

<b>Noname manuscript No.</b> (will be inserted by the editor)
--

---

# Derivative-Free Optimization of a Rapid-Cycling Synchrotron

Jeffrey S. Eldred · Jeffrey Larson ·  
Misha Padidar · Eric Stern ·  
Stefan M. Wild

August 9, 2021

**Abstract** We develop and solve a constrained optimization model to identify an integrable optics rapid-cycling synchrotron lattice design that performs well in several capacities. Our model encodes the design criteria into 78 linear and nonlinear constraints, as well as a single nonsmooth objective, where the objective and some constraints are defined from the output of Synergia, an accelerator simulator. We detail the difficulties of the 23-dimensional simulation-constrained decision space and establish that the space is nonempty. We use a derivative-free manifold sampling algorithm to account for structured non-differentiability in the objective function. Our numerical results quantify the dependence of solutions on constraint parameters and the effect of the form of objective function.

**Keywords** Numerical optimization · Simulation optimization · Particle accelerator design · Rapid cycling synchrotron

---

This manuscript is based upon work supported by the applied mathematics and the Scientific Discovery through Advanced Computing (SciDAC) programs of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357. This manuscript has been authored by Fermi Research Alliance, LLC under Contract DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. Synergia development has been supported by the Office of Advanced Scientific Computing Research and Office of High Energy Physics SciDAC program.

---

J.S. Eldred · E. Stern  
Fermi National Accelerator Laboratory, Batavia, IL, USA  
E-mail: jseldred@fnal.gov; E-mail: egstern@fnal.gov

J. Larson · M. Padidar · S.M. Wild  
Mathematics & Computer Science Division, Argonne National Laboratory, Lemont, IL, USA  
E-mail: jmlarson@anl.gov; E-mail: wild@anl.gov

M. Padidar  
Center for Applied Mathematics, Cornell University, Ithaca, NY, USA  
E-mail: map454@cornell.edu

## 1 Introduction

With the advent of the Long-Baseline Neutrino Facility (LBNF) at Fermilab [1], there is a strong motivation to follow the PIP-II 1.2 MW upgrade [28] with a 2.4 MW upgrade of the Fermilab proton accelerator complex. The construction of a new high-intensity rapid-cycling synchrotron (RCS) accelerator would provide a clear path to achieve the 2.4 MW benchmark for the LBNF program and set the stage for the next generation of particle physics experiments [2, 15, 36].

Designing such an RCS accelerator can be greatly aided by numerical optimization techniques, which have been used to design and tune particle accelerators. Particle accelerator optics and performance characteristics have been optimized by simplex methods [19, 47]; particle swarms [20, 41]; multiobjective genetic algorithms [49, 55]; and, more recently, genetic algorithms enhanced by machine learning [12, 30]. Particle accelerator operations also deploy online tuning methods, which include model-based methods [35] and local extrema-seeking methods [38, 39, 46] and methods employing Gaussian process models [11, 44]. In this paper we develop a nonsmooth, constrained optimization model for particle accelerator performance and present the first nonsmooth optimization of nonlinear integrable accelerator lattice optics.

We develop and solve an optimization model that seeks an RCS lattice design that performs well in several capacities. Our model takes the form

$$\begin{aligned} & \underset{\mathbf{x} \in \Omega}{\text{minimize}} f(\mathbf{x}) \\ & \Omega = \{ \mathbf{x} \in \mathbb{R}^{n_{\mathbf{x}}} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, c_i(\mathbf{x}) \leq 0, i = 1, \dots, n_{\text{sc}} \}. \end{aligned} \quad (1.1)$$

The objective  $f : \mathbb{R}^{n_{\mathbf{x}}} \mapsto \mathbb{R}$  and nonlinear constraints  $\mathbf{c} : \mathbb{R}^{n_{\mathbf{x}}} \mapsto \mathbb{R}^{n_{\text{sc}}}$  encode desired lattice properties through the simulation-based evaluation of a lattice defined by the  $n_{\mathbf{x}}$  decision variables  $\mathbf{x}$ , which represent the lengths, strengths, and positions of the lattice elements: quadrupoles, dipoles, radio-frequency cavity (RF) inserts, and nonlinear (NL) inserts. Describing this problem with an optimization model provides a geometrically meaningful formulation of the decision space and an understandable parametric description for the trade-offs between desired lattice properties. However, with the benefits of an interpretable problem description come difficulties in finding points that are feasible for the problem. This manuscript seeks to serve as a guide for future RCS lattice designs by describing the construction and solution of this simulation-based optimization model as well as obstacles in solving the optimization problem.

We describe a principled search of the feasible set  $\Omega$ . We are especially careful to distinguish constraint functions that are algebraically available from those that depend on simulation output. The use of Synergia, an accelerator modeling framework, is key in this study because it allows for rapid simulation of a lattice design through linearized optics. From this basic framework of linear accelerator optics, particle accelerator performance can be assessed theoretically with additional considerations for adiabatic changes, machine errors,

coupling effects, nonlinearities, parametric resonances, particle interactions, feedback systems, and operational requirements. For example, Wei [54] gives an overview of design considerations for intense hadron synchrotron accelerators.

As a case study, we develop and solve an optimization model for an RCS compatible with nonlinear integrable optics. Nonlinear integrable optics is a recent particle accelerator technology that enables strong nonlinear focusing without generating new parametric resonances [10]. A promising application of integrable optics is in high-intensity rings, where nonlinearities are known to suppress the formation of beam halos [15, 16, 52] and enhance Landau damping of charge-dominated collective instabilities [31]. Over the next several years, the application of nonlinear integrable optics for intense beams will be studied experimentally at the Fermilab Accelerator Science and Technology Integrable Optics Test Accelerator (IOTA) [5, 50], as well as smaller-scale test accelerators at the University of Maryland [45] and at the Rutherford Appleton Laboratory [32].

This paper is structured as follows. Section 2 details the RCS lattice design problem and describes the particular synchrotron features desired in the case study considered in this paper. Section 3 details the use of Synergia [3, 4] to simulate these features. Section 4 formulates a mathematical model of the decision space  $\Omega$ . Through algebraic reductions, we are able to reduce our optimization model to a minimal set of 55 linear constraints and 23 simulation-based constraints involving 32 decision variables. Section 5 and Section 6 respectively address two of the key challenges in solving eq. (1.1): finding a feasible design in  $\Omega$  and tackling nondifferentiability in objective functions of interest. In both cases we take advantage of known dependence of the constraint and objective functions on various Synergia outputs, which allows us to find and certify locally optimal solutions. Section 7 shows the results of our numerical investigation and explores the dependence of our solutions on problem parameters. Section 8 concludes the paper with a look to future work.

## 2 Rapid-cycling synchrotron design case study

For the Fermilab 2.4 MW RCS upgrade, using an RCS lattice design compatible with nonlinear integrable optics has been considered in order to enhance the capabilities of the proton facility, reduce technical risk, and dramatically reduce costs associated with constructing the superconducting linac injector [14, 15]. It is also important to demonstrate that the additional constraints associated with nonlinear integrable optics do not adversely affect conventional lattice design criteria.

An overview of fundamental particle accelerator dynamics and notation is given in [13] and [29]. Particle accelerators maintain the stability of particle beams by providing focusing in all three planes of motion. Longitudinal focusing is provided by RF resonating cavities, which accelerate particles in the beam selectively depending on their arrival time. In the transverse planes,

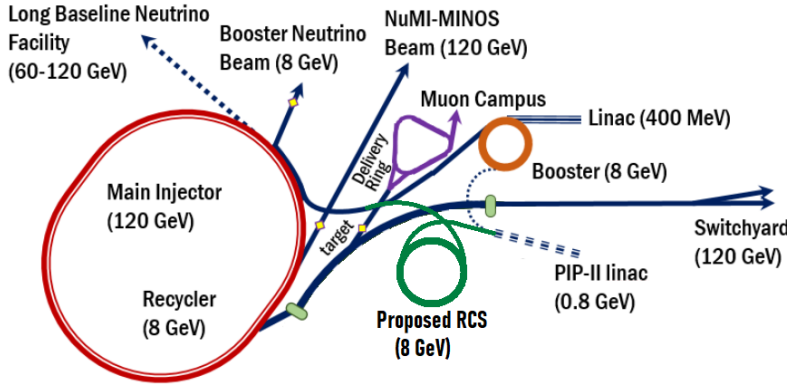
steering is provided by dipole magnets, and linear (de)focusing is provided by quadrupole magnets.

The linear transverse “optics” of the particle accelerator is described by Hill’s equations. The piecewise explicit time-dependence of Hill’s equations represent the passage of the particle beam through the sequence of focusing magnets known as the accelerator lattice. Hill’s equations are solved in their canonical form by Floquet’s theorem, which gives rise to optics functions including the amplitude-like beta function and phase advance  $\phi$  for each plane of motion. The transverse particle motion is known as a “betatron oscillation,” and the number of oscillations in one revolution around a particle accelerator ring is known as the “betatron tune.” Particle dynamics in an accelerator can also be represented by a sequence of transfer matrices, which can be thought of as mappings between snapshots of solutions to Hill’s equations at specific locations. The optics functions can be derived as properties of these matrices. Physical and engineering constraints, such as the length and strength of components, can be related to the linear accelerator optics only through the interaction of many independent accelerator magnets; consequently, nonlinear optimization methods are relied on to explore the decision space.

In application, the transverse optics are defined for a particle at a reference momentum (which can change as the beam accelerates), and additional terminology is needed to describe the relative motion of particles at other momentums. The particle trajectory that requires no restorative focusing force is known as the beam “orbit,” and the “dispersion” is defined to be the change in beam orbit with respect to beam momentum. The change in the betatron tune with beam momentum is known as “chromaticity.” The “momentum compactor factor” describes the change in the path length with momentum, aggregated over one revolution around a particle accelerator ring.

In [10] the “Danilov–Nagaitsev” criteria for integrable accelerator design first require an alternating sequence of linear and nonlinear sections. The linear sections, referred to as T-inserts, are a sequence of dipole and quadrupole magnets with an overall  $\pi$ -integer betatron phase advance in the horizontal and vertical plane. In the nonlinear sections, the lattice should be dispersion-free, and the horizontal and vertical beta functions should be equal to each other throughout. The manipulation of the beta functions and phase advances removes the time dependence of the nonlinear kick so as to avoid introducing parametric resonances. In [51, 53], the horizontal and vertical chromaticity should also be equal to maintain integrability for off-momentum particles to their lowest-order approximation.

The authors of [16] and [17] present accelerator lattices that combine Danilov–Nagaitsev integrable design criteria with features specific to the Fermilab RCS application. Those lattices were produced with a combination of Nelder–Mead simplex optimization [37] and manual manipulation; they serve as a platform for simulation experiments of nonlinear optics. However, they are bespoke lattices that do not represent a thorough or systematic investigation of the optimization landscape. The lattice given in [17] was used as a starting point for the optimization in this paper.

