNCHROTRON

L. Foldesi*, M. Giovannozzi, A. Huschauer, V. Kain, O. Naumenko, M. Schenk
CERN, Geneva, Switzerland

W. Hillert, Institute of Experimental Physics, Universität Hamburg, Germany

Abstract

The Multi-Turn Extraction (MTE) is a resonance-based technique employed in the CERN Proton Synchrotron (PS) to split the beam into multiple beamlets in the horizontal phase space before extraction to the Super Proton Synchrotron (SPS). The splitting efficiency is measured by the uniformity of the intensities across the beamlets, and optimising it requires fine-tuning several control parameters. The task is particularly challenging due to MTE's sensitivity to multiple interacting variables. In this paper, we investigate the influence of key parameters on MTE efficiency to improve the understanding of their impact on the process. Using a Gaussian Process model and various visualisation techniques, we assess the sensitivity of the MTE efficiency to parameters such as horizontal tune, horizontal excitation strength and frequency, beam intensity, and magnetic hysteresis. The results indicate a complex, non-convex relationship between the MTE performance and these parameters. In addition, external factors, such as magnetic and thermal fluctuations, are potential contributors to performance variability. The findings emphasise the need for a model-based approach to ensure consistent and optimised MTE operation. We propose a solution supported by experimental results.

INTRODUCTION

An essential technique for delivering proton beams from the Super Proton Synchrotron (SPS) to CERN's North Experimental Area for fixed-target physics is the Multi-Turn Extraction (MTE) [1–3] scheme at the Proton Synchrotron (PS). This is a resonance-based technique that enables splitting the beam into multiple smaller beamlets to transfer a uniform intensity beam with the length of five PS turns to the SPS. The beam splitting is achieved by exciting a fourth-order resonance with non-linear magnets, while crossing the resonance by ramping the horizontal tune (Q_h). Figure 1 shows the resulting phase-space topology after crossing the fourth-order resonance. The MTE technique results in the extraction of five beamlets, one corresponding to the beam distribution around the origin of the horizontal phase space and the other four corresponding to the beam trapped in non-linear stable islands [4]. Furthermore, a transverse feedback system (TFB) is used as a time-dependent external exciter that supports trapping the particles in the islands. This horizontal dipolar excitation is applied when Q_h crosses the resonance [3]. The TFB excitation is characterised by its frequency and strength. The detailed dynamics is studied




Figure 1: Example of the horizontal phase-space portrait after crossing the fourth-order resonance. The stable region around the origin is referred to as the *core*, while the non-linear stable structures are called *islands* [4].

in [5]. After the splitting process, the beam is extracted over five consecutive turns and transferred to the SPS, where it is captured, accelerated, and slowly extracted towards the SPS North Area experiments.

The beam intensity should be equally distributed over the five beamlets to provide a continuous and uniform spill to the SPS and the *splitting efficiency* η_{MTE} is the figure-of-merit describing the uniformity:

$$\eta_{\text{MTE}} = \frac{\langle I_{\text{islands}} \rangle}{I_{\text{total}}}, \quad (1)$$

where $\langle I_{\text{islands}} \rangle$ is the average intensity in the islands and I_{total} is the total beam intensity [3]. In operation, each beamlet is targeted to contain one fifth of the total intensity.

The study presented in this paper aims to provide optimal control of η_{MTE} , where the three key control parameters, namely Q_h , TFB excitation frequency (Q_{tfb}), and TFB excitation strength (k_{tfb}) are adjusted in an automated manner. Subsequent sections detail the analysis of parameter interdependencies, as well as the implementation and performance evaluation of a model-based controller.

