IPAC'26 Preliminary proceedings (edited version): MOP6001

ELIMINATING MAINS NOISE EFFECTS IN ACCELERATORS WITH MACHINE LEARNING

V. Kain[*, 1], P. A. Arrutia Sota[1], H. Bartosik[1], F. F. Ferber[2], F. Follin[1], C.-M. Genton[1],
Q. King[1], B. Rodriguez Mateos[1], O. Michels[1], M. Schenk[1]
[1]CERN, Geneva, Switzerland
[2]University of Salzburg, Salzburg, Austria

Abstract

Power supply ripple at various frequencies — characteristic to the magnet circuits or from the electrical network — has always been an issue for accelerator operation. In the CERN Super Proton Synchrotron (SPS) this is particularly detrimental at low energy for LHC-type beams and at top energy for slowly extracted beams. This contribution summarises the efforts in the CERN SPS over the last years to compensate the ripple at 50 Hz and its harmonics in the main quadrupole circuits using Machine Learning methods. The implementation required hardware modifications on the power converter electronics side, additions to the controls infrastructure and the development of adaptive algorithms that can deal with changes in the electrical distribution network throughout the day. Continuous control with tailored adaptive Bayesian Optimisation has been implemented for slow extracted spill control throughout 2024 and 2025. The improved spill quality obtained over the years is discussed, while results from R&D towards one-shot correction algorithms for beams that are only played on-demand are also briefly summarised.

INTRODUCTION

The CERN Super Proton Synchrotron (SPS) is a 7 km circumference synchrotron that serves as injector to the Large Hadron Collider (LHC) and provides beams for fixed-target physics in the North Area (NA) via chromatic slow extraction (SX) in the horizontal plane, exploiting the 2/3 resonance. Mains power supply ripple at 50 Hz and its harmonics is a significant source of perturbation in both operational modes.

For SX beams, spill quality can be degraded by several mechanisms such as macro-spill shape distortions from magnetic hysteresis or very high-frequency intensity ripple arising from slow RF debunching. The focus of this work is the time-varying ripple at $n \times 50$ Hz originating from the SPS power supplies, whose phase and amplitude drift on timescales of minutes to hours due to changes in the electrical distribution network. The operational goal is to keep normalised spill modulation amplitudes below 0.15 for more than 85% of the spill duration. For LHC-type beams, experiments with single-bunch proton beams in 2017–2018 established that 50 Hz tune ripple from the quadrupole power converters causes emittance blow-up and particle losses [1]. This effect was confirmed to persist for multi-bunch beams in dedicated Machine Development sessions in 2024, making 50 Hz compensation a priority also for LHC beams.

CORRECTION SCHEME

Hardware

The correction scheme, conceived in 2018, derives a reference signal from the 50 Hz mains voltage and injects sinusoidal corrections at 50 Hz and its harmonics into the voltage reference of the quadrupole and main dipole power converters. The gain and phase of each harmonic component are free parameters to be optimised. The FGC Class 63 digital controller provides the infrastructure for harmonic injection. One sinusoidal generator per harmonic order runs phase-locked to the 50 Hz network reference. Their outputs, scaled and shifted by the gain and phase parameters, are added to the current loop voltage reference. This superposes a time-varying correction onto the nominal magnet excitation without affecting the main regulation loop.

Measurement

The primary feedback signal in case of fixed target physics is the SPS beam intensity spill monitor (BSI), a high-bandwidth (~100 kHz) detector measuring the extracted particle flux. Its time-domain signal is Fourier-transformed in real time to extract the normalised amplitudes at 50 Hz and its harmonics, which serve as the objective for the optimisation algorithm. DC current transformers (DCCTs) on the QF, QD and main dipole magnet bus bars provide measurements of the current ripple in the corresponding circuit, enabling model-based and direct DCCT-feedback approaches for all SPS beam types, i.e. also in the case of LHC beams.

CONTROL ALGORITHM

Adaptive Bayesian Optimisation

Since 2023, Adaptive Bayesian Optimisation (ABO) provides 24/7 continuous spill quality control [2] in the SPS. Inspired by work on Gaussian Process (GP) kernels for pattern discovery and extrapolation [3,4], the algorithm builds a GP surrogate model that incorporates time t as an additional input dimension. This allows the model to predict the objective at the next time step $t + 1$ given observations up to time t, enabling continuous closed-loop control. The objective for ABO and continuous spill control was to minimise the normalised amplitude at $n \times 50$ Hz measured by the BSI.

