14th International Particle Accelerator Conference,Venice, Italy JACoW Publishing
ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-MOPA137

EXTREMUM SEEKING FOR ACCELERATOR OPTIMISATION*

B. Veglia†, I. Agapov, T. Hellert, J. Keil, L. Malina, E. Musa,
Deutsches Elektronen-Synchrotron DESY, Germany,
L. Hoummi, S.M. Liuzzo, N. Carmignani, L.R. Carver, S. White,
ESRF, Grenoble, France

Abstract

A new collaboration between ESRF and DESY is aiming at building tools and concepts that can be used for the next generation light sources. The developed tools will be applied to the ESRF-EBS and PETRA IV to validate concepts to reduce the natural horizontal emittance or improve the injection efficiency or the lifetime of storage rings. In this project framework, the bounded Extremum Seeker (ES) algorithm has been studied as a Touschek lifetime optimisation procedure. This contribution presents the experiment performed on the ESRF-EBS electron beam where several skew quadrupole knobs were tuned at the same time for vertical emittance minimisation as an initial test.

INTRODUCTION

The ESRF-EBS storage ring [1] Touschek lifetime is currently optimised routinely using an empirical approach [2]. The development of a new tool to perform a quicker and systematic maximisation would be of great benefit for operations not only at ESRF but possibly at other synchrotron radiation sources. Extremum seeking [3], with its ability to handle open-loop unstable, time-varying, nonlinear systems, is an ideal candidate for online control and optimisation of complex systems controlled by several parameters. Since accelerators are very large and complex systems with many components and time-varying beam distributions, such techniques have been deployed in few particle accelerator applications. Among others, it was used as an adaptive online model tuning for non-invasive electron beam diagnostics at the Facility for Advanced Accelerator Experimental Tests (FACET) at SLAC National Accelerator Laboratory [4] and for high voltage converter modulator optimisation at Los Alamos Neutron Science Center (LANSCE) [5]. Because it is model-independent and can tune multiple parameters simultaneously we adopted an extremum seeking algorithm [6, 7] as a first and easy-to-implement approach. Such algorithm minimises (or maximises) a user-defined objective function of an unknown system with dithered input signals towards optimal control directions. In addition to its adaptivity and robustness to measurement noise, the most promising feature of the extremum seeking approach is that no explicit model of the system is required, but only reliable measurements of the quantity to be optimised. This paper reports about the initial simulations and the preliminary results obtained during a dedicated experimental shift at the ESRF-EBS storage ring.

* This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 871072

† bianca.veglia@desy.de

ALGORITHM

The Extremum Seeker (ES) is a local, model-independent algorithm. The goal is to minimise (or maximise) a time varying, user-defined cost function which may be analytically unknown $C(\mathbf{p}, t) \in \mathbb{R}$. The cost function dependency to the system is unknown, but changes its value according to the change in n-parameters $\mathbf{p} = (p_1, \dots, p_n)$, where p_i can be any of the accelerator input such as RF voltages or magnets strengths at each time t. All of the parameters are initialized with some settings $p(1)$ based on models or experience. The algorithm adjusts the parameters according to the dynamics:

$$\frac{dp_i(t)}{dt} = \sqrt{\alpha \omega_i} \cos\left(\omega_i t + k\hat{C}(\mathbf{p}(t), t)\right) \quad (1)$$

based on possibly noise corrupted measurement of the cost function $\hat{C}(\mathbf{p}, t) = C(\mathbf{p}, t) + \nu(t)$, being $\nu(t)$ the time varying machine noise. The parameter $\alpha > 0$ is the dithering amplitude and the term $k > 0$ acts as feedback gain. The dithering frequencies ω_i must be distinct so that the perturbing functions are orthogonal in the frequency domain and large enough to dominate all of the natural time variation in the dynamics. The discrete, iterative feedback parameter update law is then implemented as:

$$p_j(n+1) = p_j(n) + \Delta\sqrt{\alpha \omega_i} \cos\left(\omega_i \Delta n + k\hat{C}(n)\right) \quad (2)$$

In this way each new parameter setting is based only on the previous parameter setting and the previous cost function measurement. We must choose the ES time step Δ small enough relative to the frequencies ω_i. Typically, Δ may be chosen as: $\Delta = \frac{2\pi}{10\times 1.75}$ so that, in the absence of other dynamics, at least ten times steps are required to complete one oscillation of parameter settings. Before starting the iterations upper and lower bounds are defined for a for all of the parameters that are being tuned. This is especially useful because these bounds can be used to normalize the parameters to all live within a common range, such as [-1,1]. The algorithm was adapted from the existing version implemented in the Ocelot simulation toolkit [8, 9] by Alexander Scheinker.

