LARGE-SCALE OPTIMISATION OF THE H4 AND M2 BEAMLINES IN THE SPS NORTH EXPERIMENTAL AREA AT CERN

G. Dal Maso*, P. Alexaki[1] D. Banerjee, E. Barber[2], J. Bernhard, N. Charitonidis, S. Esen, A. M. Goillot, F. Metzger, L. J. Nevay, B. Rae, C. Rudigier[3], S. Schuh, F. W. Stummer, R. Tietgen[4], M. Van Dijk,
CERN, Geneva, Switzerland
[1]also at University of Liverpool, Liverpool, United Kingdom,
[2]also at University of Surrey, Surrey, United Kingdom,
[3]also at Geneva Graduate Institute in International Relations & Development, Geneva, Switzerland,
[4]also at University of Mainz, Mainz, Germany

Abstract

The multi-purpose secondary beamlines of the SPS North Experimental Area are invaluable for a rich and diverse physics programme, providing many different particle species and beam characteristics to fixed-target experiments and test-beam users. Many optical elements are required to enable such flexibility, posing a significant challenge in the design and optimisation of the beamline optics. Here, we present an initial optimisation of the H4 and M2 beamlines using genetic algorithms. The aim of this study is to explore the applicability of fast optics-based tracking models (here Xsuite), to enable large-scale, multi-objective optimisation using standard computing hardware. The model is benchmarked with high-fidelity Monte Carlo simulations such as BDSIM and with feedback from experiments. The optimised optics, applied to the H4 electron configuration for NA64, achieved a 30 % increase in available electron yield and a five-fold reduction in beam-related background. In the M2 line, the optimised configurations yielded a two-fold increase in electron yield and a 67 % improvement in muon transmission.

INTRODUCTION

The North Experimental Area (NA) of the Super Proton Synchrotron (SPS) at CERN hosts a network of secondary beamlines delivering a wide variety of particle species to fixed-target experiments and test-beam users over a momentum range from 10 GeV/c to 400 GeV/c [1, 2]. To accommodate diverse experimental requirements, the NA beamlines are equipped with a large number of magnetic optical elements [3], providing significant operational flexibility but resulting in high-dimensional parameter spaces that make systematic optimisation non-trivial.

Large-scale optimisation techniques such as genetic algorithms [4] offer an effective framework to simultaneously address multiple figures of merit and operational constraints. Their practical application to realistic beamline simulations, however, typically requires thousands of evaluations, making fast simulation tools essential. Fast optics-based tracking models, such as Xsuite [5] (used here), fulfil this role, enabling large-scale optimisation on standard computing hardware. We applied this approach to two representative NA beamlines, the 100 GeV/c electron beam in H4 and the 40 GeV/c electron and 160 GeV/c muon beams in M2 (see Fig. 1), obtaining improved optics that were successfully deployed in the 2025 run. A full account of the models, their validation against high-fidelity BDSIM [6] simulations, the optimisation process, and the commissioning results are currently under preparation. The fast model reproduces key beam parameters at the few-percent level when compared to BDSIM, while absolute rates are known to be overestimated. Relative performance gains are nevertheless well captured.

The availability of fast models additionally enables systematic studies of optimiser behaviour that would be otherwise impractical with high-fidelity simulations alone. In this contribution we focus on the configuration of the optimiser: a two-stage workflow, the choice of algorithm and initialisation strategy, and a first systematic characterisation of their performance on H4.

* giovanni.dal.maso@cern.ch

OPTIMISATION WORKFLOW

Fast tracking models

The optimisation relies on fast tracking models implemented in Xsuite [5], an open-source particle tracking library developed at CERN. For the electron beams in both H4 and M2, a purely optics-based model is sufficient, as no material intercepts the beam in the optimised sections. For the M2 muon beam, the model was extended to include the physics of muon production and transport: hadrons produced at the T6 target from the 400 GeV/c proton beam of the CERN Super Proton Synchrotron (SPS) (here at a 172 GeV/c momentum) are transported for 700 m in the so-called *hadron section* (see Fig. 1) while decaying to muons. At the end of the *hadron section*, 9.9 m of beryllium are inserted inside a vertical bend, where the muons further lose energy. Downstream of this section, the muon momentum (here 160 GeV/c) is selected and the hadrons are absorbed. The following *muon section* is used to shape the beam and suppress the beam halo. The resulting speed advantage over BDSIM amounts to approximately three orders of magnitude for H4 and two for M2, as summarised below, allowing large numbers of evaluations to be run on a standard workstation with 8 parallel cores.