Initially, only the Upper Confidence Bound (UCB) acquisition function was used with fixed hyperparameters tuned

---

[*] verena.kain@cern.ch

individually per controller for the different harmonics. For details of the original implementation see [2].

ML Infrastructure

The ABO-based spill control for 50 and 100 Hz run continuously as background processes via the UCAP (Unified Controls Acquisition and Processing) framework [5], which hosts “Virtual Device Service” servers capable of executing Python transformations on data streams and performing synchronous event building tied to the SPS cycle. GeOFF (Generic Optimisation Framework and Frontend [6]) deployed on UCAP with GPU acceleration provides the optimisation engine (acc-geoff4ucap). This architecture allows GUI-less 24/7 operation without operator intervention.

RESULTS 2024

Spill Quality for Slow Extracted Beams

ABO operated in combination with Empty Bucket Channeling (EBC) [7], a higher-harmonic longitudinal RF technique that increases phase-space mixing through chromaticity and thereby reduces the broadband component of spill noise. By suppressing broadband background, EBC creates the margin that allows ABO to keep the $n \times 50$ Hz amplitudes below the operational threshold also in case of fast changes that ABO will only slowly recover. Figure 1 shows the cumulative distribution function (CDF) of normalised 50 Hz and 100 Hz spill amplitudes for the entire operational year of 2024. The 50 Hz performance is only slightly above the spill quality cumulative target of 85% below 0.15 normalised amplitude due to an accidentally wrong configuration of the 50 Hz controller for most of the operational year, which resulted in outdated feedback signals in the controller.

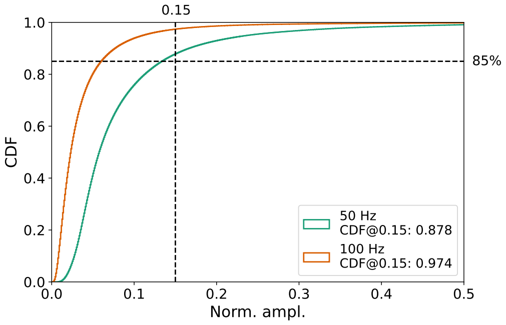

Figure 1: CDFs of normalised 50 Hz and 100 Hz spill amplitudes for the 2024 proton run.

Impact on LHC Beams

A clear reduction in flat-bottom loss rates during flat-bottom storage of 72 LHC proton bunches was observed with simply applying the fixed target correction settings, directly demonstrating the operational benefit of 50 Hz compensation even for LHC-type beams. 50 Hz correction became therefore standard for all high-intensity LHC machine development sessions from then on.

During LHC ion filling, which requires approximately one hour of LHC ion cycles in the SPS supercycle, ABO-based controllers were deployed to track the slow drift of the 50 Hz amplitude and phase in the QF and QD circuits and maintain the current ripple at acceptable levels. This yielded a 15–20% increase in SPS ion intensity transmission compared to operation without correction [8]. The 2024 ion run resulted in record LHC ion intensities per bunch, significantly surpassing the high luminosity LHC targets and increasing bunch intensities by 20-30% compared to 2023, to which the 50 Hz compensation significantly contributed [9].

TOWARDS IMPROVED ALGORITHMS

Limitations of Naive ABO

The 2023–2024 ABO implementation with fixed hyperparameters, kernel structure, and acquisition function was adequate for fixed-target physics but insufficient for applications requiring tighter control of exploration, such as LHC beam quality maintenance. Occasional large steps driven by UCB exploration can momentarily degrade the objective function, which is acceptable for NA beams but not during LHC filling. Two complementary approaches have therefore been developed for LHC beams, where the problem can be simplified to direct DCCT correction with constant phase.

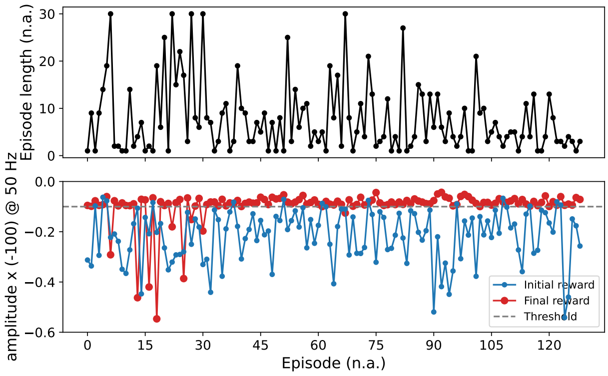

Figure 2: Training of Soft Actor-Critic (SAC) RL algorithm to control the 50 Hz power converter ripple of the SPS QD circuit. The upper plot shows the evolution of the episode length and the lower plot initial and final 50 Hz amplitudes. The threshold is −0.1 corresponding to 0.001 A ripple.

