PARASITIC OPTIMIZATION OF THE TRANSFER BEAMLINE EFFICIENCY AT ELSA

S. Witt*, K. Desch, D. Elsner, D. Proft, Physics Institute, University of Bonn, Germany

Abstract

The 3.2 GeV electron accelerator ELSA in Bonn consists of three acceleration stages each interconnected by tunable transfer beamlines. The steering of the electron beam through the transfer line from linear accelerator to the Booster Synchrotron is currently adjusted by hand, which limits a systematic improvement of the transfer efficiency. An automated optimization using the "simulated annealing" technique has been developed and integrated into the control system to improve the situation. It allows for a continuous optimization without interfering with usual beamtime for experiments by utilizing the 6s off-time in between injections into the stretcher ring. In a simulation using the actual accelerator's settings as starting parameters, transmission rates have been increased significantly. The methods and results with the accelerator hardware are presented.

ELECTRON STRETCHER FACILITY ELSA
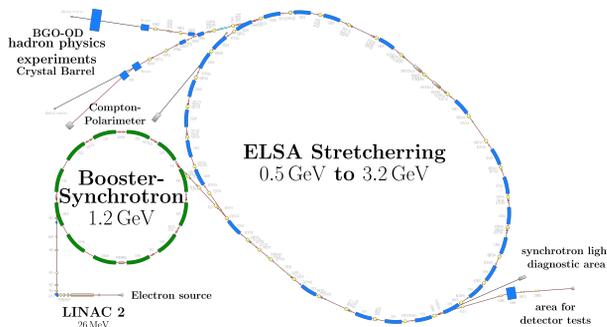
Figure 1: The Electron Stretcher Facility ELSA.

The electron stretcher facility ELSA in Bonn (Fig. 1) is a three stage electron accelerator capable of accelerating electrons to energies up to 3.2 GeV. The acceleration process is cyclical: The electron gun, linear accelerator LINAC2 and Booster Synchrotron are triggered with 50 Hz during the injection phase into the stretcher ring, which takes around 0.5 s. Then, those first two acceleration stages idle while the stretcher ring ramps up the electron energy and extracts the electrons to one of the experimentation sites. This phase typically takes 6 s, before the cycle starts over [1].

The transfer beamline between LINAC2 and the synchrotron (Fig. 2) guides the electron beam via 8 quadrupoles, 7 corrector magnets for both planes and 2 dipole magnets, all of which are controllable from the control system software, leading to a 24-dimensional parameter space. These parameters can be optimized to improve the transfer efficiency of the beamline.

---

* sebastian-witt@outlook.de
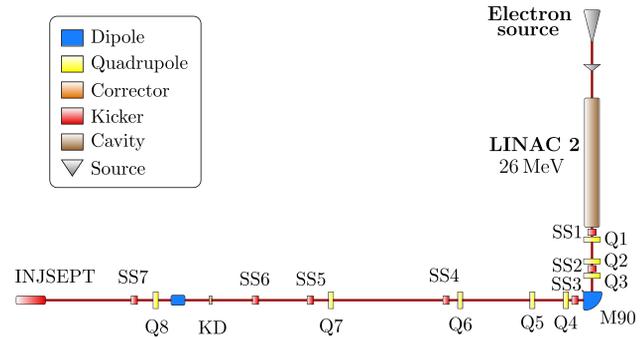
Figure 2: The transfer beamline between LINAC2 (right) and the Booster synchrotron (left).

