particles

MDPI

Article

Beamline Optimisation for High-Intensity Muon Beams at PSI Using the Heterogeneous Island Model

Eremey Valetov [1,2,*], Giovanni Dal Maso [1,3], Peter-Raymond Kettle [1], Andreas Knecht [1] and Angela Papa [1,4,5] on behalf of the HIMB Project

[1] Paul Scherrer Institut (PSI), 5232 Villigen, Switzerland; giovanni.dal-maso@psi.ch (G.D.M.); peter-raymond.kettle@psi.ch (P.-R.K.); a.knecht@psi.ch (A.K.); angela.papa@psi.ch (A.P.)
[2] Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA
[3] Institute for Particle Physics and Astrophysics-ETH Zürich, Otto-Stern-Weg 5, 8093 Zürich, Switzerland
[4] INFN Sezione di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
[5] Dipartimento di Fisica, Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
* Correspondence: eremey.valetov@psi.ch or evaletov@fnal.gov

Abstract: The High Intensity Muon Beams (HIMB) project at the Paul Scherrer Institute (PSI) will deliver muon beams with unprecedented intensities of up to 10^{10} muons/s for next-generation particle physics and material science experiments. This represents a hundredfold increase over the current state-of-the-art muon intensities, also provided by PSI. We performed beam dynamics optimisations and studies for the design of the HIMB beamlines MUH2 and MUH3 using *Graphics Transport*, *Graphics Turtle*, and *G4beamline*, the latter incorporating PSI's own measured π^+ cross-sections and variance reduction. We initially performed large-scale beamline optimisations using asynchronous Bayesian optimisation with *DeepHyper*. We are now developing an island-based evolutionary optimisation code `glyfada` based on the Paradis**eo** framework, where we implemented Message Passing Interface (MPI) islands with OpenMP parallelisation within each island. Furthermore, we implemented an island model that is also suitable for high-throughput computing (HTC) environments with asynchronous communication via a Redis database. The code interfaces with the codes *COSY INFINITY* and *G4beamline*. The code `glyfada` will provide heterogeneous island model optimisation using evolutionary optimisation and local search methods, as well as part-wise optimisation of the beamline with automatic advancement through stages. We will use the `glyfada` for a future large-scale optimisation of the HIMB beamlines.

Keywords: muon; beamline; intensity frontier; evolutionary optimization; hybrid; heterogeneous; island model; Paradis**eo**; *G4beamline*; *COSY INFINITY*

check for updates

Citation: Valetov, E.; Dal Maso, G.; Kettle, P.-R.; Knecht, A.; Papa, A., on behalf of the HIMB Project. Beamline Optimisation for High-Intensity Muon Beams at PSI Using the Heterogeneous Island Model. *Particles* **2024**, *7*, 683–691. https://doi.org/10.3390/particles7030039

Academic Editor: Armen Sedrakian

Received: 30 June 2024
Revised: 20 July 2024
Accepted: 27 July 2024
Published: 1 August 2024

Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The Paul Scherrer Institute (PSI) provides state-of-the-art facilities for intensity frontier muon experiments, including the world-leading continuous muon delivery rate of several 10^8 muons/s. To address the needs of next-generation muon experiments, the High Intensity Muon Beams (HIMB) project—a part of the IMPACT (Isotope and Muon Production using Advanced Cyclotron and Target technologies) project, which in addition to HIMB includes the TATTOOS project aiming to create radioisotopes for advanced cancer treatments in the quantities required for clinical studies —is developing a new target station TgH and the novel beamlines MUH2 and MUH3 (see Figure 1) [1–3] to provide muon intensities up to 10^{10} muons/s, two orders higher than available at present.

