

GLOBAL OPTIMIZATION OF DAMPING RING DESIGNS USING A MULTI-OBJECTIVE EVOLUTIONARY ALGORITHM*

L. Emery[†], ANL, Argonne, IL 60439, USA

Abstract

Several damping ring designs for the International Linear Collider have been proposed recently. Some of the specifications, such as circumference and bunch train, are not fixed yet. Designers must make a choice anyway, select a geometry type (dog-bone or circular), an arc cell type (TME or FODO), and optimize linear and nonlinear parts of the optics. The design process includes straightforward steps (usually the linear optics) and some steps not so straightforward (when nonlinear optics optimization is affected by the linear optics). A first attempt at automating this process for the linear optics is reported. We first recognize that the optics is defined by just a few primary parameters (e.g., phase advance per cell) that determine the rest (e.g., quadrupole strength). In addition to the exact specification of circumference, equilibrium emittance, and damping time, there are some other quantities, which could be optimized, that may conflict with each other. A multi-objective genetic optimizer solves this problem by producing a population of best-ranked solutions on a multi-dimensional surface from which one solution can be chosen by the designer. The application of the NSGA-II optimizer to a damping ring of FODO cells is presented.

INTRODUCTION

Several lattice designs for the damping rings of the International Linear Collider have been proposed [1]. They represent a sampling of possible circumference, layout (circular or dog-bone), and arc cell types (FODO or TME). Though the the above parameters have not been fixed at this time, beam specifications such as the normalized equilibrium horizontal and vertical emittances, damping time, and bunch length are fixed and not likely to change. Typically, a lattice results from a mixture of choices a designer makes (i.e., number of arc cells) and fitting routines of a lattice-matching program (strength of quadrupoles for a matching section). Then properties of the lattice are assessed, and probably more design iterations follow. Such a procedure was spelled out in [2]. For the present damping ring types and a continuum of possible circumference values, one might benefit from an automated procedure that could generate these lattices for any circumference.

We start by recognizing that a ring definition, consisting of a string of dipoles, quadrupoles, sextupoles, wigglers, and drift spaces, is very detailed but can be characterized by only a few independent primary quantities. For example, the phase advance and length for an arc cell are pri-

mary quantities while the strengths of the quadrupoles are derived from those quantities. One can define a full set of primary parameters that can then produce a complete ring with proper magnet definitions, whether or not the ring fulfills any desired criterion. Such complete rings can be created by scripts running lattice-matching programs.

The assessment of the above lattice definitions can be done automatically as well. It may consist of simple calculations of equilibrium beam properties, and could include more complex calculations such as dynamic aperture or collective effects. Each definition is characterized as either an exact specification or a desirable property that has to be maximized or minimized. The assessment is none other than a list of numbers that are functions of the original primary parameters.

Our goal is to generate a set of one or more lattices of a particular style and layout that fulfills the exact specifications and maximizes the values of the desirable properties, some of which may unfortunately oppose each other. We therefore cast the design process as an optimization of multiple objectives, for which software has been recently developed and is easily available, e.g., the Non-Dominated Sorted Genetic Algorithm, version 2 [3]. Adopting terminology from the field of optimization, the primary parameters are variables, the exact specifications are constraints, and the desirable properties are objectives.

The original approach to this work on optimizing damping rings was the summing of the objectives with individual weighting factors. This was clumsy when more than two objectives were used, and we could not get a feel for the trade-off between objectives. The inclusion of constraints added too much complexity. At that point, the author became aware of recent papers by Bazarov [4, 5] that described how the multi-objective genetic algorithm worked and how ILC main parameters and a photoinjector source were optimized.

OPTIMIZATION PROCEDURE

The multi-objective optimizer basically proceeds as follows. Some details are left out which can be found in [3]. The user provides a set of variables with their lower and upper limits. The first step is to generate a random sampling of, say, a hundred sets of variables (population of a hundred individuals). This initial random sampling is bound to produce nonsense results, but some individuals will be reasonable. The user-supplied script of functions calculates the values for the constraint violations and objectives. The algorithm reads in these values and ranks the individual based on the results and selects the best 50% to keep for mutation and crossing of the variables. The worse 50% are

* Work supported by U.S. Department of Energy, Office of Basic Energy Sciences under Contract No. W-31-109-ENG-38.

[†] emery@aps.anl.gov

thrown away. Mutation refers to slight modification of the variable values and crossing refers to exchanging the variable values between individuals, all in the hope of finding a better result.

A new population of a hundred is formed by the previous best and the new “children” individuals, and the above step is repeated. The population spread in objectives, constraints, and variables will converge to a small volume in their own multidimensional spaces. Though the convergence cannot be proved mathematically, for physical systems with no pathologies, the convergence always seems to occur.

The useful outcome of the algorithm is a distribution of best objective values lying on a multidimensional surface, which we’ll call a trade-off surface, though in the optimization field it is called a Pareto-optimum surface. The user is free to select any individual (i.e., a set of variables) on that surface and be assured that no particular objective is unnecessarily inferior.