PARAMETER SENSITIVITY ANALYSIS

A Gaussian Process (GP) model was employed to inspect the dependence of the splitting efficiency on key parameters. We start with a set of initial parameters and vary them within predefined ranges based on operational requirements and limitations. A response GP model $\text{GP}(Q_{\text{tfb}}, Q_h, k_{\text{tfb}}) = \eta_{\text{MTE}}$ is then fitted to the experimental data. For improved visualisation and interpretability, the response is plotted in terms of

* levente.foldesi@cern.ch




Figure 2: Left: Experimental data (circles) and GP interpolation (background) for selected k_{TFB} values, where $\eta_{\text{MTE}} < 0.2$ (blue) and $\eta_{\text{MTE}} > 0.2$ (coral), and $\eta_{\text{MTE}} \approx 0.2$ (black). Centre: Xsuite [6] simulations demonstrating qualitative agreement with experiments. Right: Qualitative illustration of parameter dependencies. Optimal solutions lie on the surface of the “vase-like” structure.

TFB frequency and horizontal tune for various fixed levels of TFB strength. In this way, the problem is projected into two dimensions, as shown in Fig. 2 (left). The circles depict experimental data points, whereas the response from the GP is shown in the background. The colour represents the splitting efficiency, where the black areas define the operational target. There is a good agreement between the colour of the experimental data points and the colour of the response from the GP. The grid search at $k_{\text{TFB}} = 0.16$ (bottom left of the left-hand plots) also confirms the interpolation performed based on the GP, indicating that the surrogate model can accurately capture the interaction of the parameters. It is worth noting that the relationship between Q_{TFB} and Q_h is non-linear and non-convex, highlighting the necessity for a model-based approach. Furthermore, the results indicate that an optimal η_{MTE} only exists if k_{TFB} is large enough. Selecting an appropriate k_{TFB} based on the beam intensity is a simpler task, reducing the modelling problem to two dimensions.

A qualitative illustration of the dependence of η_{MTE} on the three control parameters can be seen in Fig. 2 (right). Note that the plot serves solely as a conceptual illustration of the problem and is not derived from empirical data. Due to the dynamics of the PS, the “vase-like” shape can shift and be subject to distortion on all axes due to magnetic hysteresis, bunch intensity, and other unknown external parameters. Currently, the PS operator manually finds the optimum after a drift or sudden change. This process is challenging and time-consuming. Hence, an automated continuous control approach is proposed that monitors and corrects for performance drifts, with the aim of continuously providing optimum performance.

MODEL-BASED CONTROLLER

Acquiring a measurement of η_{MTE} takes a considerable amount of time, depending on the PS super-cycle configuration, making data collection time-consuming. Furthermore, suboptimal splitting degrades the beam performance of the SPS and fixed-target experiments. The controller must hence operate with high sample efficiency, i.e. it should provide an accurate model and an optimal splitting in as few iterations as possible, to minimise the perturbation to beam operation.

Previous attempts with the extremum seeking algorithm did not fulfil these requirements [7]. This is why a Bayesian Optimisation (BO) framework is employed based on a GP prior to construct a probabilistic surrogate model of the objective function. An acquisition function is then used to identify the most informative points for subsequent evaluations [8].

Based on the analysis discussed in the previous section (see also Fig. 2), one can conclude that an optimal splitting can be achieved as long as k_{TFB} is large enough. The latter is a function mainly of the beam intensity and was determined empirically. The GP model can now be formulated as follows: $\text{BO}(Q_{\text{TFB}}, Q_h; k_{\text{TFB}}) = -|0.2 - \eta_{\text{MTE}}|$, and the goal is to minimise the objective function to achieve optimal splitting.

The prior function is constructed using the mean and covariance of the data points Q_{TFB} and Q_h via GP, with k_{TFB} treated as a fixed parameter during optimisation. The Upper Confidence Bound (UCB) acquisition function was found to provide a good balance between exploration and exploitation. At each iteration, the BO loop proposes a new pair (Q_{TFB}, Q_h) that is applied to the PS control system, resulting in a corresponding η_{MTE} value. The UCB acquisition function combines two components: an exploitation term that maximises the posterior mean and an exploration term weighted by the model’s posterior variance to capture epistemic uncertainty. Figure 3 depicts the described process.