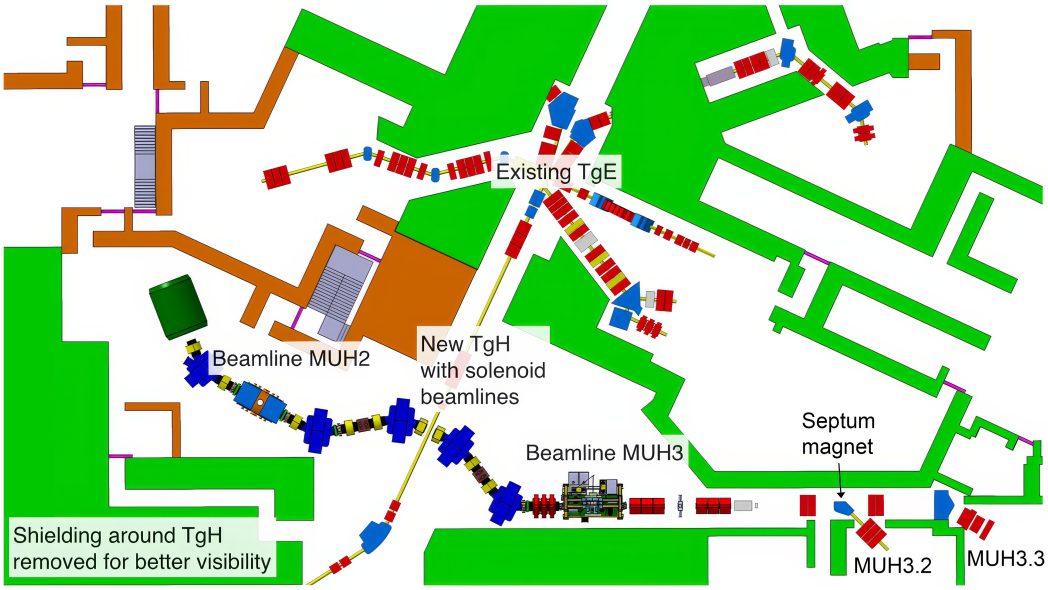

Figure 1. Layout of the new TgH target station and the HIMB beamlines MUH2 and MUH3. Blue, yellow, and red modules in the HIMB beamlines represent magnetic dipoles, solenoids, and magnetic quadrupoles, respectively. (General concept only; may not include details or revisions).

The enhanced muon delivery capabilities of the HIMB will significantly benefit a range of cutting edge experiments in particle physics. For instance, the Mu3e experiment, which aims to detect the extremely rare decay of a muon into three electrons [4], would benefit substantially from these increased muon rates. This decay mode, a form of charged lepton flavour violation (cLFV) [5], is currently virtually undetectable within the Standard Model framework at a branching ratio level of 10^{-55}.

Similarly, a future iteration of the $\mu \rightarrow e\gamma$ (MEG) experiment [6] searching for the decay of a muon into an electron and a photon, another highly suppressed cLFV process, would benefit from the significant increase in supplied muons. Various other experiments and studies [7] could potentially leverage the enhanced muon flux provided by HIMB as well.

The benefits of HIMB also extend into materials science. Muon spin rotation (μSR) measurements such as those employing pixel-based detectors or microbeams [7] require higher muon intensities to enable improved spatial resolution and sensitivity in material characterisation studies.

The hundredfold leap from the current 10^8 muons/s to the targeted 10^{10} muons/s is driven by the need for higher statistical precision in less time during experimental measurements. This substantial increase in muon flux will substantially reduce the required experimental run times, making previously impractical or lengthy measurements feasible. For example, achieving a sensitivity of 10^{-16} in phase II of the Mu3e experiment would be possible within a reasonable timeframe [7], rather than requiring over a decade of continuous operation. See [7] for details regarding the new opportunities in μSR thanks to the increased muon flux.

The HIMB project incorporates several innovative design features to achieve its goal. These include a new slanted graphite target TgH with optimised geometry to increase the surface muon flux from a proton beam by a factor of 1.4 to 2 [1,8,9], high-acceptance capture solenoids positioned in close proximity to the target, and solenoidal focusing for higher transmission relative to conventional quadrupole focusing. The project also features a redesigned layout with reduced bending angles and large-aperture dipoles to maximise muon transmission. While the effect of the design features typically varies depending on factors such as the beam optics and apertures, a comparative study of the existing μE4 beamline to a simulated solenoid-based beamline has shown a capture efficiency of ∼26% with solenoidal focusing versus ∼6% with quadrupole focusing, and

a transmission efficiency of ∼40% with solenoidal focusing over ∼7% with quadrupole focusing, providing an overall enhancement factor of ∼25 considering both capture and transmission efficiency [1].