OPTIMIZATION PROBLEM

Due to the high dimensionality of the problem, optimization of the transfer efficiency is a nontrivial problem. The interplay of different magnetic elements makes the electron transmisson rate into the synchrotron a non-convex target function featuring many local extrema. Due to the fact that the magnetic elements take up to 1 s to reach the desired magnetic field after a set command is issued, evaluations of this target function $y_i = f(x_i)$ are quite costly and need to be kept to a minimum. The optimization algorithm known as Simulated Annealing fits the problem [2], as it does not rely on gradient information or any knowledge of the underlying target function, and is designed to escape local extrema of non-convex target functions (see Fig. 3). The *iterative*
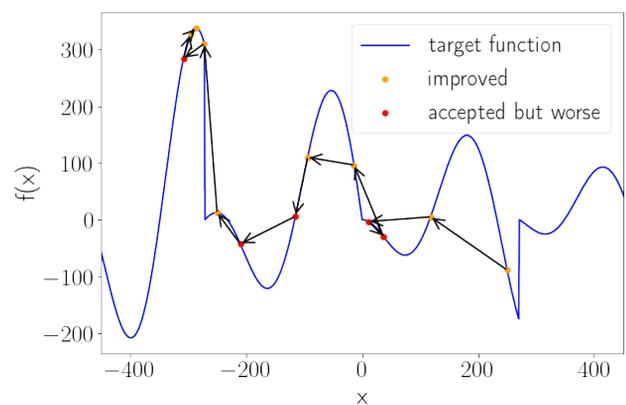
Figure 3: Example of Simulated Annealing on a one-dimensional non-convex target function $f(x)$.

algorithm (schematically depicted in Fig. 4) explores the parameter space by adding *random steps* Δx drawn from a uniform distribution to the current configuration x_n of magnet strengths before evaluating the target function (resulting transmission rate) and comparing this y_t to the previous iter-
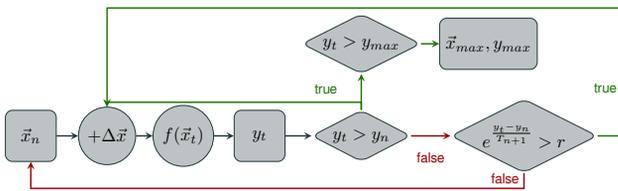
Figure 4: Simulated Annealing.

ation's result y_n. Any increase in transmission leads to the acceptance of the new configuration as initial configuration for the following iteration, while for decreased transmission the metropolis criterion

$$e^{\frac{y_t - y_n}{T_n}} \overset{?}{>} r, \quad R \sim \mathcal{U}_{[a,b]}$$

is evaluated. Here, T_n denotes the *temperature* and is taken from a function[1] falling in iterations n. If the criterion evaluates to `True`, the new configuration is accepted even though it produced a decrease in transmission, otherwise the configuration is rejected and a different random step Δx is taken from the previous configuration x_n.

Employing the metropolis criterion therefore results in a large probability of accepting configurations of decreased transmission early on into the optimization process and if the decrease is only minor, while the acceptance probability is low in the late iterations or for significant decreases. This way, both an increased tendency for exploration at the beginning as well as exploitation towards the maxium in the end can be achieved (see Fig. 3).

IMPLEMENTATION

For the implementation of the optimization procedure described above, a modular approach has been chosen (Fig. 5). This way, the optimization is an interdependent program communicating only with the accelerator control system, allowing to switch easily between a simulation of the transfer beamline and the actual accelerator by exchanging the target address from the computer hosting a virtual instance of the control system to the one of the real control system computer. In case that the simulation, which uses the tool `elegant` [3],
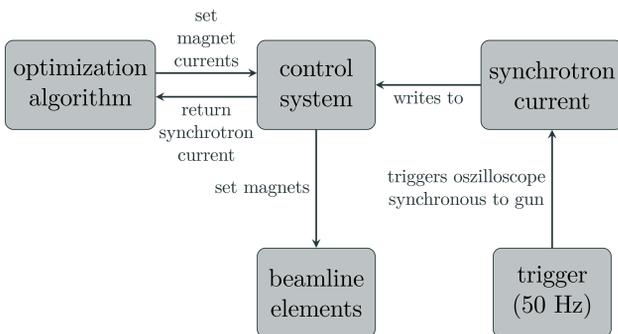
is in use, actual transmission rates can be simulated, and additional information such as the beam envelopes is generated, which allows easy visualization and cross-checking of the results. When the actual accelerator is used, this information is not availiable. So only the current in the synchrotron normalized to the emitted charge from the electron gun, which is proportional to the transmission rate and therefore the quantity to be optimized here, can be measured for each of the 50 Hz trigger pulses.