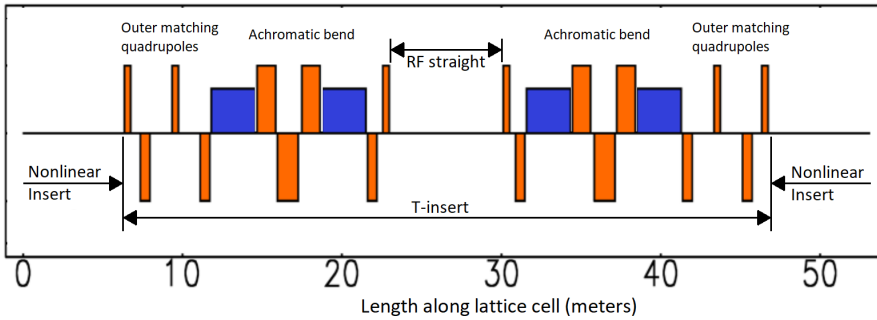


**Fig. 2.1** PIP-II and LBNF upgrades shown in dashed blue. Proposed RCS upgrade shown in green. Main injector ring (connecting RCS to LBNF) shown in red. Adapted from [48].

In addition to the Danilov–Nagaitsev criteria for integrable optics, previous RCS lattice designs were shown to fulfill conventional accelerator optics criteria that are specific to the Fermilab RCS application. The 637.0468718545753 m circumference of the integrable RCS design was not varied in our optimization; this RCS circumference was previously chosen to satisfy requirements for filling the downstream proton accelerator known as the Fermilab Main Injector. The Main Injector accumulates beam from the RCS, accelerates it to higher energies, and then delivers it to the LBNF beamline. The large circumference (relative to the Fermilab Booster) also arises as a result of the requirements for a high number of periodic cells, low-momentum compaction factor, and dispersion-free straights. Increasing the number of periodic cells (around the ring) to twelve was found to improve the performance of the nonlinear integrable optics with intense space-charge by more than a factor of 2. The low-momentum compaction factor ( $< 5.5 \times 10^{-3}$ ) is necessary to avoid a loss of longitudinal focusing when accelerating the beam from  $\approx 1$  GeV to  $\approx 8$  GeV. The dispersion-free straights eliminate transverse-longitudinal coupling from RF focusing as well as simplify injection and extraction optics.

The RCS dipole magnets are limited to a bending radius of  $\rho = 21.2$  m, equivalent to a field strength of 1.4 T at 8 GeV, comparable to the Fermilab Main Injector dipoles at 120 GeV. The RCS quadrupole magnets are limited to a focusing strength of  $K_1 = 1.2$  m $^{-2}$ , a pole radius and pole field comparable to Fermilab Main Injector dipoles at 120 GeV. A minimum of 0.2 m between elements of the magnetic lattice is required to allow clearance for the copper coil winding and vacuum flanges. These design considerations arise as constraints in our optimization model presented in Section 4.

In [15] it is argued that the maximum values of the beta functions (horizontal and vertical) play a direct role in the cost and performance of an RCS design. The size of the beam is proportional to the square root of the beta functions (neglecting dispersive effects) and the beam emittance (the phase



**Fig. 2.2** Lattice structure for each of 12 RCS cells. Unoptimized location and length of magnetic lattice elements with dipoles shown as short blue rectangles and quadrupoles shown as tall orange rectangles.

area occupied by the beam). Consequently the beta function determines the minimum diameter of a round beampipe aperture, which is a major consideration in a magnet cost. The physical size of accelerator magnets generally scales quadratically with the diameter of the aperture (transversely, to capture the return flux) or cubically (longitudinally, to maintain the same integrated field). On the other hand, within a given beampipe aperture, a smaller beta function can allow for a proportionately larger beam emittance, which in turn weakens deleterious space-charge effects. Consequently the maximum value taken by the beta functions over the beamline is considered as an objective function for our optimization model.

As shown in [Figure 2.2](#), the overall lattice structure of the ring is 12 identical lattice cells, each cell composed of two double-bend achromatic arcs that separate two dispersion-free straight sections of unequal lengths. The straight section in the “center” of the cell is for conventional accelerator components—RF accelerating cavities, injection chicane, extraction septa, and collimators. The straight section in the “outer” region of the cell is for the nonlinear optics inserts, with all other parts of the cell forming the linear T-insert. Each achromatic arc is a quadrupole triplet between two dipole bends and is tightly constrained by the dispersion-matching, low-momentum compaction factor, quadrupole strength, and overall compactness. A simple quadrupole doublet conveys the beam through the “RF” straight section. On the outer edges of the cell, at least a quadrupole triplet is needed to convey the beam through the nonlinear straight with matching beta functions; and a fourth quadrupole is included to allow additional finesse between the phase-advance, chromaticity, and maximum beta constraints.

The phase advance over the nonlinear insert is a parameter closely related to the achievable nonlinear tune shift with amplitude as well as the sensitivity of the nonlinear insert to various errors. The IOTA facility, an experimental ring for beam dynamics research [\[5\]](#), uses a phase advance across the nonlinear insert of  $0.3 \times 2\pi$ . The RCS application is focused on robust performance and may allow for a smaller phase advance.

### 3 RCS simulation with Synergia

Candidate RCS designs are evaluated by the accelerator modeling framework Synergia [3,4], which was developed to simulate both extant and proposed accelerator designs, to understand observed behavior in operational accelerators, to evaluate designs for proposed upgrades, and to probe beam physics effects [31]. Synergia combines the two major facets of accelerator physics: single particle optics and beam collective effects. To first order, particle transport in a beamline is analogous to light propagation in an optical system, thus the appellation “beam optics.” Collective effects arise due to interactions among all the charged particles in a beam bunch, primarily space charge, which is mediated by electromagnetic fields. Synergia has the capability to track bunches of particles and evaluate resulting electromagnetic fields using field solvers and applying resulting forces to the particle motions. For this study we have neglected collective effects in order to achieve a first optimization of the RCS design, but the capability in Synergia leaves open the possibility of considering them in a more detailed future study.

The particle motion in a periodic focusing system is governed locally by standard equations of motion but can be globally characterized by solutions to Hill’s equations. These are a set of linear differential equations that depend on the (de)focusing strength as a function of position. The focusing strength of an accelerator is determined by beamline element strengths and positions. These are the decision variables that define the candidate designs described below. The solutions to Hill’s equations consist of beta functions and phase advances (optical functions) that appear in the objectives and constraints of our current problem. Synergia derives the optical functions from the transfer map calculated by propagating particle coordinates through the beamline elements using the equations of motion of classical mechanics. Automatic differentiation of the coordinate vectors produces the maps [34] that are analyzed to determine stability conditions, beta functions, and phase advances. There are values for beta functions and phase advances for each beamline element. As described in Section 2, these quantities make up the bulk of typical objective and constraint functions.

Synergia is a C++ library for accelerator simulations utilizing models of physical accelerator components and charged particles organized in trains of beam bunches along with the machinery to apply the appropriate physical calculations of particle propagation to the C++ classes and objects. The computational libraries and data structures are organized so that parallel computations proceed efficiently on current and future platforms. The version used for this report uses OpenMPI and OpenMP for parallelism, which is needed when simulating collective effects among bunches containing  $10^5$ – $10^7$  particles. For flexibility and user accessibility, the core objects and methods required to perform a simulation have Python bindings provided by Boost.Python. Although simulations may be written in C++, most end users perform Synergia simulations with scripts written in Python. Selected classes of user importance may also be extended, including in Python.

This computational architecture proved ideal for this case study. Candidate RCS designs generated by the optimizer were passed to a Python script, which created the model beamline as Synergia objects using wrapped C++ methods. Synergia primitives were called to calculate beta functions and other characteristics needed to evaluate the candidate RCS designs. The results of the evaluation were returned to the optimizer as a Python dictionary.

#### 4 Mathematical model for the RCS case study

We now formulate a mathematical optimization problem based on the case study outlined in [Section 2](#). In particular, we optimize the performance of an integrable lattice by modifying the positions, lengths, and strengths of electromagnetic elements. Below we detail the key features of the problem: two-sided bound constraints and sequential linear constraints on the elements' positions, lengths, and strengths encoding engineering requirements and nonlinear black-box simulation-based constraints encoding desirable physical properties of the beamline.

##### 4.1 Decision variables

An RCS lattice is defined by the placement of various elements around a ring. We construct a ring by placing four types of elements (dipoles, quadrupoles, RF inserts, and NL inserts) on a **half-cell**, a  $\frac{1}{24}$ th slice of a full ring. This half-cell is then reflected once (see [Figure 2.2](#)) and repeated twelve times to form the periodic ring.

We let  $x_{e,p,i}$ ,  $x_{e,l,i}$ , and  $x_{e,s,i}$  denote the position, length, and strength, respectively, for the  $i$ th element of type  $e$ . We consider  $e \in \{d, q, r, n\}$ , corresponding to dipole, quadrupole, RF-insert, and NL-insert element types. In each half-cell, we place  $n_q = 9$  quadrupoles and  $n_d = 2$  dipoles along with a single RF insert and a single NL insert in the fixed ordering

$$(n, q, q, q, q, d, q, q, q, d, q, q, r), \quad (4.1)$$

as illustrated in [Figure 2.2](#).

The position (in meters) defines the center of each element. The length (in meters) is the length of the element; and thus the starting position of element  $e$  is  $x_{e,p,i} - 0.5x_{e,l,i}$ , and its end position is  $x_{e,p,i} + 0.5x_{e,l,i}$ . The quadrupole strength is a scaled focusing strength given in radians/m<sup>2</sup>, and the dipole strength is specified by a bend angle in radians. For this optimization, the effects of the RF-insert and NL-insert elements are not simulated, and so they contribute to the lattice design only through the space that they occupy and how the spacing interacts with the constraints on other elements.