As part of the overall design of the beamlines, we have performed large-scale beamline optimisation [2,3,10] using asynchronous Bayesian optimisation, running the asynchronous Bayesian optimisation code *DeepHyper* [11] on a computing cluster. These simulations were performed using the particle physics and transport code *G4beamline* [12]. The scientific code *COSY INFINITY* [13] and beam dynamics codes *Graphics Transport* [14] and *Graphics Turtle* [15] were also used for design, optimisation, and study. We use a custom build of *G4beamline* with PSI's own measured π^+ cross-sections [8] and splitting at pion production and decay vertices for variance reduction; however, for beamline optimisation, we use preproduced surface muon beam input files, with the simulation starting next to the target or further in the beamline as appropriate. The beamline design and optimisation process is iterative and involves magnet design, optimisation, and modelling ([16]; see also [2,3]). This paper is based on a part of our muon beamline beam dynamics design and optimisation efforts at PSI (see also, e.g., [10,17]).

We optimised beamline parameters such as bending and kicker dipole, solenoid, and quadrupole tuplet fields, drift lengths, dipole position and rotation offsets, and other operational and design parameters (e.g., target rotation angle and dipole magnet type). The optimisations were performed in stages, optimising overlapping sets of roughly eight to ten parameters at a time, starting from the target station and proceeding toward the final foci.

We have achieved a transmission of 1.34×10^{10} muons/s in the particle physics beamline MUH2 [2] and 3.43×10^{9} muons/s to the collimator windows in front of the septum magnet in the materials science beamline MUH3 [2,3] (the septum magnet directs the beam into branches MUH3.2 and MUH3.3 and supports "shared" and "only" modes, where the beam is supplied to both or one of the branches). We performed a variety of final focus optimisations for the MUH3.2 and MUH3.3 branches of the MUH3 beamline, maximising the simulated muon rates on the detectors.

Throughout the optimisation process, we collaborated with the magnet design and analysis team to engage in an iterative beamline element design process. As part of the standard procedure, we ensured that all optimisation parameters corresponded to feasible operating setpoints. For instance, in many optimisations, we constrained the capture solenoid field to below 0.45 T, while the transport solenoid fields were limited to a maximum of 0.4 T.

It should be noted that we employed sufficiently high effective statistics of 10^{11} to 10^{13} protons on target for the design optimisations. Rigorous uncertainty quantification (UQ) is typically unnecessary for beamline transmission optimisation. Early in the project, we decided not to formally perform UQ for quantities such as beamline transmission.

While asynchronous Bayesian optimisation has been generally effective in our optimisations of the HIMB beamlines, we have noticed that, unsurprisingly, evolutionary optimisation (EO, a quasi-global method inspired by biological evolution using mechanisms such as mutation and selection to iteratively explore the search space and converge toward solutions) typically provides somewhat better results. We started in 2023 by developing an EO code called `glyfada` (see Figure 2) based on the Paradis**eo** [18] EO framework for optimisation of charged particle optical systems in general, with a specific application to large-scale beamline optimisation. The current version of this code supports the heterogeneous island model, multi-objective optimisation, and parallelisation using MPI, OpenMP, and Redis, and can interface with the codes *COSY INFINITY* and *G4beamline*. The code can use a *Python* model file that has the same structure and format as a model file (often called `model.py`) for the asynchronous Bayesian optimisation code *DeepHyper*. We are developing a feature for optimisation in stages (e.g., multiple stages from upstream to downstream) with automatic advancement using dynamic reconfiguration of the islands.

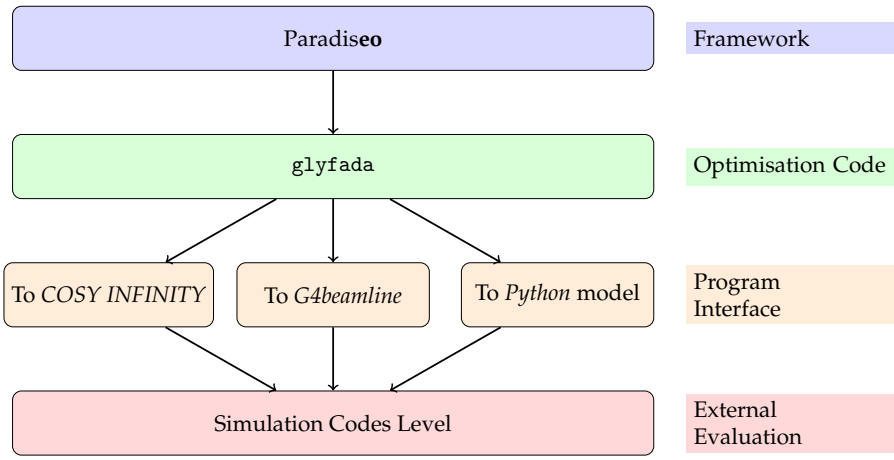

Figure 2. Conceptual diagram showing the optimiser `glyfada` in relation to the Paradis**eo** optimisation framework that it is based on as well as its interfaces with the simulation codes used to evaluate the objective function.