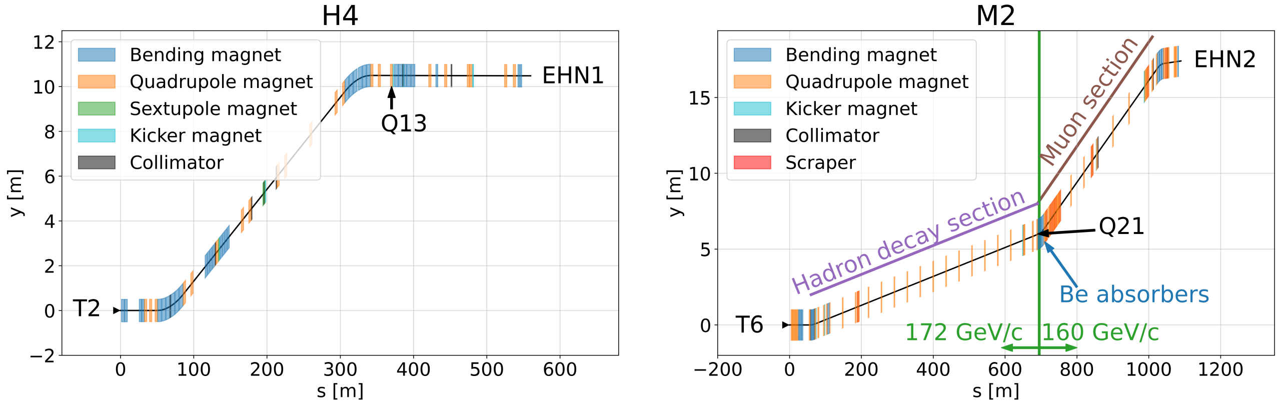

Figure 1: Vertical projections of the H4 (left) and M2 (right) beamline layouts from their target stations to the experimental halls. The breaking points at Q13 (H4) and Q21 (M2) are indicated.

Beamline	Xsuite [ms/part.]	BDSIM [ms/part.]
H4 (e^-)	0.003	3
M2 (π^+/K^+)	0.15	18

Double-staged optimisation strategy

The optimisation strategy is structured in two sequential stages to manage the dimensionality of the problem. In the *mono-objective stage*, all quadrupoles are optimised to maximise the particle yield at the experimental location. At the end of this stage the settings upstream of a breaking point are frozen, reducing the free parameters for the subsequent stage. For M2, the quadrupole Q21, located at 690 m from the target, is a natural boundary as it separates the hadron decay section from the muon section, each requiring distinct optical designs. For H4, the quadrupole Q13, located at 370 m from the target, was chosen so that the number of free parameters in the *multi-objective stage* equals the number of optimisation objectives plus one. A dedicated high-fidelity simulation then generates a particle distribution at the breaking point as input to the next stage.

In the *multi-objective stage*, the downstream section is optimised simultaneously for maximum yield and minimum transverse beam size and divergence at the experiment (located at 1079 m from T6 in Fig. 1), with constraints encoding the experimental requirements, such as the preferred beam spot size, divergence and momentum spread. For M2, the scraper apertures (magnetic collimators) in the muon section are included as additional free parameters. The figures of merit (yield, beam size, and divergence) are evaluated from the tracked particle distributions as the yield and the standard deviations of the marginal phase-space distributions.

Optimiser configuration

The NSGA-II algorithm [7, 8] is used for all optimisation stages using a population size of 100 individuals per generation. As an alternative initialisation strategy, the initial population can be generated using a Tree-structured Parzen Estimator (TPE) Bayesian optimiser [9], as implemented in the `NSGAIIWithTPEWarmupSampler` available via OptunaHub [10]. The TPE samples promising regions of the parameter space based on the history of previous evaluations, providing the evolutionary algorithm with a better-than-random starting population. This *warm-up* phase is continued until the population contains only solutions transmitting at least one particle. Both algorithms are accessed through the Optuna optimisation package [11].

OPTIMISER PERFORMANCE

Mono-objective stage

The performance of the mono-objective stage was characterised by running 500 independent optimisation experiments for each of three configurations: *NSGA-II + TPE warmup* with fixed upstream polarity, *NSGA-II* with fixed upstream polarity, and *NSGA-II + TPE warmup* with unconstrained upstream polarity. Here, *fixed polarity* constrains the quadrupoles upstream of the breaking point to retain their polarity, while *unconstrained polarity* lifts this constraint; in all cases the downstream quadrupoles are free to assume either polarity. The study is performed on H4, whose deterministic objective function isolates the optimiser behaviour directly; in M2 the stochastic nature of the muon production introduces an additional source of variance in the objective evaluations that would be convolved with the optimiser performance.