We collect all decision variables in the ( $n_{\mathbf{x}} = 3n_q + 2n_d + 1 = 32$ )-dimensional vector  $\mathbf{x}$  with components

$$(x_{q,l,i}, x_{q,p,i}, x_{q,s,i})_{i=1}^{n_q}, (x_{d,l,i}, x_{d,p,i})_{i=1}^{n_d}, x_{d,s,1}. \quad (4.2)$$



Notably, the decision variables we work with do not include five terms ( $x_{n,p}$ ,  $x_{n,l}$ ,  $x_{r,p}$ ,  $x_{r,l}$ , and  $x_{d,s,2}$ ), which are eliminated from our formulation by virtue of additional design considerations discussed next. We assume the NL insert starts at the beginning of the half-cell:

$$x_{n,p} - 0.5x_{n,l} = 0. \quad (4.3)$$

The end of the NL insert coincides with the start of the first quadrupole:

$$x_{q,p,1} - 0.5x_{q,l,1} - (x_{n,p} + 0.5x_{n,l}) = 0. \quad (4.4)$$

There is a fixed circumference  $C = 637.0468718545753$ :

$$x_{r,p} + 0.5x_{r,l} = \frac{C}{24}. \quad (4.5)$$

Since the ring must close, the bend angle in a half-cell must be  $\pi/12$ :

$$x_{d,s,1} + x_{d,s,2} = \frac{\pi}{12}; \quad (4.6)$$

and the start of the RF insert must coincide with the end of the last quadrupole:

$$x_{r,p} - 0.5x_{r,l} - (x_{q,p,9} + 0.5x_{q,l,9}) = 0. \quad (4.7)$$

These linear equalities allow us to eliminate the terms  $x_{n,p}$ ,  $x_{n,l}$ ,  $x_{r,p}$ ,  $x_{r,l}$ , and  $x_{d,s,2}$ , which can be recovered via

$$\begin{aligned} x_{n,p} &= 0.5(x_{q,p,1} - 0.5x_{q,l,1}) \\ x_{n,l} &= x_{q,p,1} - 0.5x_{q,l,1} \\ x_{r,p} &= \frac{C}{24} - 0.5 \left( \frac{C}{24} - x_{q,p,9} - 0.5x_{q,l,9} \right) \\ x_{r,l} &= \frac{C}{24} - x_{q,p,9} - 0.5x_{q,l,9} \\ x_{d,s,2} &= \frac{\pi}{12} - x_{d,s,1}. \end{aligned}$$

Following the taxonomy in [27], we separate constraints on the remaining  $\mathbf{x}$  into simulation-based constraints (i.e., those for which (in)feasibility can be verified without a Synergia call) and algebraic constraints (i.e., those that do not require a Synergia call).

## 4.2 Algebraic constraints

Our model's algebraic constraints consist of two-sided bound constraints and linear inequality constraints. Together, these algebraic constraints define a compact polytope in  $\mathbb{R}^{n_x}$ .

Two-sided bound constraints (with unequal lower and upper bounds) have the benefit of not contributing twice to the combinatorial complexity of an optimization problem: at most one of the sides can be active, and thus the two

bounds will not contribute to potentially violating a constraint qualification (e.g., LICQ) unless the objective's derivative with respect to a decision variable vanishes. We assume that the  $n_q + n_d$  lengths are bounded:

$$\begin{aligned} 0 \leq x_{q,l,i} \leq 5, & \quad i = 1, \dots, n_q \\ x_{d,l,i} \leq 5, & \quad i = 1, \dots, n_d. \end{aligned}$$

We also assume that the  $n_q + 1$  strengths are bounded. By noting that the lower bound on the dipoles can be tightened, since the dipole strengths must obey the bend angle constraint in eq. (4.6), we arrive at

$$\begin{aligned} -1.2 \leq x_{q,s,i} \leq 1.2, & \quad i = 1, \dots, n_q \\ \pi/12 - 0.2 \leq x_{d,s,1} \leq 0.2. \end{aligned}$$

The linear inequality constraints encode the fixed relative positioning of the elements in the order given in eq. (4.1) with spacing, as well as the reduced circumference and bend angle constraints. We have

$$\begin{aligned} x_{q,p,1} - 0.5x_{q,l,1} &\geq 1 \\ x_{q,p,i} - 0.5x_{q,l,i} - (x_{q,p,i-1} + 0.5x_{q,l,i-1}) &\geq 0.2, & i = 2, \dots, 4 \\ x_{d,p,1} - 0.5x_{d,l,1} - (x_{q,p,4} + 0.5x_{q,l,4}) &\geq 0.2 \\ x_{q,p,5} - 0.5x_{q,l,5} - (x_{d,p,1} + 0.5x_{d,l,1}) &\geq 0.2 \\ x_{q,p,i} - 0.5x_{q,l,i} - (x_{q,p,i-1} + 0.5x_{q,l,i-1}) &\geq 0.2, & i = 6, 7 \\ x_{d,p,2} - 0.5x_{d,l,2} - (x_{q,p,7} + 0.5x_{q,l,7}) &\geq 0.2 \\ x_{q,p,8} - 0.5x_{q,l,8} - (x_{d,p,2} + 0.5x_{d,l,2}) &\geq 0.2 \\ x_{q,p,9} - 0.5x_{q,l,9} - (x_{q,p,8} + 0.5x_{q,l,8}) &\geq 0.2. \end{aligned}$$

We note that the first constraint combines eq. (4.3) with eq. (4.4) and a lower bound of 1 on the NL-insert length.

We include two additional linear inequalities that combine the bounds [3, 3.75] on the RF-insert lengths, the fixed circumference constraint eq. (4.5), and the spacing constraint eq. (4.7) between the last quadrupole and RF insert:

$$x_{q,p,9} + 0.5x_{q,l,9} \leq \frac{C}{24} - 3 \quad (4.8)$$

$$-x_{q,p,9} - 0.5x_{q,l,9} \leq 3.75 - \frac{C}{24}. \quad (4.9)$$

Lastly, we apply two linear constraints to proportionally lower bound the dipole lengths,  $\rho x_{d,s,i} \leq x_{d,l,i}$  for  $i = 1, \dots, n_d$ . By using the change of variables defined in eq. (4.6), we have

$$\rho x_{d,s,1} - x_{d,l,1} \leq 0 \quad (4.10)$$

$$-\rho x_{d,s,1} - x_{d,l,2} \leq -\frac{\rho\pi}{12}, \quad (4.11)$$

where  $\rho = 21.2$  meters.

These algebraic constraints (i.e., the  $2n_q + 1$  two-sided bound constraints,  $n_d$  one-sided bound constraints, and the 15 other linear inequalities) can be expressed compactly by  $\mathbf{Ax} \leq \mathbf{b}$ , where  $\mathbf{A}$  is of size  $n_{ac} = 4n_q + 19$  by  $n_{\mathbf{x}} = 3n_q + 2n_d + 1$  (55 by 32), which we use for conciseness below.

For the remainder of this study we rescale the decision variables  $\mathbf{x}$  to the unit cube. The appropriate rescaling is by the side lengths of the smallest box that encloses the polytope  $\{\mathbf{x} \in \mathbb{R}^{n_{\mathbf{x}}} : \mathbf{Ax} \leq \mathbf{b}\}$ .

### 4.3 Simulation-based constraints

Numerous criteria (described in [Section 2](#)) are desired in an RCS lattice. In our model we opt for a constraint-based approach in which simulation-based constraints encode the majority of the physical design criteria for the RCS lattice described in [Section 2](#). With this approach a key concern is establishing joint feasibility across all the constraints; we are also interested in the dependence of solutions on the constraint parameters, which we study in [Section 6.2](#). In the remainder of this section we describe these  $n_{sc} = 23$  constraints, all of which require a Synergia call in order to verify (in)feasibility.

A fundamental constraint to the RCS design and mathematical model is the requirement that the lattice have a periodic orbit (and is therefore stable). Many quantities of interest are undefined for physically infeasible solutions where no periodic orbit exists. In those situations Synergia is unable to evaluate the candidate solution. Mathematically, the periodic orbit exists when the ‘‘one turn map’’  $M(\mathbf{x})$ , corresponding to direction of motion  $(x, y, z)$ , has eigenvalues of roughly unit norm. We constrain only the transverse directions  $x$  and  $y$  so we require two conforming eigenvalue pairs. We enforce the norm unity conditions with the two two-sided constraints  $c_1, c_2$  on the modulus of the single relevant eigenvalue  $\lambda(\mathbf{x})$  of  $M(\mathbf{x})$ :

$$c_1(\mathbf{x}) := |\lambda(\mathbf{x})| - 1 - \epsilon_\lambda \leq 0 \quad (4.12)$$

$$c_2(\mathbf{x}) := -|\lambda(\mathbf{x})| + 1 - \epsilon_\lambda \leq 0, \quad (4.13)$$

with  $\epsilon_\lambda = 10^{-10}$ . In practice, we find that this constraint is always satisfied at  $\mathbf{x}$  that are feasible with respect to all other constraints.

The next constraint enforces that the momentum compaction factor  $\alpha_c(\mathbf{x})$  should be at most  $5.5 \times 10^{-3}$ :

$$c_3(\mathbf{x}) := \alpha_c(\mathbf{x}) - 5.5 \times 10^{-3} \leq 0. \quad (4.14)$$

The RCS design seeks a dispersion-free RF insert, which in practice is relaxed by a numerical tolerance  $10^{-3}$ :

$$c_4(\mathbf{x}) := D_{x,\text{rf}}(\mathbf{x}) - 10^{-3} \leq 0 \quad (4.15)$$

$$c_5(\mathbf{x}) := -D_{x,\text{rf}}(\mathbf{x}) - 10^{-3} \leq 0. \quad (4.16)$$

Similarly, constraints for dispersion-free NL inserts are relaxed by a numerical tolerance  $10^{-5}$ :

$$c_6(\mathbf{x}) := D_{x,\text{nl}1}(\mathbf{x}) - 10^{-5} \leq 0 \quad (4.17)$$

$$c_7(\mathbf{x}) := -D_{x,\text{nl}1}(\mathbf{x}) - 10^{-5} \leq 0 \quad (4.18)$$

$$c_8(\mathbf{x}) := D_{x,\text{nl}2}(\mathbf{x}) - 10^{-5} \leq 0 \quad (4.19)$$

$$c_9(\mathbf{x}) := -D_{x,\text{nl}2}(\mathbf{x}) - 10^{-5} \leq 0. \quad (4.20)$$

To ensure sufficient similarity of the beta functions over the NL insert, we enforce the four constraints

$$c_{10}(\mathbf{x}) := \beta_{y,1}(\mathbf{x}) - (1 + \epsilon_\beta)\beta_{x,1}(\mathbf{x}) \leq 0 \quad (4.21)$$

$$c_{11}(\mathbf{x}) := -\beta_{y,1}(\mathbf{x}) + (1 - \epsilon_\beta)\beta_{x,1}(\mathbf{x}) \leq 0 \quad (4.22)$$

$$c_{12}(\mathbf{x}) := \beta_{y,2}(\mathbf{x}) - (1 + \epsilon_\beta)\beta_{x,2}(\mathbf{x}) \leq 0 \quad (4.23)$$

$$c_{13}(\mathbf{x}) := -\beta_{y,2}(\mathbf{x}) + (1 - \epsilon_\beta)\beta_{x,2}(\mathbf{x}) \leq 0, \quad (4.24)$$

where we use  $\epsilon_\beta = 0.01$  and where  $\beta_{\cdot,1}$  indicates the beta function value at the start of the NL insert and  $\beta_{\cdot,2}$  indicates its value at the end of the NL insert.