Reinforcement Learning

For DCCT-based correction on a single harmonic with known phase, the state and action spaces for reinforcement learning (RL) simplify to $\vec{s} = [A_{\mathrm{DCCT}}, A_{\mathrm{corr}}]$ and $\vec{a} = [\Delta A_{\mathrm{corr}}]$, where A_{DCCT} is the current ripple amplitude and A_{corr} the total correction setting for harmonic injection.

A Soft Actor-Critic (SAC) [10] agent was trained online on the SPS machine, exploiting the fact that each quadrupole station houses two sub-converters per QD and QF circuit: one was used to inject randomised noise, providing diverse initial states for efficient episode-based training, while the other applied the compensation. After 1300 iterations of overnight online training, the agent achieved robust closed-loop control, see Fig. 2. (Note: the limited action step size used during the setup meant that correction of large initial ripples required a few iterations; this is acceptable in operational use as the ripple evolves slowly.)

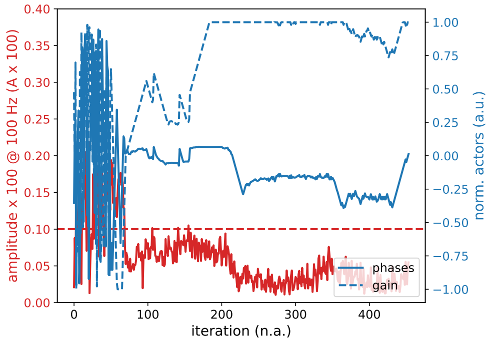

Figure 3: Running the 2025 ABO algorithm for 100 Hz control with initial random exploration phase to build the first GP model. The red dashed line indicates the limit above which beam quality degradation is expected.

An alternative model-based approach that models the one-step system dynamics as a Gaussian Process (GP) and solves a model predictive control (MPC) problem over a finite horizon, selecting the action sequence that maximises predicted cumulative reward while accounting for GP uncertainty, was also tested [11]. The GP is updated continuously using a circular conditioning buffer. This approach is highly sample-efficient, achieving performance comparable to the SAC agent with only 300 total iterations, making it attractive for deployment in scenarios with limited beam time.

Robust Adaptive Bayesian Optimisation

While the tests with RL gave promising results, robustness studies and behaviour under strict constraints could not be evaluated during the short 2025 SPS commissioning period. Instead an improved universal ABO algorithm was put in place for 2025. The new “Adaptive^2 BO” algorithm was commissioned in early 2025, initially targeting only the LHC beam quality use case. The algorithm makes both the controller and the GP algorithm itself state-dependent, with three adaptive mechanisms triggered by the current model uncertainty and the trend of recent objective values:

- **Action clipping:** smaller exploration steps when the system is stable and well-corrected; larger steps when the objective is deteriorating.
- **Acquisition function switching:** PosteriorMean for exploitation in stable conditions; UCB to recover from deterioration.
- **Buffer truncation:** when the objective deteriorates rapidly, only the most recent 25% of the circular buffer is used to condition the time-dependent GP, discarding stale observations that no longer reflect current network conditions.

Test runs demonstrated that the algorithm stabilises performance and converges sufficiently fast under major changes in $n \times 50$ Hz ripple in the main SPS power converters. Figure 3 shows the evolution of 100 Hz (gain and phase) correction on QF after the initial random exploration phase needed to build the first GP model, with the acquisition function switch from UCB to PosteriorMean visible as the system locks onto the optimum. This algorithm became the work horse for 2025 across all beam types; see Figs. 4 and 5 for the 2025 slow extracted spill control performance.

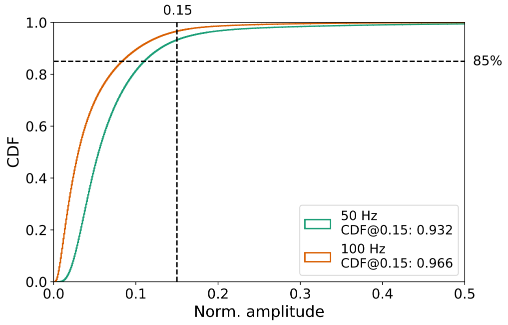

Figure 4: CDFs of normalised 50 Hz and 100 Hz spill amplitudes for the 2025 proton run. The overall performance was well above target, despite running without EBC for weeks 25-33, see Fig. 5.