2. Materials and Methods

EO has been found to be effective for the design and optimisation of beam dynamics systems. For instance, the GATool Evolutionary Algorithm code implemented using *COSY INFINITY* was successfully used for the design of the front-end subsystem of the Neutrino Factory [19].

*2.1. Paradis**eo** Framework*

Our optimisation code `glyfada` is based on the Paradis**eo** framework, which is a modular and customisable open-source C++ framework with a large number of EO, particle swarm optimisation (PSO), and local search (LS) algorithms. PSO is an optimisation technique inspired by the social behaviour of bird flocking or fish schooling, where potential solutions "moving" through the problem space guided by their own and the swarm's best-known positions. LS methods search for improved solutions by exploring the immediate neighbourhoods of the current best solutions. LS is often used to fine-tune results obtained using quasi-global optimisation methods.. The framework has a codebase that supports multi-objective optimisation, parallelisation, and heterogeneous island model implementation. Having compared several options, Paradis**eo** is likely the best library to build an optimiser with the features of `glyfada`.

For our purposes, the main limitation of Paradis**eo** is that while it has been supported by about fifteen institutions over many years of development, the support for this project has decreased in recent years. This was seemingly because, similar to *Geant4*, such C++ frameworks have a somewhat substantial learning curve and require relatively high development effort to build an application; thus, easily configurable solutions that work out of the box are preferred in many cases .

While developing `glyfada`, we also implemented a number of new classes and features in the Paradis**eo** code, and we plan to contribute the framework code enhancements to the Paradis**eo** project. This will help to ensure that `glyfada` can be easily built using future versions of Paradis**eo**, helping to support the long-term sustainability of both codes.

2.2. Heterogeneous Island Model

The heterogeneous island model (see Figure 3) implemented in `glyfada` is an application of advanced concepts in optimisation to solve complex optimisation problems in charged particle optics, where the objective function is computed using a CPU-intensive simulation involving beam dynamics and potentially, particle physics. This model uses partitions called "islands" configured with different optimisation methods such as EO, LS, or other algorithms. Islands can also use the same algorithm with different parameters, such different mutation probabilities in case of NSGA-II or the neighbourhood structure in case

of LS algorithms. The heterogeneous model can overcome the limitations of individual methods and exploit their complementary strengths. For instance, EO algorithms excel at exploring broad solution spaces and avoiding convergence to local optima, while LS methods can efficiently fine-tune solutions in promising regions.

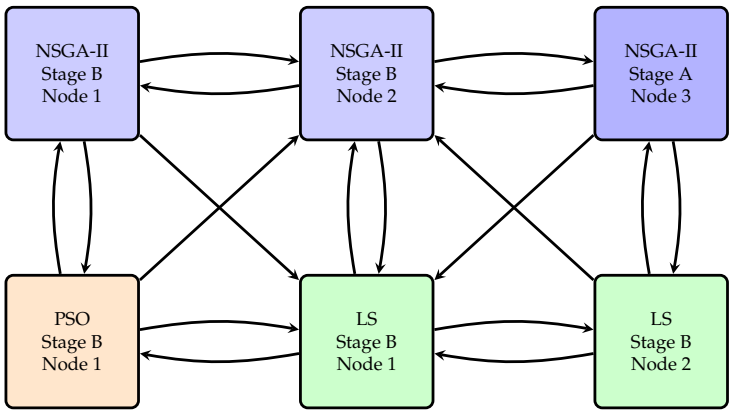

Figure 3. Illustration of the heterogeneous island model concept with EO algorithm, NSGA-II, PSO, and LS. In this example, the algorithm state is primarily in Stage B, with one Stage A node, demonstrating dynamic multistage execution. Arrows between islands represent the periodic migration of individuals.