We enforce integrability constraints by using the modulo operator to ensure that the T-insert phase advances  $\Phi_x$  and  $\Phi_y$  are multiples of  $\pi$ . We apply a shift of  $\pi/2$  to move away from the discontinuity in the  $\text{mod}(\pi)$  function:

$$c_{14}(\mathbf{x}) := (\pi/2 + \Phi_x(\mathbf{x})) \bmod(\pi) - (\pi/2 + \epsilon_\phi) \leq 0 \quad (4.25)$$

$$c_{15}(\mathbf{x}) := -(\pi/2 + \Phi_x(\mathbf{x})) \bmod(\pi) + (\pi/2 - \epsilon_\phi) \leq 0 \quad (4.26)$$

$$c_{16}(\mathbf{x}) := (\pi/2 + \Phi_y(\mathbf{x})) \bmod(\pi) - (\pi/2 + \epsilon_\phi) \leq 0 \quad (4.27)$$

$$c_{17}(\mathbf{x}) := -(\pi/2 + \Phi_y(\mathbf{x})) \bmod(\pi) + (\pi/2 - \epsilon_\phi) \leq 0, \quad (4.28)$$

where we use the tolerance  $\epsilon_\phi = 10^{-4}$ .

We also seek an RCS design where the ring-wide betatron tunes  $\nu_x(\mathbf{x}), \nu_y(\mathbf{x})$  are sufficiently close. Here we use a tolerance of  $10^{-4}$  and thus have

$$c_{18}(\mathbf{x}) := \nu_x(\mathbf{x}) - \nu_y(\mathbf{x}) - 10^{-4} \leq 0 \quad (4.29)$$

$$c_{19}(\mathbf{x}) := -(\nu_x(\mathbf{x}) - \nu_y(\mathbf{x})) - 10^{-4} \leq 0. \quad (4.30)$$

The phase advance over the NL inserts  $\Psi_x$  and  $\Psi_y$  should be at least  $0.2 \times 2\pi$ :

$$c_{20}(\mathbf{x}) := -\Psi_x(\mathbf{x}) + 0.2 \times 2\pi \leq 0 \quad (4.31)$$

$$c_{21}(\mathbf{x}) := -\Psi_y(\mathbf{x}) + 0.2 \times 2\pi \leq 0. \quad (4.32)$$

The horizontal and vertical chromaticity should differ by no more than 0.1 globally:

$$c_{22}(\mathbf{x}) := C_x(\mathbf{x}) - C_y(\mathbf{x}) - 0.1 \leq 0 \quad (4.33)$$

$$c_{23}(\mathbf{x}) := -(C_x(\mathbf{x}) - C_y(\mathbf{x})) - 0.1 \leq 0. \quad (4.34)$$

Collectively, we denote these 23 simulation-based constraints by  $\mathbf{c}(\mathbf{x}) \leq \mathbf{0}$ , so that nonpositive values indicate feasibility.

With a complete description of the decision variables and constraints in hand, we have a mathematical description of our feasible decision space, which aims to capture RCS lattice designs satisfying the intent of [Section 2](#). In the following section we summarize challenges for finding feasible decision points and establish that such feasible decisions exist.


## 5 Phase-one solution: finding a feasible point

Most optimization approaches to problems of the form [eq. \(1.1\)](#) benefit from being initialized with a feasible starting point. Nonlinear simulation-based constraints can make finding such a point difficult, and “phase-one” (i.e., feasibility seeking) optimization approaches may not cope well with the nonlinear, black-box form of the constraints. Finding a feasible point is a key challenge of this problem since it validates (up to simulation fidelity) the existence of an RCS lattice with the desired properties outlined in [Section 2](#), which has until now not been shown to exist. In our search for a feasible point we take advantage of the polytope structure in our algebraic constraints (since this is available without querying the simulation) and a solution that is known to satisfy some (but not all) of the constraints from previous synchrotron design problems.

The polytope defined by  $\mathbf{Ax} \leq \mathbf{b}$  bounds the feasible region  $\Omega$ , and thus sampling within the polytope may appear to be a promising method of finding feasible points. Approaches for random sampling from the polytope include taking convex combinations of the polytope vertices or employing Markov chain Monte Carlo (MCMC) methods such as hit-and-run sampling or the Vaidya walk [\[9, 33\]](#). MCMC methods require only an interior point to the polytope. Although vertex enumeration is difficult in general, it is tractable in our case because our linearly constrained polytope has some dimensions (e.g., the element strengths) that appear only in bound constraints. This allows an easier enumeration of vertices, for example, by enumerating the vertices of the bound-constrained region separately from the other vertices of the linearly constrained region [\[8, 21\]](#).

[Figure 5.1](#) illustrates that taking uniformly selected convex combinations of vertices is not a useful approach in finding feasible points in our case study: in 67,810 samples, not a single point was feasible. Of the 67,810 points sampled, roughly 90% resulted in unsuccessful simulation evaluations (i.e., because the stability conditions  $c_1, c_2$  were not satisfied). [Figure 5.1](#) shows that even the successfully evaluated points still tend to violate most simulation-based constraints. Furthermore, MCMC methods were observed to have an even higher fraction of unsuccessful simulation evaluations than sampling convex combinations of vertices have. We conclude that the points that are feasible for the simulation-based constraints are “needles in the haystack” of the points satisfying the linear constraints.


On the other hand, a phase-one optimization procedure initialized at an (infeasible) “expert-designed” point, denoted  $\mathbf{e}_0$ , was successful in finding



**Fig. 5.1** Fraction of 6,182 successfully simulated points randomly sampled within the polytope at which each simulation-based constraint (or both sides of a two-sided constraint pair) is satisfied. The two-sided constraint pairs  $(c_{18}, c_{19})$ ,  $(c_{16}, c_{17})$ , and  $(c_{14}, c_{15})$  are never satisfied.

a feasible point. The point  $\mathbf{e}_0$  was lacking for two reasons. First, it violated algebraic constraints: the spacings between Dipole 2 and Quadrupole 7, Dipole 2 and Quadrupole 8, Dipole 1 and Quadrupole 4, Dipole 1 and Quadrupole 5, and Quadrupoles 5 and Quadrupole 6 had a separation distance of 0.1 m rather than the desired 0.2 m. Second, the simulation-based constraints  $c_7, c_9, c_{14}, c_{16}, c_{19}$  were violated; see Figure 7.3. To ameliorate these issues, we solved two optimization problems: one to move into the algebraic-constraint-defined polytope while not worsening violation of the simulation-based constraints and a subsequent one to improve the violation of the simulation-based constraints. Working with simulation-based constraints can be challenging because of discontinuities or numerical noise in a constraint function, especially when these are present when a constraint is nearly active (see, e.g., [25, Section 7], [6, Chapter 12]). In our problem, we found that the simulation outputs were remarkably smooth; for example, as shown in Figure 5.2, we see that numerical noise is not readily apparent for changes in  $\mathbf{x}$  larger than  $10^{-10}$ . This is expected since the system being simulated is a smooth physical system with forces that vary only linearly with particle amplitudes, which should not exhibit chaotic behavior with small deviations.

Starting from  $\mathbf{e}_0$  and letting  $\mathbf{A}_i$  indicate the  $i$ th row of  $\mathbf{A}$ , we minimize the algebraic constraint violation  $\sum_{i=1}^{n_{ac}} \max\{\mathbf{A}_i \mathbf{x} - b_i, 0\}^2$  while constraining the simulation-based constraints and algebraic constraints that hold at  $\mathbf{e}_0$  to continue to hold and constraining the simulation-based constraints that are violated at  $\mathbf{e}_0$  by the initial value  $c_i(\mathbf{e}_0)$  to prevent them from worsening. A globally optimal solution to this problem (i.e., satisfying the stated constraints and all algebraic constraints) was found by using the COBYLA local optimization routine [42, 43].



**Fig. 5.2** Examples of smoothness in simulation outputs momentum compaction  $\alpha_c$  (blue), NL-Dispersion  $D_{x,nl1}$  (orange) and betatron tunes  $\nu_x - \nu_y$  (green) for changes in  $\mathbf{x}$  smaller than  $10^{-9}$  along a ray in parameter space. The decision vectors  $\mathbf{x}$  and  $\mathbf{x}_0$  are on the unit cube. Notably, Synergia shows high fidelity under small changes in the decision variables  $\mathbf{x}$ , free of any significant noise to roughly a  $10^{-10}$  change in  $\mathbf{x}$ . In the vicinity of feasible solutions there are no nonlinearities in the optics that would lead to chaotic behavior.


From this point,  $\mathbf{e}_1$ , in the polytope we initialize a second optimization problem in order to find simulation-based constraint satisfaction. We minimize the simulation-based constraint violation  $\sum_{i=1}^{n_{sc}} \max\{c_i(\mathbf{x}), 0\}^2$  while constraining the simulation-based constraints and algebraic constraints that hold at  $\mathbf{e}_1$  to continue to hold and constraining the simulation-based constraints  $c_i$  that are violated at  $\mathbf{e}_1$  by their initial value  $c_i(\mathbf{e}_1)$  to prevent their infeasibility from worsening. We also constrain the values of all components of  $\beta_x$  and  $\beta_y$  by the nominal value of 50 so as to not unnecessarily affect objective function values (discussed in the next section) when searching for a feasible point. COBYLA solved this problem to global optimality (i.e., finding a point in  $\Omega$  that also satisfied constraints imposed on  $\beta_x$  and  $\beta_y$ ) in roughly 60,000 function evaluations.

A key to the success of the phase-one optimization was starting the feasibility restoration at an expert-designed point that was relatively close to the feasible region  $\Omega$ . In part because of the nonconvexity of the simulation-based constraints, we found that beginning phase-one from points with significant simulation-based constraint violation did not produce feasible solutions.

## 6 Nondifferentiable composite objective functions and local optimality

Finding points satisfying all of the constraints posed in Section 4 allows us to consider minimization of the beta functions. A minimal beta profile is desired because it indicates a high level of beam concentration.

Because the accelerator community does not have a single standard objective function, we investigate the solution of eq. (1.1) with three different objectives, consisting of different nonsmooth compositions of the simulation



**Fig. 6.1** Visible nonsmoothness of the objective defined by  $h_b$  in a random two-dimensional slice of  $\mathbb{R}^{n_{\mathbf{x}}}$  around an arbitrary feasible point. In this plot,  $h_b$  is defined by the active indexes  $\beta_{x,48}$  (bottom left),  $\beta_{x,180}$  (top right), and  $\beta_{y,111}$  (bottom right),  $\beta_{y,158}$  (top left).

outputs  $\beta_x(\mathbf{x})$  and  $\beta_y(\mathbf{x})$ . Expert solicitation indicated that a leading criterion is a low maximum in the beta functions. A second criterion is a symmetry in the peaks of the beta functions. Based on these criteria and a composite objective of the form  $f(\mathbf{x}) = h(\beta(\mathbf{x}))$ , we consider the following three compositions:

$$h_a(\beta(\mathbf{x})) = \max_j \{\beta_{x,j}(\mathbf{x})\} + \max_j \{\beta_{y,j}(\mathbf{x})\} = \|\beta_x(\mathbf{x})\|_{\infty} + \|\beta_y(\mathbf{x})\|_{\infty} \quad (6.1)$$

$$h_b(\beta(\mathbf{x})) = \max \left\{ \max_j \{\beta_{x,j}(\mathbf{x})\}, \max_j \{\beta_{y,j}(\mathbf{x})\} \right\} = \|\beta(\mathbf{x})\|_{\infty} \quad (6.2)$$

$$h_c(\beta(\mathbf{x})) = \max_j \{\beta_{x,j}(\mathbf{x}) + \beta_{y,j}(\mathbf{x})\} = \|\beta_x(\mathbf{x}) + \beta_y(\mathbf{x})\|_{\infty}. \quad (6.3)$$

Since the beta functions are nonnegative, we have that, for any  $\mathbf{x} \in \Omega$ ,  $h_b(\beta(\mathbf{x})) \leq h_c(\beta(\mathbf{x})) \leq h_a(\beta(\mathbf{x}))$ . Although smooth approximations to these objectives can be formulated, we directly employ these nonsmooth forms in order to preserve the interpretability of each objective form.