SIMULATIONS

The ES optimiser was tested with the Python interface of the Accelerator Toolbox (pyAT) module [10] for the EBS storage ring [1]. The chosen target of the optimisation is the vertical emittance, to be minimised by tuning skew quadrupole knobs. Two distinct cases were tested:

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-MOPA137

1. In the first case the vertical emittance was increased intentionally by setting a random value for the strength of the skew quadrupole component of a sextupole of the SF2A family;
2. In the second case all skew quadrupoles were set to random values.

The ES algorithm was used in both cases to minimise the vertical emittance calculated with the pyAT program. The controlled parameters used for the minimisation were 32 skew quadrupoles of the same SF2A family including the detuned one.

In the first case the algorithm would ideally identify the single magnet that was moved and restore the nominal conditions. In the second case there was no particular expectation on the pattern of corrections. Figure 1 shows the result of these two simulated optimisation scenarios. In both cases the ES optimiser finds an improved solution for the vertical emittance.

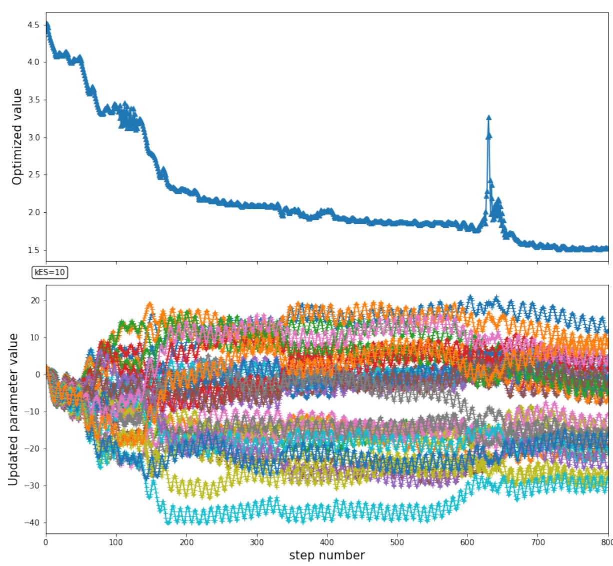

Figure 1: Only one skew quad of the SF2A in the pyAT simulated lattice is excited to induce an increase on the vertical emittance (optimised value, top) which is minimised by the ES algorithm tuning 32 SF2A knobs (updated parameter value, bottom). Each colour corresponds to a different magnet.

The values of the dithering amplitude α and the feedback gain k are non trivial to select but strongly affect the ability of the algorithm to converge or even to trigger the minimisation to start. The optimal values need to be identified through empirical observations. In the simulations the optimisation operated with $\alpha = \alpha_0 = 0.1$ while k was slowly increased from 1 to 10 in order to finally observe a satisfactory convergence towards a minimum. Larger values showed slower or no convergence at all. Additional tests in preparation for studies on the actual accelerator were performed also using the EBS control system simulator [11].

ES STUDIES AT ESRF

Following the promising simulations, the ES algorithm was tested on the ESRF-EBS storage ring during a four-hour machine dedicated time (MDT) on November 2022. As proof of principle and to work on a low-current beam (set at 4.60 mA), the ES was adapted to the vertical emittance minimisation tested in simulations rather than lifetime optimisations. In order to ensure that the magnets reached the set-point value calculated by the algorithm and that the read value of the vertical emittance has stabilized, a waiting period of 5 seconds is introduced. The emittances are calculated at two pinholes in the storage ring, for the studies only measurement from ID07 are considered. The ESRF-EBS electron beam has a low equilibrium vertical emittance of 0.5 ± 0.1 pm rad. The Extremum Seeking algorithm was run through a Jupyter notebook in the ESRF-EBS control room. The MDT started with the introduction of a $+0.005\ \mathrm{m}^{-1}$ correction strength in one skew, located at a randomly chosen sextupole of the SF2A family (see Fig. 2), which increased the vertical emittance to 3.9 ± 0.1 pm rad. The minimisation will be conducted using all the 32 knobs acting on as many skew quadrupoles located at the SF2A sextupoles. The parameter updates were carried out on normalized values that were then unnormalized to physical set-points that could be sent to the EBS controls.

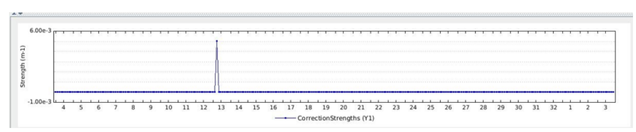

Figure 2: Skew quadrupole correction strengths at the beginning of the MDT. The application shows the difference between the current strengths and the nominal strengths.