The *NSGA-II + TPE warmup* configuration reaches the nominal, simulated, H4 electron yield of 3.77×10^{-6} e^- per proton on target with a median of 4.8×10^3 evaluations, compared to 10^4 for *NSGA-II*. The faster initial progress of the warm-up comes at the cost of a less diverse population and leads to less robust long-run exploration; *NSGA-II* performs better on average for larger evaluation budgets. The preference between the two therefore depends on the available evaluation budget.

The *NSGA-II + TPE warmup (unconstrained polarity)* configuration fails to exceed the nominal electron yield in approximately 95 % of experiments within 10^4 evaluations, and in 90 % within 2×10^4. This indicates that the sampling strategy adopted here is not sufficient to efficiently explore

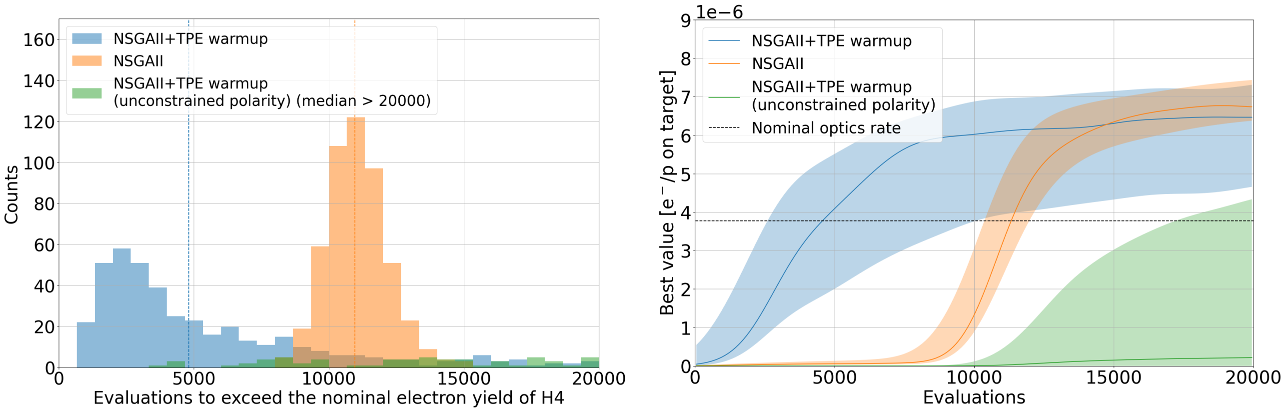

Figure 2: Statistical characterisation of the H4 mono-objective stage over 500 independent trials per configuration. *Left*: distribution of the number of evaluations required to first exceed the nominal H4 electron yield of 3.77×10^{-6} e^- per proton on target, for NSGA-II with TPE warm-up and fixed polarity (blue), NSGA-II only with fixed polarity (orange), and NSGA-II with TPE warm-up and unconstrained polarity (green). Trials that never exceed the threshold are counted at the maximum budget. *Right*: evolution of the mode of the best yield found across trials as a function of evaluations; bands cover 66 % of trials.

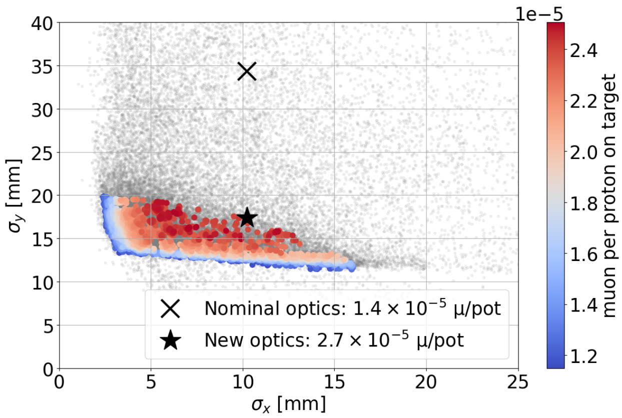

Figure 3: Solutions from the M2 multi-objective stage projected onto the STD(x)-STD(y) plane. Grey points show all solutions; coloured points are selected via a weighted dominance criterion (see text), coloured by the muon yield. The simulated nominal operating point is shown as a black cross; the commissioned solution as a black star.

the unconstrained polarity space. Alternative parameter sampling strategies are under investigation.