Figure 6.1 illustrates that the resulting objective functions readily exhibit nondifferentiable behavior in  $\Omega$ .

### 6.1 Manifold sampling from feasible points

To address the structured nondifferentiability illustrated in Figure 6.1, we employ a manifold sampling trust-region method [23, 24, 26]. Since the nonsmoothness occurs because of a composition of a known mapping  $h$  with a (smooth) simulation output  $\beta(\mathbf{x})$ , manifold sampling can exploit the known form of each  $h$ . Manifold sampling is a model-based approach that categorizes points  $\beta(\mathbf{x})$  in the domain of  $h$  as belonging to different manifolds that occur due to the nonsmoothness in  $h$ ; this information is used when determining search directions [23, 24, 26].

Instead of enforcing all 78 inequality constraints, we use a quadratic penalty formulation [40], where we penalize the square of the constraint violation  $[\cdot]_+ =$



$\max\{\cdot, 0\}$  through minimization of the unconstrained objective

$$f(\mathbf{x}) = h(\boldsymbol{\beta}(\mathbf{x})) + \alpha \sum_{i=1}^{n_{sc}} [c_i(\mathbf{x})]_+^2 + \alpha \sum_{i=1}^{n_{ac}} [\mathbf{A}_i \mathbf{x} - b_i]_+^2 \quad (6.4)$$

for the various nonsmooth definitions of  $h$  above. A sequence of solutions of eq. (6.4) as  $\alpha \rightarrow \infty$  converge to a Karush—Kuhn—Tucker (KKT) point of eq. (1.1), where  $f(\mathbf{x}) = h(\boldsymbol{\beta}(\mathbf{x}))$  provided that certain conditions are satisfied: feasibility of the limit point, appropriate constraint qualification, and sufficient accuracy of solutions of eq. (6.4) [40, Thm. 17.2].

Manifold sampling can address the fact that the  $f$  in eq. (6.4) is a nonsmooth function with additional potential nonsmoothness coming from constraint violations. For example, if the constraint and beta functions are continuously differentiable, the Clarke subdifferential of eq. (6.4) when  $h = h_b$  can be computed by taking the convex hull of the gradients of the active beta functions

$$\begin{aligned} \partial h_b(\boldsymbol{\beta}(\mathbf{x})) = & \operatorname{co}_{j \in J^*(\mathbf{x})} \{ \nabla \beta_j(\mathbf{x}) \} \\ & + 2\alpha \left( \sum_{i=1}^{n_{sc}} [c_i(\mathbf{x})]_+ \nabla c_i(\mathbf{x}) + \sum_{i=1}^{n_{ac}} [\mathbf{A}_i \mathbf{x} - b_i]_+ \mathbf{A}_i^T \right), \end{aligned} \quad (6.5)$$

where  $J^*(\mathbf{x}) = \arg \max_j \{ \beta_j(\mathbf{x}) \}$ . It is precisely the knowledge of which manifolds (in this case, the indices of the beta functions) are active that allows manifold sampling to approximate this subdifferential and determine descent directions. Since the quantities  $\nabla_{\mathbf{x}} \beta_{x,\cdot}$  and  $\nabla_{\mathbf{x}} \beta_{y,\cdot}$  are unavailable, manifold sampling builds local models  $m_{\beta_j}$  of both beta functions at each of a finite number of indices  $j$  and uses  $\nabla_{\mathbf{x}} m_{\beta_j}$  in place of the corresponding (unavailable) gradients.

Manifold sampling notes what indices define the set  $J^*(\mathbf{x})$  at (or near) candidate points  $\mathbf{x}$  and uses this information in order to infer when the subdifferential of  $f$  changes. Given a set  $\mathbf{Y}$  of past points evaluated during the course of the algorithm (either iterates of the algorithm or from building models of the components of  $\boldsymbol{\beta}$ ) and a trust-region radius  $\Delta$ , the manifold sampling implementation used in our numerical tests approximated the set  $J^*(\mathbf{x})$  by

$$\bigcup_{\mathbf{y} \in \mathbf{Y}: \|\mathbf{y} - \mathbf{x}\| \leq \Delta} \left\{ j : \beta_j(\mathbf{y}) \geq \max_i \{ \beta_i(\mathbf{y}) \} - 10^{-8} \right\}.$$

That is, the indices that are considered to define  $h_b(\boldsymbol{\beta}(\mathbf{y}))$  are those where the max is within  $10^{-8}$  of the true maximum; we take the union of these sets for any point  $\mathbf{y}$  in the trust region around  $\mathbf{x}$ .

Such an approach to approximating the subdifferential at  $\mathbf{x}$  may seem excessive, but it is in fact essential for algorithmic performance. The logic of the manifold sampling algorithm requires that the manifolds active at any putative iterate be known before a step can be taken/rejected. This requires an inner set of manifold sampling iterations where the set of active manifolds

grows iteratively. This process is guaranteed to terminate because the number of indices that can be active is finite. Warm-starting the manifold sampling process with past information means the algorithm can make progress sooner.

## 6.2 Local optimality conditions

A key benefit of optimization models in the RCS design is providing insight into how the performance of the RCS design is dependent on tunable parameters. From a physical standpoint this brings interpretability to how design adjustments alter performance. We are particularly interested in the effect of parameters  $\boldsymbol{\varepsilon} \in \mathbb{R}^{n_{sc}}$  defining a feasible region based on a parameterized right-hand side of the simulation-based constraints:

$$\Omega_{\boldsymbol{\varepsilon}} = \{\mathbf{x} \in \mathbb{R}^{n_x} : \mathbf{Ax} \leq \mathbf{b}, c_i(\mathbf{x}) \leq \varepsilon_i, i = 1, \dots, n_{sc}\}; \quad (6.6)$$

$\varepsilon_i = 0$  corresponds to our nominal formulation, whereas a positive (negative)  $\varepsilon_i$  value corresponds to a relaxation (tightening) of the  $i$ th simulation-based constraint. We look at the KKT optimality conditions [40], which indicate how changes in the violation of active constraints locally improve or degrade performance in the objective. Solving the KKT conditions yields Lagrange multipliers  $\lambda_i$  for each constraint; these multipliers quantify how each parameter  $\varepsilon_i$  affects the optimal objective value. With the Lagrange multipliers in hand, an RCS designer can weigh tradeoffs between constraint violation and objective quality when considering different designs.

Assuming that the constraint and beta functions are continuously differentiable, the nonsmooth KKT conditions [7, 18] state that under a constraint qualification, such as LICQ, there exist a subgradient  $\mathbf{g}(\mathbf{x}) \in \partial h(\boldsymbol{\beta}(\mathbf{x}))$  and Lagrange multipliers  $\boldsymbol{\lambda} = (\boldsymbol{\lambda}^s, \boldsymbol{\lambda}^a)$  such that

$$\mathbf{g}(\mathbf{x}) + \sum_{i \in \mathcal{A}_{sc}(\mathbf{x})} \lambda_i^s \nabla c_i(\mathbf{x}) + \sum_{j \in \mathcal{A}_{ac}(\mathbf{x})} \lambda_j^a \mathbf{A}_j^T = 0 \quad (6.7)$$

$$c_i(\mathbf{x}) \leq \varepsilon_i, \quad i = 1, \dots, n_{sc} \quad (6.8)$$

$$\mathbf{Ax} \leq \mathbf{b} \quad (6.9)$$

$$\boldsymbol{\lambda}^s, \boldsymbol{\lambda}^a \geq \mathbf{0} \quad (6.10)$$

$$\lambda_i^s c_i(\mathbf{x}) = 0, \quad i = 1, \dots, n_{sc} \quad (6.11)$$

$$\lambda_j^a (\mathbf{A}_j \mathbf{x} - b_j) = 0, \quad j = 1, \dots, n_{ac}, \quad (6.12)$$

where  $\mathcal{A}_{ac}(\mathbf{x})$  and  $\mathcal{A}_{sc}(\mathbf{x})$  indicate the active sets of algebraic and simulation-based (i.e.,  $c_i(\mathbf{x}) = \varepsilon_i$ ) constraints, respectively. The stationarity condition eq. (6.7) provides us the handle for understanding parametric dependence since it expresses the objective gradient in terms of active constraint gradients. We can use this to approximate the objective value under changes to constraint values, up to linearization. For example, suppose there exists (in fact there does exist by LICQ) a direction  $\mathbf{d}$  that is orthogonal to all active constraint

gradients aside from, say,  $\nabla c_i(\mathbf{x})$ . Then, the change in objective value in the direction  $\mathbf{d}$  is approximately  $-\lambda_i^s \mathbf{d}^T \nabla c_i(\mathbf{x})$ .

When considering the objective  $h = h_b$ , the subgradients are defined by the convex combination  $\mathbf{g}(\mathbf{x}) = \sum_{j \in J^*(\mathbf{x})} \gamma_j \nabla \beta_j(\mathbf{x})$ , where  $\boldsymbol{\gamma} \geq \mathbf{0}$  satisfies  $\sum_j \gamma_j = 1$  and  $J^*(\mathbf{x}) = \arg \max_j \{\beta_j(\mathbf{x})\}$  indexes the set of active beta indices. When  $\mathbf{x}$  is optimal for eq. (1.1), the Lagrange multipliers  $\boldsymbol{\lambda}$  solve

$$\min_{\boldsymbol{\lambda}, \boldsymbol{\gamma}} \left\| \sum_{j \in J^*(\mathbf{x})} \gamma_j \nabla \beta_j(\mathbf{x}) + \sum_{i \in \mathcal{A}_{sc}(\mathbf{x})} \lambda_i^s \nabla c_i(\mathbf{x}) + \sum_{i \in \mathcal{A}_{ac}(\mathbf{x})} \lambda_i^a \mathbf{A}_i^T \right\|^2 \quad (6.13)$$

$$\boldsymbol{\lambda}, \boldsymbol{\gamma} \geq \mathbf{0} \quad (6.14)$$

$$\lambda_i^s (c_i(\mathbf{x}) - \varepsilon_i) = 0 \quad i = 1, \dots, n_{sc} \quad (6.15)$$

$$\lambda_i^a (\mathbf{A}_i \mathbf{x} - b_i) = 0 \quad i = 1, \dots, n_{ac} \quad (6.16)$$

$$\sum_{j \in J^*(\mathbf{x})} \gamma_j = 1. \quad (6.17)$$


We use this method to compute  $\boldsymbol{\lambda}$  in Section 7.3. In practice, we determine the active set of beta indices to be those functions  $\beta_j$  within  $10^{-8}$  of the max objective value. The active set of constraints is determined by first finding the steepest descent direction and then collecting the constraints that are violated when traveling in such a direction. In this way we find a sequential list of constraints that are limiting motion. The constraint right-hand sides are then adjusted so that the active constraints exactly equal zero at the optima.