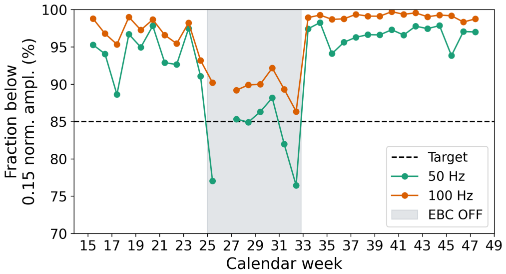

Figure 5: Weekly CDFs for normalised 50 Hz and 100 Hz spill amplitudes for the 2025 proton run.

CONCLUSION

Mains 50 Hz noise has been identified as a limiting factor for SPS slow-extraction spill quality and, recently, also for LHC beam quality and ion transmission performance in the SPS. A combination of power converter hardware modifications, dedicated controls infrastructure, and tailored adaptive ML algorithms now enables 24/7 compensation of $n \times 50$ Hz ripple for slow extracted beams. The 2024 fixed target proton run met the operational spill quality target, and record LHC ion intensities were achieved with 50 Hz compensation. New algorithms – SAC-based RL, GP-MPC, and Robust ABO – were successfully tested as potential improvements, particularly for the more demanding LHC beam quality applications. Owing to its reduced complexity and flexibility without lengthy retraining, the improved ABO algorithm was selected for 2025 operation, achieving spill quality significantly surpassing the target for both 50 and 100 Hz ripples.

REFERENCES

[1] H. Bartosik, "Challenges in Understanding Space Charge Dynamics", in *Proc. HB'18*, Daejeon, Korea, Jun. 2018. `doi:10.18429/JACoW-HB2018-MOA1PL01`

[2] V. Kain *et al.*, "Slow extracted spill ripple control in the CERN SPS using adaptive Bayesian optimisation", in *Proc. IPAC'24*, Nashville, TN, USA, May 2024, pp. 1790–1793. `doi:10.18429/JACoW-IPAC2024-TUPS55`

[3] A. G. Wilson and R. P. Adams, "Gaussian Process Kernels for Pattern Discovery and Extrapolation," in *Proc. 30th Int. Conf. Machine Learning (ICML'13)*, Atlanta, GA, USA, Jun. 2013, pp. 1067–1075. `doi:10.48550/arXiv.1302.4245`

[4] N. Kuklev *et al.*, "Robust Adaptive Bayesian Optimization," in *Proc. IPAC'23*, Venice, Italy, May 2023, pp. 4428–4431. `doi:10.18429/JACoW-IPAC2023-THPL007`

[5] L. Cseppentő and M. Büttner, "UCAP: A Framework for Accelerator Controls Data Processing @ CERN," in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 230–235. `doi:10.18429/JACoW-ICALEPCS2021-MOPV039`

[6] P. Madysa *et al.*, "Geoff: The generic optimization framework & frontend for particle accelerator controls," *SoftwareX*, vol. 32, p. 102335, Dec. 2025. `doi:10.1016/j.softx.2025.102335`

[7] P. A. Arrutia Sota *et al.*, "Empty-bucket techniques for spill-quality improvement at the CERN Super Proton Synchrotron," *Phys. Rev. Accel. Beams*, vol. 27, p. 074001, 2024. `doi:10.1103/PhysRevAccelBeams.27.074001`

[8] R. Bruce *et al.*, "Exceeding High-Luminosity LHC Performance Targets During the 2024 Pb-Pb Ion Run," in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 899–902. `doi:10.18429/JACoW-IPAC2025-TUBN3`

[9] E. Waagaard *et al.*, "Investigating Lead Ion Beam Loss Mechanisms at the Super Proton Synchrotron Flat Bottom," *J. Phys.: Conf. Ser.*, vol. 3094, p. 012034, 2025. `doi:10.1088/1742-6596/3094/1/012034`

[10] T. Haarnoja *et al.*, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," in *Proc. 35th Int. Conf. Machine Learning (ICML)*, Stockholm, Sweden, Jul. 2018, pp. 1861–1870. `doi:10.48550/arXiv.1801.01290`

[11] J. Kocijan *et al.*, "Gaussian Process Model Based Predictive Control," in *Proc. American Control Conference (ACC'04)*, Boston, MA, USA, Jun. 2004, pp. 2214–2219. `doi:10.23919/ACC.2004.1383790`