To bootstrap the BO, the model requires at least one data point to construct the surrogate model. This is typically achieved by performing n initial exploration steps, in which the model randomly samples the input space. However, this could cause great volatility in the splitting efficiency. Hence, instead of performing a random search, five predefined points are sampled in the parameter space to ensure sufficient coverage and limit extreme values. Furthermore, super-cycle changes are very disruptive due to effects caused mainly by magnetic hysteresis. Therefore, the collected data are grouped according to the two cycles preceding the MTE beam. Historical data are used as a “warm start” for the model to avoid significant fluctuations. As a result, while the controller is running, the parameters and the respective objective values are collected for a certain super-cycle configuration. Whenever a known configuration occurs, the

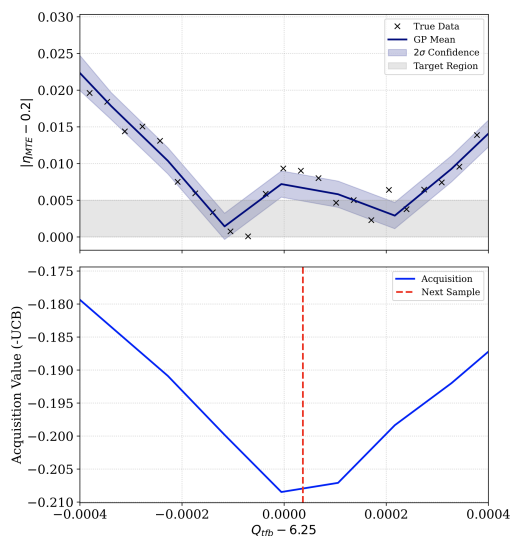


Figure 3: Illustration of the BO building blocks in one dimension. After constructing the initial GP posterior over the objective function with fixed k_{tfb} and Q_h (top), UCB is used to select the next Q_{fb} candidate for evaluation (bottom), and the posterior is updated with the new measurement.

saved data are used to construct the model. This avoids the initial exploration steps, as the GP already has enough data to build a posterior estimate of the objective function.

The approach needs to adapt to the high sensitivity of the PS ring; even small modifications of the parameters can have a considerable effect on the splitting efficiency. This means that once an optimal parameter setting is found by the model, instead of continuously changing the parameters around the optimum, the parameters are frozen. Furthermore, when the controller has found settings that provide near-optimal splitting, UCB with Proximal Biasing [9] is used to prevent large deviations from current settings. From here on, the controller monitors η_{MTE} and intervenes only if the optimum has shifted, which means that it changes machine parameters only when necessary.

RESULTS

Figure 4 (top) shows the result of a seven-hour run during PS beam commissioning on 31st March 2025. The first anchor box shows the model performance with frozen parameters. The variance of η_{MTE} clearly demonstrates the sensitivity of the problem. Note the spikes where the splitting efficiency drops significantly. This is caused by the initial exploration steps, that is, when there are no historical data available to condition the GP. Such abrupt changes and reactions of the controller typically occur when the optimum shifts suddenly due to a super-cycle change. The second anchor box illustrates how the use of historical data for a warm start of the BO routine reduces the number of suboptimal shots from ~ 10 to 1-2. The third box presents a general case where, following a super-cycle change, the model performs five exploration steps before initiating optimisation, finding a solution, and subsequently freezing the parameters.

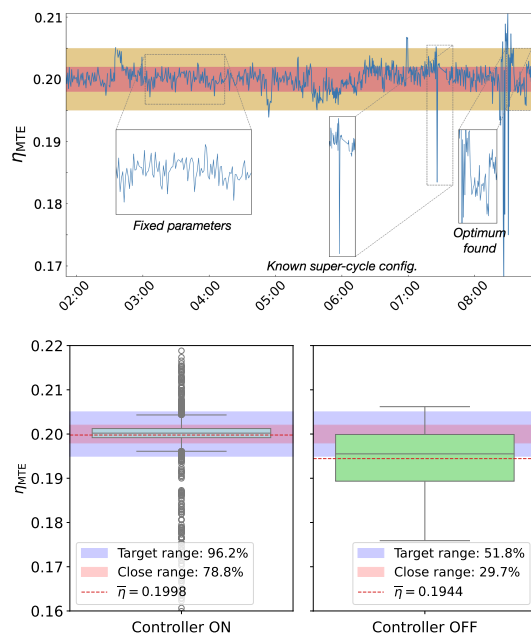


Figure 4: Top: Evolution of η_{MTE} over time. The red area indicates the parameter freeze threshold, while the yellow band defines the operational requirements for the MTE beam. Bottom: Comparison of the splitting performances on an MTE cycle with and without controller, respectively. Data were collected over a 60-hour period.