## 7 Numerical results

Here we numerically explore properties of the optimization model eq. (1.1), described in Section 2 and developed in Section 4. We discuss the physical properties, physical validity, sensitivity, and effects of integrability at various local optima and compare them with unoptimized points. Furthermore, to understand the dependence of optima on constraint parameters, we analyze the KKT conditions and Lagrange multipliers, which provide a lens to elucidate the dependence of the objective value on the set of restrictive constraints. The principal finding is that the thorough development of an optimization model provides a powerful tool for finding high-performance RCS lattice designs that meet the desired criteria as well as analyzing the design criteria and their effect on solutions.

### 7.1 Studying different objectives

Since it is not immediately clear which beta profiles are preferable, we optimize all three composite objective functions  $h_a(\boldsymbol{\beta}(\mathbf{x}))$ ,  $h_b(\boldsymbol{\beta}(\mathbf{x}))$ ,  $h_c(\boldsymbol{\beta}(\mathbf{x}))$ , described in Section 6 subject to the constraints described in Section 4. The three



**Fig. 7.1** Beta functions at the local minima  $\mathbf{x}_a, \mathbf{x}_b, \mathbf{x}_c$  after manifold sampling optimization under the respective objectives  $h_a, h_b, h_c$ . The  $h_b$  and  $h_c$  objectives produce minima with nearly identical beta functions; the  $h_a$  minimizer also has a similar profile.


constrained optimization problems are reformulated with an unconstrained quadratic penalty method and solved by using the manifold sampling trust-region method [22]. Figure 7.1 shows the resulting beta profiles for the optima under each objective.

We find that optimization under the three objectives given the same initial point yields solutions that have notably similar beta profiles and correspondingly similar element positions, lengths, and strengths. While the accelerator design experts did not prefer one set of solutions over the other for their physical properties, they did prefer the solutions from  $h_b$  because of the physical interpretability of the  $\max_j(\beta_j(x))$  objective. For the remainder of the exposition, we focus on the objective  $h_b$  and the corresponding minimizer  $\mathbf{x}_b$  found, which is illustrated in Figure 7.2.


## 7.2 A preferred lattice

The beta profile for  $\mathbf{x}_b$  significantly improves on the beta profile of the initial lattice, reducing the max  $\beta$  value from 29.91 to 17.80, a 40% reduction. Furthermore, as shown in Figure 7.3, the  $\mathbf{x}_b$  lattice satisfies the constraints, whereas the initial lattice does not. Although several constraints are nearly active, this is in part due to the tight tolerances on the two-sided constraints, which imply that a small change in  $\mathbf{x}$  can easily change the constraint activities.

While the low beta functions and feasibility of our solution indicate that we have satisfied the design criteria, the true value of the candidate design is still unclear since the simulation used in this study captures only linear optics. A full validation of our model and candidate solution leading up to a



**Fig. 7.2** (top) Visualization of the placement and strengths of the elements of the design  $\mathbf{x}_b$ . Also shown are the corresponding beta profile (middle) and dispersion profile (bottom).




**Fig. 7.3** Value of  $c_i(\mathbf{x})$  for the initial lattice configuration  $\mathbf{x} = \mathbf{e}_0$  (blue) and the optima  $\mathbf{x} = \mathbf{x}_b$  (orange). Points left of the dotted black line indicate constraint satisfaction.

full project plan would require further investigation, particularly including a collective particle simulation.

To satisfy ourselves that the optimized lattice would be able to propagate particles properly, we performed particle-tracking simulations using the  $\mathbf{x}_b$  lattice. The beam parameters are shown in Table 7.1. Over the course of

Parameter	Value	Units
Particles	1,000,000	
Normalized emittance	24.0	mm-mr
Beta functions $x$ and $y$	2.15	m
Transverse RMS beam size	2.34	mm
Width $\Delta p/p$	2.5%	uniform

**Table 7.1** Beam parameters for optimized lattice ( $\mathbf{x}_b$ ) simulation.




**Fig. 7.4** The  $x$  and  $y$  emittances of a representative beam distribution simulated using  $\mathbf{x}_b$ . The beamline is very stable: even after 20,000 turns, the emittances have a relative change less than 0.1%.

20,000 simulated turns, there was no particle loss. As seen in [Figure 7.4](#), the emittance was stable within 0.1% with no oscillations or long-term growth.

In applications, the constraints on the lattice optics should also be robust to small perturbations in the quadrupole strength. Perturbing the quadrupole strength simulates uncontrolled variation in magnet construction and power supplies. Any variation in dipole strength can be corrected with greater precision and has only indirect effect on the accelerator optics.

[Figure 7.5](#) shows that the  $\mathbf{x}_b$  lattice is indeed robust to uniformly distributed perturbations  $U([-10^{-3}, 10^{-3}]^{n_q})$  to the  $n_q$  quadrupole strengths. These perturbations reflect nontrivial changes of roughly 0.08% relative to the original quadrupole strength bounds. The chromaticity matching criteria ( $c_{22}, c_{23}$ ) and the ratio of the beta functions in the center of the NL insert ( $c_{10}, c_{11}$ ) show the most significant variation in the constraint values and are easily restored under tuning. Those constraints, along with the dispersion-matching constraints ( $c_4, \dots, c_9$ ), exhibit perturbative behavior which naturally spans the constrained values. The betatron tune-matching ( $c_{18}, c_{19}$ ) and T-insert phase advances ( $c_{14}, \dots, c_{17}$ ) are Danilov–Nagaitsev integrability criteria that appeared to be violated under nearly any small random perturbation. Fortunately, those parameters can be measured with great precision in application and can be restored in dedicated tuning. The beam stability cri-



**Fig. 7.5** Empirical distribution of constraint values  $c_i(\mathbf{x}_b + \mathbf{I}_q \boldsymbol{\xi})$  around  $\mathbf{x}_b$  under  $\boldsymbol{\xi} \sim U([-10^{-3}, 10^{-3}]^{n_q})$  perturbations to the  $n_q = 9$  quadrupole strengths. When applicable, the maximum is taken over constraint pairs (upper and lower bounds) to show the worse of the two constraint violations. Orange squares indicate the constraint value (or maximum value for constraint pairs) at  $\mathbf{x}_b$ . The distribution of constraint values under the random perturbations is robust since they are within a tolerable range to be corrected in tuning.


terion ( $c_8, c_9$ ), the momentum-compaction factor ( $c_3$ ), the beta matching at the edge of the NL insert ( $c_{12}, c_{13}$ ), and the phase advance over the NL insert ( $c_{20}, c_{21}$ ) are all insensitive to the small perturbations.

The robustness of the lattice design  $\mathbf{x}_b$  may in part be due to the stringent simulation-based constraints required at the solution: we observe that points within the feasible region tend to be robust to large perturbations, while points violating simulation-based constraints tend to demonstrate an instability in their constraint values when perturbed.

### 7.3 Sensitivity and tradeoff analysis

The role of constraints in this problem cannot be overstated: all of the optima found are constrained optima in that the objective value would continue to decrease if the active constraints were removed. This begs the question of how constraint parameters, such as the right-hand side  $\boldsymbol{\varepsilon}$ , play a role in selection of solutions. We explore this question through the KKT conditions at our constrained optima. The KKT conditions give rise to Lagrange multipliers, which quantify the improvement in objective value when active constraints are relaxed. We numerically solve [eq. \(6.13\)](#) to compute the Lagrange multipliers.

In order to be physically meaningful, a design must only be resolved so that the positions, lengths, and strengths are prescribed to a precision of roughly  $10^{-3}$  or  $10^{-4}$ . Element attributes cannot practically be tuned any tighter than this in a particle accelerator, so optimizing further is not meaningful from this



**Fig. 7.6** (Top) Decrease of the optimal objective value as active constraints at  $\mathbf{x}_b$  are relaxed. (Bottom) The minuscule difference between  $h_b$  and its predicted value using the linearization eq. (7.1) while moving in the direction  $\nabla c_i(\mathbf{x})$  for the active constraints. The high order of accuracy implies that the Lagrange multipliers capture the objective-constraint tradeoff tightly.

perspective. However, in order to validate the KKT conditions and compute the Lagrange multipliers, solutions must be resolved to greater precision.


The Lagrange multipliers for active constraints at constrained optima give a prediction of how relaxation of these active constraints (i.e., increasing  $\varepsilon_i$  in eq. (6.6)) improves the objective value. For instance, at the optima  $\mathbf{x}_b$ , where only simulation-based constraints indexed by  $\mathcal{A}_{sc}$  are active, the change in objective value along a direction  $\mathbf{d}$  can be predicted locally with the linearization

$$h_b(\boldsymbol{\beta}(\mathbf{x})) \approx \hat{h}_b(\boldsymbol{\beta}(\mathbf{x})) := h_b(\boldsymbol{\beta}(\mathbf{x})) - \sum_{i \in \mathcal{A}_{sc}(\mathbf{x}_b)} \lambda_i^s \mathbf{d}^T \nabla c_i(\mathbf{x}_b). \quad (7.1)$$

While this shows that relaxing active constraints can improve the objective value, it also shows that relaxing inactive constraints will have no effect on the solution. In Figure 7.6 we see the actual and predicted decrease of the objective as we relax active constraints.

Figure 7.6 shows that a small relaxation of constraint  $c_9$  (the constraint on dispersion in the nonlinear section) yields a correspondingly small improvement in the objective. This suggests that a redesign of the achromatic bend may allow for smaller peak beta functions. Similarly a small relaxation of constraint  $c_{14}$  (phase advance over the T insert) or  $c_{18}$  (matching the betatron tunes) also has a small impact on the objective function. Because many of the Danilov–Nagaitsev integrability constraints are not active, the effect on the optimal solution of relaxing the complete set  $(c_6, \dots, c_{23})$  remains small.





**Fig. 7.7** Decrease in the objective  $h_b$  and the three essentially active beta functions  $\beta_{x,48}, \beta_{x,180}, \beta_{y,112}$  along the line segment  $\mathbf{x}_b + t\mathbf{d}^*$  until  $\beta_{y,159}$  becomes active. At  $\mathbf{x}_b$ , relaxing  $\beta_{x,48}, \beta_{x,180}, \beta_{y,112}$  by roughly  $\epsilon = 0.015$  would result in a change in activities, making  $\beta_{y,159}$  the sole active index.