Furthermore, the heterogeneous island model facilitates part-wise optimisation of large systems with extensive parameter spaces, dividing the problem into more manageable partitions. The island model periodically migrates selected solutions between islands, exchanging information between optimisers and facilitating coordination of the overall optimisation process. We plan to implement automatic advancement between stages of optimisation, corresponding to different parts of a beamline, with some feedback from later stages to earlier stages to account for the fact that in certain cases reaching an optimal solution for the complete system may include seemingly suboptimal solutions for its parts. Optimisation stages may be implemented using dynamic reconfiguration of the islands.

The heterogeneous island model is particularly useful for model parameter optimisation in high-statistics simulations using codes such as *G4beamline*, while high-order transfer maps using the differential–algebraic (DA) data type in *COSY INFINITY* typically enable highly efficient and highly accurate solutions using internal optimisers. [1] In certain particularly complex cases, such as can arise in the design of a fixed-field alternating gradient (FFAG) ring, for example, a practical approach may be to combine optimisation using `glyfada` as an external optimiser with internal optimisation using *COSY INFINITY*, along with its FFAG module `COSYFFAG` where applicable.

2.3. Hybrid Parallelism

The parallelised optimisation (PEO) module of Paradis**eo** v.1 has been deprecated, and as such is no longer operable or distributed. The developers of Paradis**eo** re-implemented PEO's shared memory part in v.2 of the framework as the Shared-Memory Multiprocessing (SMP) module; however, the MPI part was not re-implemented, rendering the heterogeneous island model unusable for CPU-intensive applications.

We have implemented a new Message Passing Interface (MPI)-based island model class for Paradis**eo**, additionally making OpenMP parallelisation possible within each node. An island can be equivalent to a node with a specialised configuration or can represent multiple nodes. Furthermore, considering the option of running optimisations on high-throughput computing (HTC) resources, we implemented an island model class where each node connects to a Redis database for parallelisation and where OpenMP parallelisation is available internally within each node. See Figure 4 for an illustration of the implemented parallelisation approaches. (HTC refers to the use of large amounts of computing resources

over extended periods to accomplish computational tasks while emphasising throughput rather than coupled or advanced parallelisation. It is often used for large-scale simulations or data analysis in scientific research).

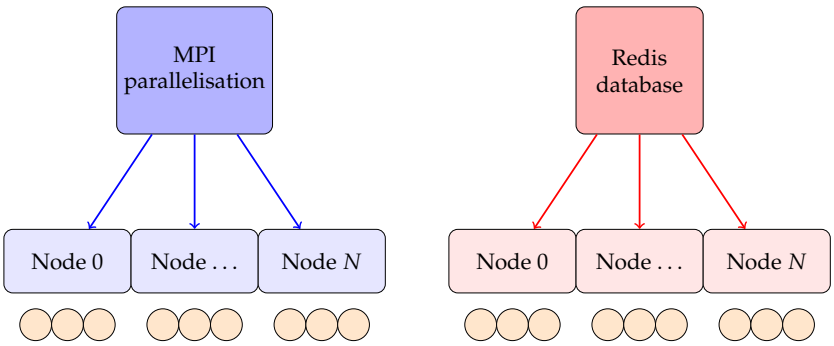

Figure 4. The diagram on the left shows hybrid parallelisation using MPI for inter-node communication and OpenMP for intra-node parallelism. The diagram on the right illustrates the use of a Redis database for asynchronous inter-node communication combined with OpenMP. The orange circles below the nodes represent OpenMP threads.

2.4. Simulation Interfaces

For evaluation of the objective function, we implemented interfaces to perform simulations using *G4beamline* or *COSY INFINITY*. We also implemented an interface to execute a *DeepHyper*-compatible *Python* model. This is particularly useful because, in addition to the same model file being reusable for *DeepHyper* optimisation, the *Python* file can encapsulate the formal model definition with the full logic for simulation runs, analyses, objective function computation, and quality checks whether locally or on a cluster. The *Python* model file can perform simulations by calling any code, including *G4beamline* or *COSY INFINITY*.