As beamtime at the actual accelerator is expensive and largely allocated to the experiments receiving the electron beam, the optimization is best be done in a parasitic way. To this end, the cyclicity of the acceleration (Fig. 6) as described before can be taken advantage of: instead of idling the gun, LINAC2 and Booster Synchrotron during the ramp up and extraction phase, the algorithm keeps them running, simply shutting off the injection from synchrotron into the stretcher ring.
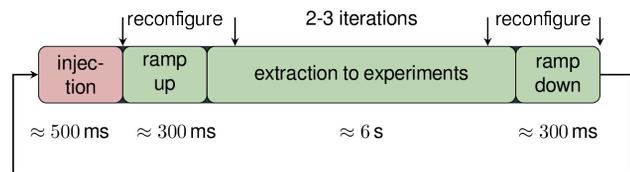
Figure 6: Acceleration cycle at ELSA.

A measurement of the magnetic field response to changes in setpoint current for the beamline elements reveals a delay time of almost 1 s (see Fig. 7). Around 5 such changes can be performed each cycle during the extraction phase without interfering with the ongoing operation. Of those 5, one change needs to be allocated just before the start of the next injection phase to reconfigure the transfer beamline into a workable state for the experiments, and one just after the injection to turn everything back to where the optimization left off, leaving us with just enough time for 3 iterations of the optimization procedure in each cycle. Concerning the reconfiguration before injection, the best currently found magnet settings can be employed each cycle to enable for a continuous improvement of the beamline efficiency during operation.

Figure 5: Code schematics.
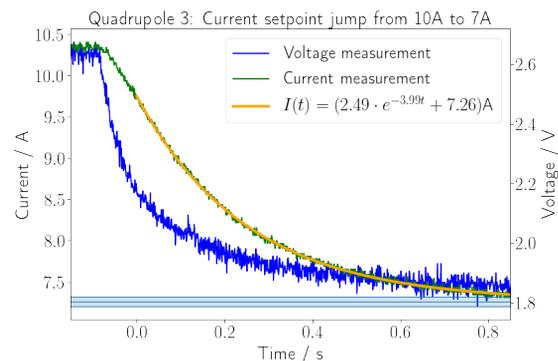
Figure 7: Current delay of magnet powersupply.

---

[1] An example temperature function is given by $T_n = \frac{T_0}{1+n\alpha}$ with $\alpha = 0.01$.

FURTHER APPLICATIONS

The nature of the algorithm, particularly the fact that no knowledge of the target function or any underlying physical problem is needed allows it to be flexibly used for different beamlines and optimization problems. For example, by re-defining the target function, the beam shape and position as observed on optionally connectible beam monitors can be optimized towards a predefined target. The input data would then be a Gaussian fit performed to the beam image, thus re-sulting in values μ_x, μ_z for the beam position and σ_x, σ_z for the width, which need to be condensed to a one-dimensional value by use of a target function like

$$f(\mu_x, \mu_z, \sigma_x, \sigma_z) = a \cdot \left((\mu_x - \mu_x^{\text{target}})^2 + (\mu_z - \mu_z^{\text{target}})^2 \right)$$
$$+ b \cdot \left((\sigma_x - \sigma_x^{\text{target}})^2 + (\sigma_z - \sigma_z^{\text{target}})^2 \right),$$

where the additional parameters a and b allow to modify the weighting of the squared deviations between beam shape and position. In this case the target function is to be minimized. This optimization problem has been tested and shown to reach the target beam parameters at an availiable beamline, proving the adaptibility and flexibility of the algorithm.