Multi-objective stage

The multi-objective stage yields a family of solutions, or Pareto front, trading off yield against transverse beam size at the experimental location. In H4, for the NA64 [12] experiment, this led to a 30 % improvement in electron yield and a five-fold reduction in beam-related background. For M2, Fig. 3 shows the solutions projected onto the STD(x) - STD(y) plane. As the objective evaluations are inherently stochastic, to ease interpretability, the Pareto front is shown for solutions that are not dominated in yield up to a 5 % tolerance, corresponding to approximately twice the single-evaluation uncertainty. The nominal operating point is shown for reference. The plot shows that the *NSGA-II + TPE warmup* optimiser identifies a family of configurations simultaneously improving yield and beam size with respect to the nominal optics. The operating point chosen for the MUonE [13] experiment improves the nominal muon yield by a factor of 2 and halves the vertical beam spot size. During dedicated tests, the improvement in muon yield was measured to be by two-thirds, in good agreement with the expectation. Similarly, this approach was applied to the M2 electron beam and led to an 80 % increase in the electron yield.

CONCLUSION

We have presented a two-stage optimisation workflow based on fast tracking models implemented in Xsuite, enabling large-scale multi-objective optimisation of high-dimensional beamline parameter spaces on standard computing hardware. Applied to the H4 and M2 beamlines of the CERN SPS North Area, an initial optimisation campaign yielded a 30 % increase in electron yield and a five-fold reduction in beam-related background for H4, an improvement in muon transmission by two-thirds for M2, and an 82 % increase in M2 electron yield deployed for the MUonE [13] calorimeter calibration run. A systematic study on H4, conducted over 500 independent experiments per configuration, shows that TPE warm-up provides faster initial convergence at the cost of population diversity, making its advantage over plain NSGA-II dependent on the available evaluation budget. The sampling strategies investigated here are not sufficient to explore unconstrained polarity configurations within a practical budget; alternative approaches are under investigation. These results demonstrate that, despite the significant improvements already achieved, margin for further optimisation remains.

REFERENCES

[1] L. Gatignon, "Design and Tuning of Secondary Beamlines in the CERN North and East Areas," CERN-ACC-NOTE-2020-0043, 2026, v3.

[2] D. Banerjee *et al.*, "The CERN SPS North Area: overview and consolidation plans," CERN-ACC-NOTE-2019-0040, 2019.

[3] P. Coet and N. T. Doble, "An introduction to the design of high-energy charged particle beams," CERN Report 86-69, 1986, `https://cds.cern.ch/record/172394`.

[4] A. S. Hofler, "Genetic algorithms and their applications in accelerator physics," in *Proc. PAC'13*, Pasadena, CA, USA, Sep.–Oct. 2013, paper THTB1, pp. 1111–1115, `https://www.osti.gov/biblio/1113658`.

[5] G. Iadarola *et al.*, "Xsuite: an integrated beam physics simulation framework," in *Proc. IPAC'23*, Venice, Italy, May 2023, pp. 4982–4985, `doi:10.18429/JACoW-IPAC2023-THPL164`.

[6] L. J. Nevay *et al.*, "BDSIM: An accelerator tracking code with Geant4," *Comput. Phys. Commun.*, vol. 252, p. 107200, 2020, `doi:10.1016/j.cpc.2020.107200`.

[7] K. Deb *et al.*, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, 2002, `doi:10.1109/4235.996017`.

[8] R. Saravanan, S. Ramabalan, and C. Balamurugan, "Evolutionary optimal trajectory planning for industrial robot with payload constraints," *Int. J. Adv. Manuf. Technol.*, vol. 38, pp. 1213–1226, 2008, `doi:10.1007/s00170-007-1169-7`.

[9] J. Bergstra *et al.*, "Algorithms for hyper-parameter optimization," in *Proc. NeurIPS'11*, Granada, Spain, Dec. 2011, pp. 2546–2554.

[10] OptunaHub, `https://hub.optuna.org/samplers/nsgaii_with_tpe_warmup_sampler/`.

[11] T. Akiba *et al.*, "Optuna: A next-generation hyperparameter optimization framework," in *Proc. KDD'19*, Anchorage, AK, USA, Aug. 2019, pp. 2623–2631, `doi:10.1145/3292500.3330701`.

[12] M. Yu. Andreev *et al.*, "Search for a Light Dark Matter Candidate with the NA64 Experiment at CERN," Phys. Rev. Lett., 131 161801, `doi:10.1103/PhysRevLett.131.161801`.

[13] G. Abbiendi *et al.*, "Letter of Intent: the MUonE project," CERN-SPSC-2019-026, SPSC-I-252, `https://cds.cern.ch/record/2677471`.