3. Results

In a comparison for final focus optimisation of the MUH3.3 branch of the materials science MUH3 beamline with modified collimating square windows immediately upstream of the septum magnet, Bayesian optimisation using *DeepHyper* provided a solution with a rate of 8.67×10^2 kHz/mA on an 8-mm-diameter detector [3], while evolutionary optimisation using an early version of `glyfada` provided 1.76×10^3 kHz/mA. Quadrupole, bending magnet, and kicker magnet currents and the offset of a square collimating window were all optimised. For reference, the existing rates are only approximately 70 kHz/mA in the πM3.2, πM3.3 branches of the πM3 beamline, which will be replaced by the new MUH3 beamline. We provide these rates to compare the performance of Bayesian optimisation using *DeepHyper* and evolutionary optimisation using `glyfada` on a specific relevant example.

For reference, the nominal proton beam current is 2.3 mA to 2.4 mA. The rate on an 8-mm-diameter detector is lower than the full transmission in the 320-mm-diameter beam pipe, considering the substantial transverse size of the muon beam and the momentum spread of the beam of roughly 25 MeV/c to 29.79 MeV/c, varying depending on factors including beam optics and the location in the beamline. Optimising the final focus in a beamline, especially with two concurrently used branches, is a detailed and complex matter that often benefits from multi-objective optimisation algorithms. The MUH3 materials science beamline has lower transmission efficiency requirements compared to the MUH2 particle physics beamline. The preliminary detector rates presented here meet the current design requirements but are subject to further improvements. Notably, `glyfada` has undergone extensive testing throughout its development.

The early version of `glyfada` provided a significant improvement over Bayesian optimisation in a relatively challenging final focus optimisation problem. We note that such a large difference in optimality is likely not fully representative of most cases in the

optimisation of the HIMB beamlines, where we expect a smaller improvement on average. In a connected but somewhat separate line of optimisation and study, we have noticed that evolutionary optimisation using the NSGA-II method typically provides better objective function values than Bayesian optimisation when optimising the MUH2 beamline [10].

We will perform new large-scale optimisations of the MUH3 beamline (see Figure 5) using the heterogeneous island optimisation code `glyfada`, leveraging an automatic advancement through stages and a combination of evolutionary optimisation and local search islands. We will present the results of this optimisation when they become available.

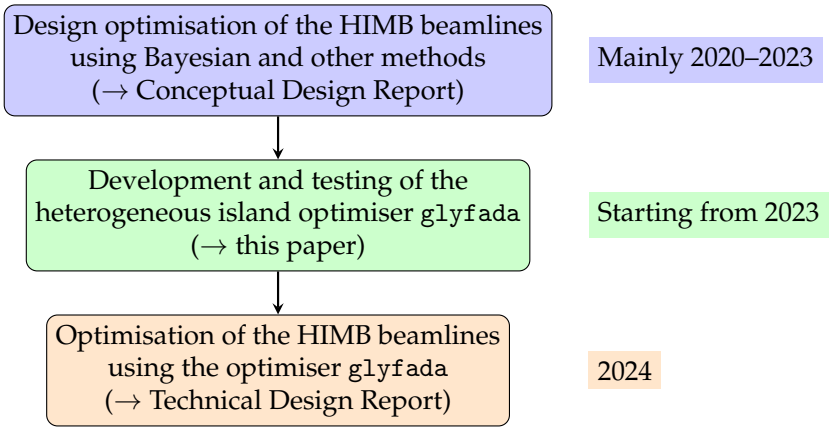

Figure 5. Timeline of the HIMB beamline optimisation process from the perspective of this paper: initial design optimisation (see, e.g., [1–3]), mainly using Bayesian methods (blue); development of the `glyfada` optimiser (green); and planned HIMB beamline optimisation using `glyfada` (orange). The arrows ("→") roughly indicate correspondence to a publication.

4. Discussion

We are developing an optimiser called `glyfada` for charged particle optics optimisation problems, with a focus on applications with CPU-intensive simulations involving beam dynamics and particle physics processes. The program features the heterogeneous island model, a set of optimisation algorithms including evolutionary and local search optimisation, and OpenMP/MPI and Redis as the parallelisation options.

A comparison with asynchronous Bayesian optimisation has shown significant improvement in the objective: a simulated rate of 1.76×10^3 kHz/mA over 8.67×10^2 kHz/mA on an 8-mm-diameter detector in a final focus optimisation of the MUH3.3 beamline branch. This comparison is a specific, relevant example of the advantages of optimisation using `glyfada` over Bayesian optimisation.