Furthermore, none of the bound or linear constraints are active at  $\mathbf{x}_b$ , and so relaxing those will similarly not affect the solution.

In addition to interpreting the dependence of the optimal objective value on the active constraints, we can interpret the dependence of the optimal objective value on the active beta functions (i.e., the beta functions  $\beta_j$  for which  $h_b(\beta(\mathbf{x}_b)) = \beta_j(\mathbf{x}_b)$ ). At the optima  $\mathbf{x}_b$ , the three beta functions  $\beta_{x,48}, \beta_{x,180}, \beta_{y,112}$  are all essentially active, achieving values within  $10^{-4}$  of the objective value  $h_b(\beta(\mathbf{x}_b))$ . Mathematically, we can relax the dependence of the objective  $h_b$  on these three beta function values by considering the alternative objective  $h_b(\beta(\mathbf{x}) - \sum_{j \in J^*(\mathbf{x}_b)} \epsilon \mathbf{e}_j)$ , where  $\mathbf{e}_j$  is the  $j$ th column of the identity matrix and  $\epsilon > 0$  is a small parameter to make the active beta function values inactive. The goal of relaxing active beta function values is to allow the objective to have an increased value at the beta function's associated element in order to have an improved performance along the remainder of the cell.

Figure 7.7 shows the value of  $h_b$ , the three essentially active beta functions  $\beta_{x,48}, \beta_{x,180}, \beta_{y,112}$ , and  $\beta_{y,159}$  along the line segment  $\mathbf{x}_b + t\mathbf{d}^*$ , where  $\mathbf{d}^*$  is the normalized solution to  $\text{proj}\{0, \text{co}\{\nabla\beta_{x,48}(\mathbf{x}_b), \nabla\beta_{x,180}(\mathbf{x}_b), \nabla\beta_{y,112}(\mathbf{x}_b)\}\}$ . At  $\mathbf{x}_b$ , relaxing  $\beta_{x,48}, \beta_{x,180}, \beta_{y,112}$  by roughly  $\epsilon = 0.015$  would result in a change in activities, making  $\beta_{y,159}$  the sole active index. Furthermore under smaller relaxations the activities would not change, and so motion along the descent direction  $\mathbf{d}^*$  would yield a brief descent in  $h_b$  before  $\beta_{y,159}$  becomes active. This suggests a tradeoff between the peak value of the vertical beta function in the bending section ( $\beta_{y,112}$ ) and the inner quadrupole doublet ( $\beta_{y,159}$ ). As the peak vertical beta function in the bending section is reduced, the two peaks in the horizontal beta function within the bending section ( $\beta_{x,91}$  and  $\beta_{x,130}$ ) may be increased without becoming active. Similarly the peaks in the

horizontal beta function within the bending section may be increased while the peak horizontal beta functions at the outer matching section ( $\beta_{x,48}$ ) and inner quadrupole doublet ( $\beta_{x,180}$ ) are reduced. The peak beta function can be allowed to be higher for a contiguous section of an accelerator lattice (such as the inner doublets and RF section) as a design choice in which the aperture is increased (which raises the cost and reduces the field strength of magnets in that section).

## 8 Conclusion

Through the development and solution of an optimization model we explored the design of an integrable rapid cycling synchrotron. The model leverages linear optics to rapidly simulate the properties of lattice designs. Challenges inherent to the model include its sizable dimension (32 decision variables, 55 linear constraints, and 23 simulation-based constraints); lack of derivative availability of the simulation-based quantities; and a nonsmooth, simulation-based objective function. By judicious navigation of the feasible region and exploitation of known compositions of the simulation-based quantities, application of a manifold sampling algorithm yielded solutions that not only verify (up to numerical tolerances) the viability of the integrable lattice design but also perform well. This study aims to serve as a foundation for further accelerator optimization studies by methodically formulating a model, taking advantage of problem structure, and studying key sensitivities.

Additional improvements can be made toward a more comprehensive and general-purpose design of synchrotrons. For example, this optimization relied on specific optics-based lattice criteria that prior Synergia simulations of intense nonlinear beams have shown to be associated with achieving benchmarks in machine performance, such as the beam quality and the loss-limited beam intensity. An analysis complementary to our work could integrate nonlinear space-charge simulations into the optimization model directly to explore the relationship between machine parameters and those performance benchmarks. However, the outcome of such an optimization would necessarily be contingent on assumptions regarding the initial beam distribution, the allocation of machine errors, the availability of beam tuning, and the particle loss model.

The lattice optimization also takes place within the context of an overall lattice cell structure (the sequence of magnets as well as the number of periodic cells). In prior design work, this overall lattice cell structure was generated manually through a combination of domain knowledge and elementary operations (splitting, merging, transposing, and changing the number of periodic cells). If the lattice optimization presented here could be paired with an algorithm for selecting and rejecting cell structures, then a truly general accelerator optimization result could be generated. A key challenge in such a method is that it would require efficiently managing the vast majority of cell structures that are unstable, directly incompatible with constraints, or highly ineffective.

**Acknowledgements** We gratefully acknowledge the computing resources provided on Be-bop, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

## References

1. Abi, R., et al.: Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics. Fermilab, Batavia, FERMILAB-PUB-20-025-ND (2020). URL <https://arxiv.org/abs/2002.03005>
2. Ainsworth, R., Dey, J., Eldred, J., Harnik, R., Jarvis, J., Johnson, D.E., Kourbanis, I., Neuffer, D., Pozdeyev, E., Syphers, M.J., Valishev, A., Yakovlev, V.P., Zwaska, R.: An upgrade path for the Fermilab accelerator complex. Fermilab, Batavia, FERMILAB-TM-2754-AD-APC-PIP2-TD (2021). URL <https://arxiv.org/abs/2106.02133>
3. Amundson, J., Goldhaber, S., Lebrun, P., Lu, Q., Macridin, A., Michelotti, L., Park, C.S., Spentzouris, P., Stern, E.: Synergia. URL <https://synergia.fnal.gov/>
4. Amundson, J.F., Spentzouris, P., Qiang, J., Ryne, R.: Synergia: an accelerator modeling tool with 3-D space charge. *Journal of Computational Physics* **211**, 229–248 (2006). doi:[10.1016/j.jcp.2005.05.024](https://doi.org/10.1016/j.jcp.2005.05.024)
5. Antipov, S., Broemmelsiek, D., Bruhwiler, D., Edstrom, D., Harms, E., Lebedev, V., Leibfritz, J., Nagaitsev, S., Park, C., Piekarz, H., Piot, P., Prebys, E., Romanov, A., Ruan, J., Sen, T., Stancari, G., Thangaraj, C., Thurman-Keup, R., Valishev, A., Shiltsev, V.: IOTA (Integrable Optics Test Accelerator): Facility and experimental beam physics program. *Journal of Instrumentation* **12**(03), T03002–T03002 (2017). doi:[10.1088/1748-0221/12/03/t03002](https://doi.org/10.1088/1748-0221/12/03/t03002)
6. Audet, C., Hare, W.L.: *Derivative-Free and Blackbox Optimization*. Springer (2017). doi:[10.1007/978-3-319-68913-5](https://doi.org/10.1007/978-3-319-68913-5)
7. Bagirov, A., Karmitsa, N., Mäkelä, M.M.: *Introduction to Nonsmooth Optimization: Theory, Practice and Software*. Springer (2014). doi:[10.1007/978-3-319-08114-4](https://doi.org/10.1007/978-3-319-08114-4)
8. Bieler, B., Enge, A.: Vinci (2000). URL <https://www.math.u-bordeaux.fr/~aenge>
9. Chen, Y., Dwivedi, R., Wainwright, M.J., Yu, B.: Fast MCMC sampling algorithms on polytopes. *The Journal of Machine Learning Research* **19**(1), 2146–2231 (2018). URL <https://jmlr.org/papers/v19/18-158.html>
10. Danilov, V., Nagaitsev, S.: Nonlinear accelerator lattices with one and two analytic invariants. *Physical Review Accelerators and Beams* **13**(8), 084002 (2010). doi:[10.1103/PhysRevSTAB.13.084002](https://doi.org/10.1103/PhysRevSTAB.13.084002)
11. Duris, J., Kennedy, D., Hanuka, A., Shtalenkova, J., Edelen, A., Baxevanis, P., Egger, A., Cope, T., McIntire, M., Ermon, S., Ratner, D.: Bayesian optimization of a free-electron laser. *Physical Review Letters* **124**(12) (2020). doi:[10.1103/physrevlett.124.124801](https://doi.org/10.1103/physrevlett.124.124801)
12. Edelen, A., Neveu, N., Frey, M., Huber, Y., Mayes, C., Adelman, A.: Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems. *Physical Review Accelerators and Beams* **23**(4) (2020). doi:[10.1103/physrevaccelbeams.23.044601](https://doi.org/10.1103/physrevaccelbeams.23.044601)
13. Edwards, D.A., Syphers, M.J.: *An Introduction to the Physics of High Energy Accelerators*. Wiley-VCH (1993). doi:[10.1002/9783527617272](https://doi.org/10.1002/9783527617272)
14. Eldred, J.: Novel approaches to high-power proton beams. *PoS NuFact2019*, 055 (2020). doi:[10.22323/1.369.0055](https://doi.org/10.22323/1.369.0055)
15. Eldred, J., Lebedev, V., Valishev, A.: Rapid-cycling synchrotron for multi-megawatt proton facility at Fermilab. *Journal of Instrumentation* **14**(07), P07021 (2019). doi:[10.1088/1748-0221/14/07/P07021](https://doi.org/10.1088/1748-0221/14/07/P07021)
16. Eldred, J., Valishev, A.: Space-charge simulation of integrable rapid cycling synchrotron. *Proceedings of the 8th Int. Particle Accelerator Conf.* (2017). doi:[10.18429/JACOW-IPAC2017-THPVA032](https://doi.org/10.18429/JACOW-IPAC2017-THPVA032)
17. Eldred, J., Valishev, A.: Simulation of integrable synchrotron with space-charge and chromatic tune-shifts. *Proceedings of the 9th Int. Particle Accelerator Conf.* (2018). doi:[10.18429/JACOW-IPAC2018-TUPAF073](https://doi.org/10.18429/JACOW-IPAC2018-TUPAF073)
18. Fletcher, R.: *Practical Methods of Optimization*, second edn. John Wiley & Sons (1987). doi:[10.1002/9781118723203](https://doi.org/10.1002/9781118723203)