The first minimisation is launched with the hyper-parameters $\alpha = 3 * \alpha_0$ and $k = 40$ which provided a good response in simulations. The result of the 87 iterations as well as the sent amplitudes are displayed in Fig. 3. Two limitations are spotted: the amplitudes reach values too large, and the first observed reduction from 3.9 to 2.9 pm rad doesn't seem steep enough to ensure a minimisation and after the first 50 steps starts growing again blowing up to almost the double of the initial value. Therefore, the parameter boundaries were decreased from [-8,8] to [-5,5], and k decreased to 30 to limit the impact of the read emittance value in the perturbation function.

The applied changes are followed by small variations in the vertical emittance around the starting point. The minimisation was stopped when the emittance blew up with increased amplitudes of the skew quadrupole strengths. The first step was not successful: first, the very localized error on one skew, although simple to correct, still triggers a global response of the selected knobs. Instead of perturbing one skew, random errors were sent to all skew quadrupoles, limited to 1% in strength change, which generated a vertical emittance of 5.7 pm rad. The next ES run used the same

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-MOPA137

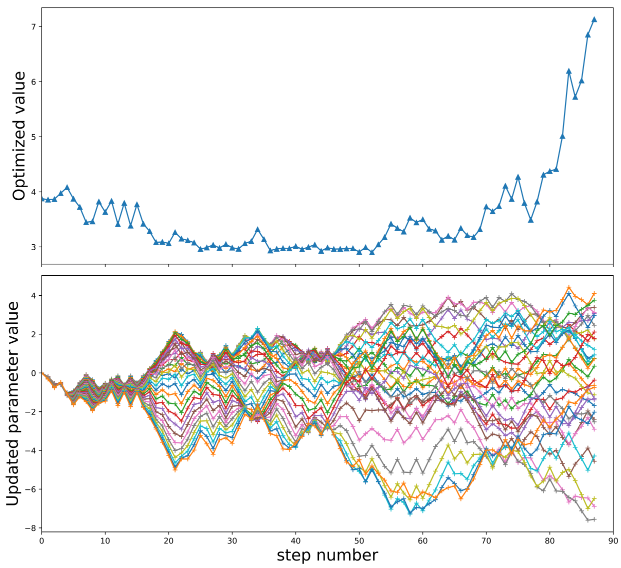

Figure 3: Evolution of the vertical emittance (top) and knob amplitudes (bottom) sent to the skew quadrupoles on the SF2A during the first attempt of the Extremum Seeking sessions.

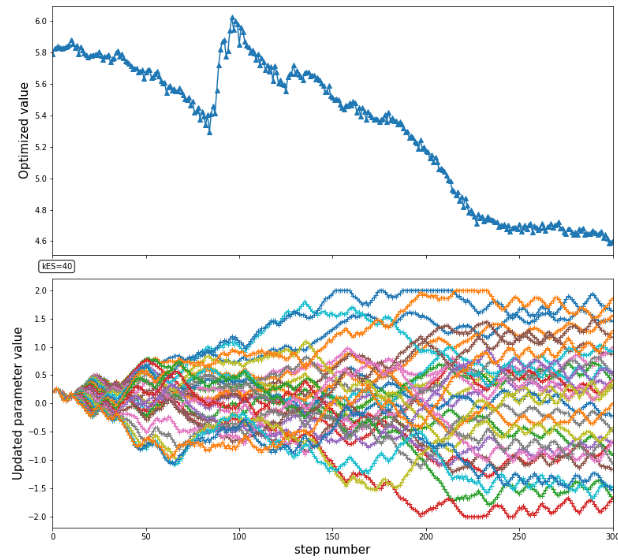

Figure 4: vertical emittance (optimised value, top) which is minimised by the ES algorithm tuning 32 SF2A knobs (updated parameter value, bottom).

hyper-parameters as the latest run, and the same knobs. The variation in vertical emittance were of the 10^{-3} order, so the parameter α was increased to $\alpha = 6 * \alpha_0$. Despite presenting small and promising variations in the vertical emittance, this minimisation ended with large amplitudes in the skew quadrupoles exceeding the input boundaries, and a blow up of the emittance from 5 pm rad to about 30 pm rad. This phenomenon is to be avoided in a future version of the algorithm, implementing for instance boundaries on the maximum emittance allowed with the possibility to reset the parameters to a previously found optimum.