After implementing automatic advancement through optimisation stages, such as part-wise from the target station to the final foci, we will apply `glyfada` for a new large-scale optimisation of the revised version of the MUH3 beamline model, supporting the next-generation high-intensity muon experiments and measurements at PSI.

Author Contributions: Conceptualisation, P.-R.K. and A.K. (the HIMB project, target station, and beamlines), E.V. (optimisation methods, optimiser `glyfada`); methodology (creation of models), E.V.; software (simulation programs and scripts, optimiser `glyfada`), E.V.; validation, A.K. (beamline models, simulations, and studies) and E.V. (optimiser; beamline models, simulations, and studies); formal analysis (beamline optimisation and studies), E.V.; investigation (the HIMB project), P.-R.K., A.K., A.P., E.V., G.D.M.; investigation (optimiser `glyfada`), E.V.; resources (the HIMB project), A.K.; data curation (beamline optimisation, simulation, and studies), E.V.; writing—original draft preparation, E.V.; writing—review and editing, E.V., P.-R.K., A.K., A.P. and G.D.M.; visualisation (beamline optimisation, simulation, and studies; does not include the original version of Figure 1), E.V.; supervision (the HIMB project; beamline optimisation, simulation, and studies), A.K.; project administration, A.K. (the HIMB project; beamline optimisation, simulation, and studies) and E.V. (optimiser `glyfada`); funding acquisition, A.K. (the HIMB project) and E.V. (Marie Skłodowska-Curie grant). All authors have read and agreed to the published version of the manuscript. This paper

does not cover a significant part of HIMB beamline design, optimisation, and studies by P-R.K. and G.D.M.; that research will be covered by separate publications.

Funding: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 884104.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors upon request, subject to any applicable institutional or regulatory limitations.

Acknowledgments: We gratefully acknowledge the contributions of the entire IMPACT team, including colleagues involved in the HIMB project. We are thankful to Hubertus Luetkens (PSI) for productive and interesting discussions regarding aspects of the MUH3 beamline design. We gratefully acknowledge Rebecca Riccioli (PSI) for developing and iterating magnet designs and providing field maps, which are integral to the beamline design process. Development of the optimiser `glyfada` used the resources of the Wilson computing cluster at Fermi National Accelerator Laboratory, which is managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study, in the collection, analysis, or interpretation of the data, in the writing of the manuscript, or in the decision to publish the results.

Note

[1] *COSY INFINITY*'s DA data type is an advanced computational tool that extends truncated power-series algebra (TPSA) with integral and derivative operators [20]. It implements a form of automatic differentiation, enabling efficient computation of high-order transfer maps and rigorous error bounds. This makes DA computations particularly powerful for accurate and efficient beam dynamics analyses and simulations, among other applications.