19. Huang, X.: Robust simplex algorithm for online optimization. *Physical Review Accelerators and Beams* **21**(10), 104601 (2018). doi:[10.1103/PhysRevAccelBeams.21.104601](https://doi.org/10.1103/PhysRevAccelBeams.21.104601)
20. Huang, X., Safranek, J.: Nonlinear dynamics optimization with particle swarm and genetic algorithms for SPEAR3 emittance upgrade. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **757**, 48–53 (2014). doi:[10.1016/j.nima.2014.04.078](https://doi.org/10.1016/j.nima.2014.04.078)
21. Khachiyan, L., Boros, E., Borys, K., Elbassioni, K., Gurvich, V.: Generating all vertices of a polyhedron is hard. *Discrete & Computational Geometry* **39**(1-3), 174–190 (2008). doi:[10.1007/s00454-008-9050-5](https://doi.org/10.1007/s00454-008-9050-5)
22. Khan, K.A.: Branch-locking AD techniques for nonsmooth composite functions and nonsmooth implicit functions. *Optimization Methods and Software* **33**(4-6), 1127–1155 (2017). doi:[10.1080/10556788.2017.1341506](https://doi.org/10.1080/10556788.2017.1341506)
23. Khan, K.A., Larson, J., Wild, S.M.: Manifold sampling for optimization of nonconvex functions that are piecewise linear compositions of smooth components. *SIAM Journal on Optimization* **28**(4), 3001–3024 (2018). doi:[10.1137/17m114741x](https://doi.org/10.1137/17m114741x)
24. Larson, J., Menickelly, M., Wild, S.M.: Manifold sampling for  $\ell_1$  nonconvex optimization. *SIAM Journal on Optimization* **26**(4), 2540–2563 (2016). doi:[10.1137/15M1042097](https://doi.org/10.1137/15M1042097)
25. Larson, J., Menickelly, M., Wild, S.M.: Derivative-free optimization methods. *Acta Numerica* **28**, 287–404 (2019). doi:[10.1017/s0962492919000060](https://doi.org/10.1017/s0962492919000060)
26. Larson, J., Menickelly, M., Zhou, B.: Manifold sampling for optimizing nonsmooth nonconvex compositions. *SIAM Journal on Optimization* (2021). URL <https://arxiv.org/abs/2011.01283>. To appear
27. Le Digabel, S., Wild, S.M.: A taxonomy of constraints in black-box simulation-based optimization. Preprint ANL/MCS-P5350-0515, Argonne National Laboratory, Mathematics and Computer Science Division (2015-01). URL <http://www.mcs.anl.gov/papers/P5350-0515.pdf>
28. Lebedev, V., et al.: The PIP-II conceptual design report. Fermilab, Batavia, FERMILAB-TM-2649-AD-APC (2017). URL [https://pxie.fnal.gov/PIP-II\\_CDR/PIP-II\\_CDR\\_v.0.3.pdf](https://pxie.fnal.gov/PIP-II_CDR/PIP-II_CDR_v.0.3.pdf)
29. Lee, S.Y.: *Accelerator Physics*, third edn. World Scientific (2011). doi:[10.1142/8335](https://doi.org/10.1142/8335)
30. Li, Y., Cheng, W., Yu, L.H., Rainer, R.: Genetic algorithm enhanced by machine learning in dynamic aperture optimization. *Physical Review Accelerators and Beams* **21**(5), 054601 (2018). doi:[10.1103/PhysRevAccelBeams.21.054601](https://doi.org/10.1103/PhysRevAccelBeams.21.054601)
31. Macridin, A., Burov, A., Stern, E., Amundson, J., Spentzouris, P.: Simulation of transverse modes with their intrinsic Landau damping for bunched beams in the presence of space charge. *Physical Review Accelerators and Beams* **18**(7), 074401 (2015). doi:[10.1103/PhysRevSTAB.18.074401](https://doi.org/10.1103/PhysRevSTAB.18.074401)
32. Martin, L.K., Kelliher, D.J., Sheehy, S.L.: Can a Paul ion trap be used to investigate nonlinear quasi-integrable optics? *Journal of Physics: Conference Series* **1350**, 012132 (2019). doi:[10.1088/1742-6596/1350/1/012132](https://doi.org/10.1088/1742-6596/1350/1/012132)
33. Mete, H.O., Zabinsky, Z.B.: Pattern hit-and-run for sampling efficiently on polytopes. *Operations Research Letters* **40**(1), 6–11 (2012). doi:[10.1016/j.orl.2011.11.002](https://doi.org/10.1016/j.orl.2011.11.002)
34. Michelotti, L., Ostiguy, J.F.: CHEF: A framework for accelerator optics and simulation. In: *Proceedings of 9th International Computational Accelerator Physics Conference*, pp. 2–6 (2006)
35. Minty, M.G., Zimmermann, F.: *Measurement and Control of Charged Particle Beams*. Springer (2003). doi:[10.1007/978-3-662-08581-3](https://doi.org/10.1007/978-3-662-08581-3)
36. Nagaitsev, S., Lebedev, V.: A cost-effective rapid-cycling synchrotron. *Reviews of Accelerator Science and Technology* **10**(01), 245–266 (2019). doi:[10.1142/s1793626819300135](https://doi.org/10.1142/s1793626819300135)
37. Nelder, J.A., Mead, R.: A simplex method for function minimization. *The Computer Journal* **7**(4), 308–313 (1965). doi:[10.1093/comjnl/7.4.308](https://doi.org/10.1093/comjnl/7.4.308)
38. Neveu, N., Hudson, S., Larson, J., Spentzouris, L.: Comparison of model-based and heuristic optimization algorithms applied to photoinjectors using libEnsemble. In: *Proceedings of the 13th International Computational Accelerator Physics Conference*, pp. 22–24 (2019). doi:[10.18429/JACoW-ICAP2018-SAPAF03](https://doi.org/10.18429/JACoW-ICAP2018-SAPAF03)
39. Neveu, N., Larson, J., Power, J.G., Spentzouris, L.: Photoinjector optimization using a derivative-free, model-based trust-region algorithm for the Argonne Wakefield Accelerator. *Journal of Physics: Conference Series* **874**, 012062 (2017). doi:[10.1088/1742-6596/874/1/012062](https://doi.org/10.1088/1742-6596/874/1/012062)

40. Nocedal, J., Wright, S.J.: Numerical Optimization, second edn. Springer (2006). doi:[10.1007/978-0-387-40065-5](https://doi.org/10.1007/978-0-387-40065-5)
41. Pang, X., Rybarczyk, L.: Multi-objective particle swarm and genetic algorithm for the optimization of the LANSCE linac operation. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **741**, 124–129 (2014). doi:[10.1016/j.nima.2013.12.042](https://doi.org/10.1016/j.nima.2013.12.042)
42. Powell, M.J.D.: A direct search optimization method that models the objective and constraint functions by linear interpolation. In: S. Gomez, J.P. Hennart (eds.) Advances in Optimization and Numerical Analysis, *Mathematics and its Applications*, vol. 275, pp. 51–67. Springer (1994). doi:[10.1007/978-94-015-8330-5\\_4](https://doi.org/10.1007/978-94-015-8330-5_4)
43. Ragonneau, T.M., Zhang, Z.: PDFO: Cross-platform interfaces for Powell’s derivative-free optimization solvers (version 1.0). doi:[10.5281/zenodo.3887569](https://doi.org/10.5281/zenodo.3887569)
44. Roussel, R., Hanuka, A., Edelen, A.: Multiobjective Bayesian optimization for online accelerator tuning. Phys. Rev. Accel. Beams **24**, 062801 (2021). doi:[10.1103/PhysRevAccelBeams.24.062801](https://doi.org/10.1103/PhysRevAccelBeams.24.062801)
45. Ruisard, K., Komkov, H.B., Beaudoin, B., Haber, I., Matthew, D., Koeth, T.: Single-invariant nonlinear optics for a small electron recirculator. Physical Review Accelerators and Beams **22**(4), 041601 (2019). doi:[10.1103/PhysRevAccelBeams.22.041601](https://doi.org/10.1103/PhysRevAccelBeams.22.041601)
46. Scheinker, A., Pang, X., Rybarczyk, L.: Model-independent particle accelerator tuning. Physical Review Accelerators and Beams **16**(10), 102803 (2014). doi:[10.1103/PhysRevSTAB.16.102803](https://doi.org/10.1103/PhysRevSTAB.16.102803)
47. Shang, H., Borland, M.: A parallel simplex optimizer and its application to high-brightness storage ring design. In: Proceedings of the 2005 Particle Accelerator Conference, pp. 4230–4232. IEEE (2005). doi:[10.1109/pac.2005.1591774](https://doi.org/10.1109/pac.2005.1591774)
48. Shiltsev, V.: Fermilab proton accelerator complex status and improvement plans. Modern Physics Letters A **32**(16) (2017). doi:[10.1142/S0217732317300129](https://doi.org/10.1142/S0217732317300129)
49. Sun, Y.: Multi-objective online optimization of beam lifetime at APS. In: Proceedings of North American Particle Accelerator Conference, 3, pp. 913–915 (2017). doi:<https://doi.org/10.18429/JACoW-NAPAC2016-WEPOB12>
50. Valishev, A.: Research at FAST/IOTA: Strategy and priorities. IOTA Collaboration Meeting (2020). URL [https://indico.fnal.gov/event/43231/contributions/187342/attachments/129553/157411/2020-06-15\\_Strategy\\_CollaborationMeeting.pdf](https://indico.fnal.gov/event/43231/contributions/187342/attachments/129553/157411/2020-06-15_Strategy_CollaborationMeeting.pdf)
51. Webb, S., Cook, N., Eldred, J.: Averaged invariants in storage rings with synchrotron motion. Journal of Instrumentation **15**(12), 12032 (2020). doi:[10.1088/1748-0221/15/12/p12032](https://doi.org/10.1088/1748-0221/15/12/p12032)
52. Webb, S.D., Bruhwiler, D.L., Abell, D.T., Sishlo, A., Danilov, V., Nagaitsev, S., Valishev, A., Danilov, K., Cary, J.R.: Effects of nonlinear decoherence on halo formation. Tech. Rep. 1205.7083, ArXiv (2012). URL <https://arxiv.org/abs/1205.7083>
53. Webb, S.D., Bruhwiler, D.L., Valishev, A., Nagaitsev, S.N., Danilov, V.V.: Chromatic and dispersive effects in nonlinear integrable optics. Tech. Rep. 1504.05981, ArXiv (2015). URL <https://arxiv.org/abs/1504.05981>
54. Wei, J.: Synchrotrons and accumulators for high-intensity proton beams. Reviews of Modern Physics **75**(1383), 1383–1432 (2003). doi:[10.1103/RevModPhys.75.1383](https://doi.org/10.1103/RevModPhys.75.1383)
55. Yang, L., Li, Y., Guo, W., Krinsky, S.: Multiobjective optimization of dynamic aperture. Physical Review Accelerators and Beams **14**(5), 054001 (2011). doi:[10.1103/PhysRevSTAB.14.054001](https://doi.org/10.1103/PhysRevSTAB.14.054001)

The U.S. Government retains for itself, and others acting on its behalf, a paid-up non-exclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. <http://energy.gov/downloads/doe-public-access-plan>.