RESULTS

The quality of the optimization process depends heavily on an appropriate choice of the hyperparameters *number of iterations*, *step interval size* and *temperature function*. During the testing, these were found by trial and error, such that finally the following results could be obtained:

When using the simulation with the magnets configured as it was the case with the actual accelerator at the time as initial state, the transmission pre optimization was found to be around 14%. Looking at the beam envelopes (Fig. 8), which surpass the radius of the beampipe significantly, the large losses are explained. An optimization run of 1500 iterations increased the transmission over 6-fold to around 84%. The corresponding beam envelopes are nicely confined in the beampipe after the optimization is finished. First results using the actual accelerator look promising as well. In an optimization run of 1500 iterations, the synchrotron current could be increased from around 1.4 mA initially to over 5.3 mA (Fig. 9).
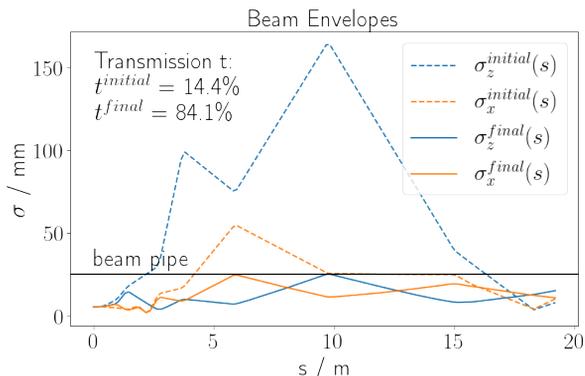
Figure 8: Optimization of simulated beamline.
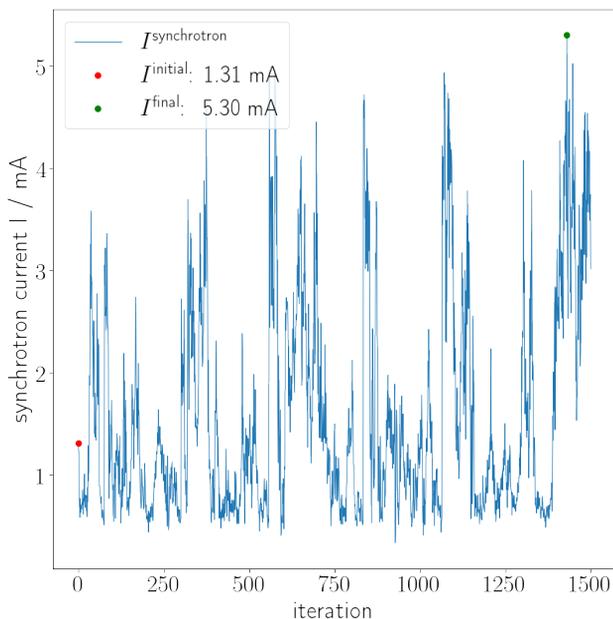
Figure 9: Optimization of accelerator beamline.

CONCLUSION

The optimization of transfer beamline transmission rate is a problem of high dimensionality. Together with the large cost of evaluating the target function this restricts the ap-proches for optimization. Simulated Annealing fits the crite-ria for this kind of problem. The modular implementation allows for testing with a simulated beamline or use with the actual accelerator, and in the second case the cyclicity of the acceleration process can be utilized to allow for parasitic op-timization without interference with the ongoing accelerator operation. Both with the simulation and the actual acceler-ator, the algorithm has been demonstrated to successfully improve transmittance into the synchrotron.

In the future, the optimization can be extended to non-magnetic parameters influencing the transmittance, such as the synchrotron injection timing or LINAC2 energy. By im-plementing an autostart feature, the optimization can be run continually and automatically during beamtime. Longterm evaluations of the found optima might also be used to judge drifts due to temperature and other external factors.

Besides, further optimization problems at different beam-lines (e.g. between Booster Synchrotron and stretcher ring or to the experiments) and with different target functions can be tackled by adapting the existing algorithm.

REFERENCES

[1] Wolfgang Hillert, "The Bonn electron stretcher accelerator ELSA: past and future", *Eur. Phys. J. A*, vol. 28, pp. 139–148, 2006. doi:10.1007/3-540-36754-3_15

[2] Bastien Chopard & Marco Tomassini, *An Introduction to Meta-heuristics for Optimization*, Springer International Publishing, 2018.

[3] M. Borland, "elegant: A flexible SDDS-compliant code for accelerator simulation", Advanced Photon Source, ANL, Ar-gonne, IL, USA, Rep. LS-287, Sep. 2000.