References

1. Barth, L.; Baumgarten, C.; Caminada, L.; Chang, J.; Dal Maso, G.; Dölling, R.; Forss S.; Gabard A.; Geiselhart, C.; Häfeli, R.; et al. IMPACT conceptual design report. In *Paul Scherrer Institut Report No. 22-01*; Eichler, R., Kiselev, D., Koschik, A., Knecht, A., van der Meulen, N., Scheibl, R., Eds.; Paul Scherrer Institut: Villigen PSI, Switzerland, 2022.
2. Valetov, E. Beamline Design and Optimisation for High Intensity Muon Beams at PSI. *J. Phys. Conf. Ser.* **2023**, *2420*, 012053. [CrossRef]
3. Valetov, E. Beamline Optimization Methods for High Intensity Muon Beams at PSI. In Proceedings of the NAPAC'22, Albuquerque, NM, USA, 7–12 August 2022; JACoW Publishing: Geneva, Switzerland; pp. 63–66.
4. Arndt, K.; Augustin, H.; Baesso, P.; Berger, N.; Berg, F.; Betancourt, C.; Bortoletto, D.; Bravar, A.; Briggl, K.; vom Bruch, D.; et al. Technical design of the phase I Mu3e experiment. *Nucl. Instrum. Methods A* **2021**, *1014*, 165679. [CrossRef]
5. de Gouvêa, A.; Vogel, P. Lepton flavor and number conservation, and physics beyond the Standard Model. *Prog. Part. Nucl. Phys.* **2013**, *71*, 75–92. [CrossRef]
6. Baldini, A.M.; Baracchini, E.; Bemporad, C.; Berg, F.; Biasotti, M.; Boca, G.; Cattaneo, P.W. The design of the MEG II experiment. *Eur. Phys. J. C* **2018**, *78*, 1–60. [CrossRef]
7. Aiba, M.; Amato, A.; Antognini, A.; Ban, S.; Berger, N.; Caminada, L.; Chislett, R.; Crivelli, P.; Crivellin, A.; Maso, G.D.; et al. Science Case for the new High-Intensity Muon Beams HIMB at PSI. *arXiv* **2021**, arXiv:2111.05788.
8. Berg, F.; Desorgher, L.; Fuchs, A.; Hajdas, W.; Hodge, Z.; Kettle, P.R.; Knecht, A.; Lüscher, R.; Papa, A.; Rutar, G.; et al. Target studies for surface muon production. *Phys. Rev. Accel. Beams* **2016**, *19*, 024701. [CrossRef]
9. Kettle, P.R. *HiMB Technical Note*; Technical Report; Paul Scherrer Institut: Villigen, Switzerland, 2019.
10. Dal Maso, G. Paul Scherrer Institut. Personal communication, 2023.
11. Balaprakash, P.; Salim, M.; Uram, T.D.; Vishwanath, V.; Wild, S.M. DeepHyper: Asynchronous Hyperparameter Search for Deep Neural Networks. In Proceedings of the 2018 IEEE 25th International Conference on High Performance Computing (HiPC), Bengaluru, India, 17–20 December 2018; pp. 42–51.
12. Muons, Inc. *G4beamline (Version 3.06) [Computer Software];* Muons, Inc.: Batavia, IL, USA, 2019.
13. Makino, K.; Berz, M. COSY INFINITY Version 9. *Nucl. Instrum. Methods A* **2006**, *558*, 346–350. [CrossRef]
14. Rohrer, U. PSI Graphic TRANSPORT Framework by U. Rohrer Based on a CERN-SLAC-FERMILAB Version by K.L. Brown et al. [Computer Software] 2007. Available online: http://aea.web.psi.ch/Urs_Rohrer/MyWeb/trans.htm (accessed on 26 July 2024).
15. Rohrer, U. PSI Graphic TURTLE Framework by U. Rohrer Based on a CERN-SLAC-FERMILAB Version by K.L. Brown et al. [Computer Software] 2007. Available online: http://aea.web.psi.ch/Urs_Rohrer/MyWeb/turtle.htm (accessed on 26 July 2024).

16. Riccioli, R.; Calzolaio, C.; Gabard, A.; Kettle, P.R.; Kiselev, D.; Knecht, A.; Martinie, R.P.; Dal Maso, G.; Papa, A.; Rauber, T.; et al. Magnet Design for the High-Intensity Muon Beams Project (HIMB) at PSI's Accelerator Complex HIPA. *IEEE Trans. Appl. Supercond.* **2024**, *34*, 4004305. [CrossRef]
17. Dal Maso, G.; Kettle, P.R.; Knecht, A.; Leonetti, F.; Lospalluto, G.; Papa, A.; Valetov, E. Bayesian tuning of the Compact Muon Beam Line for the Mu3e experiment. In Proceedings of the 16th Pisa Meeting on Advanced Detectors, Isola d'Elba, Italy, 26 May–1 June 2024. Submitted to Nucl. Instrum. Methods A.
18. Dreo, J.; Liefooghe, A.; Verel, S.; Schoenauer, M.; Merelo, J.J.; Quemy, A.; Bouvier, B.; Gmys, J. Paradiseo: From a Modular Framework for Evolutionary Computation to the Automated Design of Metaheuristics: 22 Years of Paradiseo. In Proceedings of the Genetic and Evolutionary Computation Conference Companion, New York, NY, USA, 10–14 July 2021; pp. 1522–1530.
19. Poklonskiy, A.; Neuffer, D. Evolutionary Algorithm for the Neutrino Factory Front End Design. *Int. J. Mod. Phys. A* **2009**, *24*, 959–973. [CrossRef]
20. Berz, M.; Makino, K. New approaches for the validation of transfer maps using remainder-enhanced differential algebra. *Nucl. Instrum. Methods A* **2004**, *519*, 53–62. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.