Figure 4 (bottom) compares the splitting performance in a cycle with the controller running vs. an uncorrected cycle. During the 60-hour test period, 96.2% of the cycles played had a splitting efficiency within requirements, compared to 51.8% for the uncorrected cycle. The increased variance in the controller results from exploration steps that are required when no historical data are available. This diminishes as more super-cycle configurations are encountered.

CONCLUSIONS

This work demonstrates that Gaussian Process models can effectively capture the dependencies of complex physical processes, such as the multi-turn extraction scheme employed at the CERN PS, even when fitted to a limited number of samples. The resulting understanding of the dynamics of the process allowed a reduction in the dimensionality of the control problem by identifying the two key parameters that dominate the behaviour of the system. The proposed Bayesian Optimisation based controller maintained the splitting efficiency within specifications in 96.2% of the cycles, significantly outperforming the uncorrected baseline at 51.8%. Importantly, this approach also addressed limitations observed with numerical optimisers, extremum seeking, and hybrid strategies, such as insufficient sample efficiency, variance due to continuous exploration during stable beam production periods, and disruptive exploration, the latter being particularly problematic due to the resonant nature and sensitivity of the process. As a next step, long-term tests are planned in 2025 before deploying the approach in routine operation.

REFERENCES

- [1] R. Capi and M. Giovannozzi, “Novel Method for Multiturn Extraction: Trapping Charged Particles in Islands of Phase Space”, *Phys. Rev. Lett.*, vol. 88, no. 10, p. 104801, 2002.
doi:10.1103/PhysRevLett.88.104801
- [2] J. Borburgh *et al.*, “First implementation of transversely split proton beams in the CERN Proton Synchrotron for the fixed-target physics programme”, *EPL*, vol. 113, no. 3, 34001. 6 p, 2016.
doi:10.1209/0295-5075/113/34001
- [3] A. Huschauer *et al.*, “Transverse beam splitting made operational: Key features of the multiturn extraction at the CERN Proton Synchrotron”, *Phys. Rev. Accel. Beams*, vol. 20, no. 6, p. 061001, Jun. 2017.
doi:10.1103/PhysRevAccelBeams.20.061001
- [4] A. Huschauer, “Beam Dynamics Studies for High-Intensity Beams in the CERN Proton Synchrotron”, Ph.D. thesis, Technischen Universität Wien, Vienna, Austria, 2016. <https://cds.cern.ch/record/2194332>
- [5] A. Bazzani, F. Capoani, and M. Giovannozzi, “Analysis of double-resonance crossing in adiabatic trapping phenomena for quasi-integrable area-preserving maps with time-dependent excitors”, *Phys. Rev. E*, vol. 109, no. 5, p. 054212, May 2024.
doi:10.1103/PhysRevE.109.054212
- [6] G. Iadarola *et al.*, “Xsuite: an integrated beam physics simulation framework”, *arXiv*, 2023.
doi:10.48550/arXiv.2310.00317
- [7] C. Uden, A. Huschauer, H. Pahl, M. Schenk, N. Madysa, and V. Kain, “Multi-objective extremum seeking to control drifts in the transverse beam splitting efficiency of the Multi-Turn Extraction at the CERN Proton Synchrotron”, in *Proc. IPAC’23*, Venice, Italy, 2023, pp. 1675–1678.
doi:10.18429/JACoW-IPAC2023-TUPA159
- [8] P. I. Frazier, “A Tutorial on Bayesian Optimization”, *arXiv*, 2018. doi:10.48550/arXiv.1807.02811
- [9] R. Roussel and A. Edelen, “Proximal Biasing for Bayesian Optimization and Characterization of Physical Systems.”, presented at NeurIPS 2021, 2021, unpublished. https://ml4physicalsciences.github.io/2021/files/NeurIPS_ML4PS_2021_82.pdf