A last test with drastically reduced boundaries on the parameter amplitudes to [-1,1] and increased k to 40 (hoping to start observing the minimisation sooner), and a reduced $\alpha = 2 * \alpha_0$ to slow down the amplitude variations. The previous skew quadrupole corrections are replaced with new small (less than 1% variation) random strengths, for a vertical emittance of 5.8 pm rad. The evolution of the amplitudes and vertical emittance during this minimisation is shown in Fig. 4. The boundaries were respected and triggered a stronger minimisation after the 200 steps. The minimisation ended after reaching its maximum iteration number which resulted in a total emittance reduction of 1.2 pm rad.

CONCLUSION

A vertical emittance minimisation using the ES algorithm on the ESRF-EBS storage ring was tested. Firstly, a long time is required to select the optimal hyper-parameters α and k of the ES, and due to the limited availability of allocated MDTs, a simple scan is not efficient enough. Additionally, the optimal hyper-parameters identified during simulations didn't produce the same outcome when implemented in the control room. Future studies should aim at reducing the time allocated for the hyper-parameter selection. Secondly, the MDT outcome was a 20% reduction of the emittance, achieved after a 30-minute run (corresponding to 300 iterations) with no convergence of the algorithm. To make the algorithm more robust and avoid stagnation to local minimum, a tool evaluating the convergence during the minimisation could be implemented. This additional module would be required to:

- Identify a possible local minimum;
- Modify the hyper-parameter to increase the amplitude of the perturbation and scan a larger scale;
- Increment k to increase the impact of the cost function on the parameter update.

In this way the convergence could be sped up and make the ES a more time-efficient algorithm for online optimisation. Adapted the aforesaid specifications to the algorithm, the tests could be extended to the final goal of improving the Touschek life-time. The results would then be compared to the current lifetime optimisation conducted on the ESRF-EBS after each shutdown of the accelerators, considering efficiency, time required and stability. Other techniques are also being considered, such as applications of the simplex algorithm as well as Bayesian optimization methods.

REFERENCES

[1] S. M. White *et al.*, "Commissioning and Restart of ESRF-EBS", in *Proc. IPAC'21*, Campinas, Brazil, May 2021, pp. 1–6. doi:10.18429/JACoW-IPAC2021-MOXA01

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-MOPA137

[2] N. Carmignani *et al.*, “Online Optimization of the ESRF-EBS Storage Ring Lifetime”, in *Proc. IPAC’22*, Bangkok, Thailand, Jun. 2022, pp. 2552–2555. doi:10.18429/JACoW-IPAC2022-THPOPT001

[3] E. Lavretsky, N. Hovakimyan, and A. Calise, “Adaptive extremum seeking control design,” in *Proc. of the 2003 American Control Conference*, vol. 1, 2003, pp. 567–572. doi:10.1109/ACC.2003.1239079

[4] A. Scheinker, X. Huang, and J. Wu, “Minimization of Betatron Oscillations of Electron Beam Injected Into a Time-Varying Lattice via Extremum Seeking,” *IEEE Trans. Control Syst. Technol.*, vol. 26, pp. 336–343, 2018. doi:10.1109/TCST.2017.2664728

[5] A. Scheinker and M. J. Bland, “High Voltage Converter Modulator Optimization”, in *Proc. IPAC’13*, Shanghai, China, May 2013, paper MOPWA074, pp. 852–854.

[6] A. Scheinker and D. Scheinker, “Bounded extremum seeking with discontinuous dithers,” *Automatica*, vol. 69, pp. 250–257, 2016. doi:10.1016/j.automatica.2016.02.023.

[7] A. Scheinker and D. Scheinker, “Extremum seeking for optimal control problems with unknown time-varying systems and unknown objective functions,” *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, p. 1143–1161, 2021. doi:10.1002/acs.3097

[8] I. Agapov, G. Geloni, S. Tomin, and I. Zagorodnov, “OCELOT: A software framework for synchrotron light source and FEL studies,” *Nucl. Instrum. Meth. A*, vol. 768, pp. 151–156, 2014. doi:10.1016/j.nima.2014.09.057

[9] A. Scheinker, “Ocelot Optimizer ES method”, 2019. https://github.com/ocelot-collab/optimizer/tree/master/op_methods

[10] B. Nash *et al.*, “New Functionality for Beam Dynamics in Accelerator Toolbox (AT)”, in *Proc. IPAC’15*, Richmond, VA, USA, May 2015, pp. 113–116. doi:10.18429/JACoW-IPAC2015-MOPWA014

[11] S. Liuzzo, L. Carver, J. Chaize, L. Farvacque, A. Götz, D. Lacoste, N. Leclercq, F. Poncet, E. Taurel, and S. White, “The ESRF-EBS Simulator: A Commissioning Booster”, in *Proc. ICALEPCS’21*, Shanghai, China, Oct. 2021, pp. 132–137. doi:10.18429/JACoW-ICALEPCS2021-MOPV012