Another aspect of genetic algorithms in general is that they are suited to cluster computing where the evaluations of individuals can take place independent of each other.

APPLICATION TO DAMPING RING

A Tc1 script was written to create a circular damping ring with FODO arc cells and FODO wiggler cells in the straight sections. Scripts for other types will be written later. For the moment, injection straight sections were omitted. The variables of the problem are: number of arc cells and their length, number of wiggler cells and their length, filling factor of dipole in FODO cell, wiggle period and field, wiggler length in each wiggler cell, rf gap voltage. Note that we must somehow include a wiggler model that relates the period, field, and aperture.

The constraints are: normalized emittance, damping time, circumference, bunch length, and wiggler aperture. Some of these constraints have a close relation to variables, such as the damping time and the wiggler length, typically. External models, such as the wiggler magnetic field model, can be inserted into the optimization as a constraint.

Possible objectives are: nonlinear kick from strongest sextupole (to maximize dynamic aperture), total wiggler length (a cost), magnet count (a cost), wiggler aperture (maximize physical aperture), wiggler nonlinearity (strength times period), and rf gap voltage. Note that one can have as an objective any variables. This is not an exhaustive list. One can add any characteristic of a damping ring that may be relevant. In this implementation all objectives are to be minimized. An objective that really needs to be maximized (wiggler aperture) is multiplied by -1 .

It takes about three minutes for a 2.4-GHz PC in our Linux cluster to develop a full lattice and to calculate the objective and constraint functions. Most of the time is actually used to fit the matching sections of the ring. Given that the optimization problem is of very high dimension and requires about a thousand iterations, we adopted a formula-

based model for calculating the objective and constraint functions, which takes a few milliseconds. A population of a thousand and a thousand generations takes about three minutes to complete.

In order to better visualize the optimization, a simplified version of the problem was formulated. For a 6-km FODO cell ring, we reduced the set of variables, constraints, and objectives in order to produce a trade-off surface that becomes a line in two-dimensional space. We decided on four variables: the number of arc cells, their length, the wiggler period, and the number of wiggler cells. The other variables were fixed to reasonable values. The three constraints are the normalized emittance (8×10^{-6} m-rad), damping time (25 ms), and circumference (6 km). The two objectives are the nonlinear kick, which pushes for fewer arc cells, and wiggler nonlinearity, which pushes for longer periods. Both of these objectives tend to increase the quantum excitation, which is, however, constrained by the emittance value. So we expect a trade-off between the objectives. The sum (5) of number of constraints and objectives minus the variables (4) gives the dimensionality (1) of the space spanned by the trade-off. After 400 generations of 400 individuals, we get the trade-off shown in Figure 1. Each point on the graph corresponds to a different set of variables.

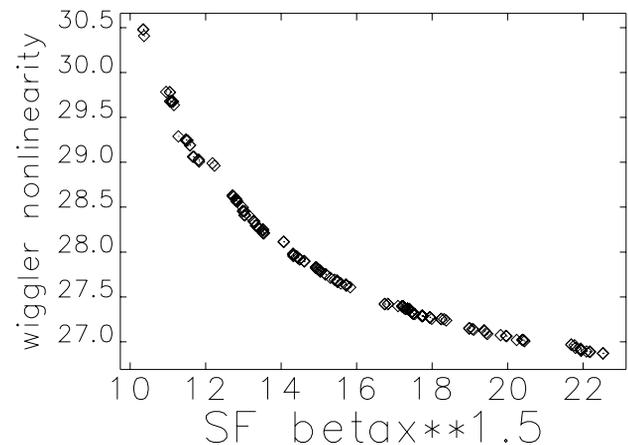


Figure 1: Trade-off between the two objectives measuring nonlinearities. The x-axis is a measure of sextupole nonlinearity.

For example, the solution corresponding to the lowest sextupole nonlinearity is at a value of 10, which corresponds to 92 arc cells, 58-m arc cell length, 8 wiggler cells, 10-m wiggler cell length (making the total length about 80 m), and a 0.36-m wiggler period. The algorithm doesn’t “know” which calculated nonlinearity is more important; it is up to the user to decide. In this particular case, it is safe to assume that the nonlinearity from the sextupole is more important. The range of solutions for the sextupole nonlinear kick is a factor of two, and for the wiggler it is a factor 1.1, which is small. Therefore one would select the solution at the bottom of the sextupole nonlinear kick range.

We present now the full problem with nine variables, five

constraints and seven objectives, while noting that some of the objectives are not really independent, nor do they have to be. Not surprisingly, this high-dimension problem requires a larger population and more iterations for convergence. We expect additional trade-offs among the objectives, though they may be complex. The dimensionality of the trade-off surface is expected to be 3, so pair-wise plots of objectives may only reveal a smeared projection of the surface. See Figures 2 through 4 for some plots.

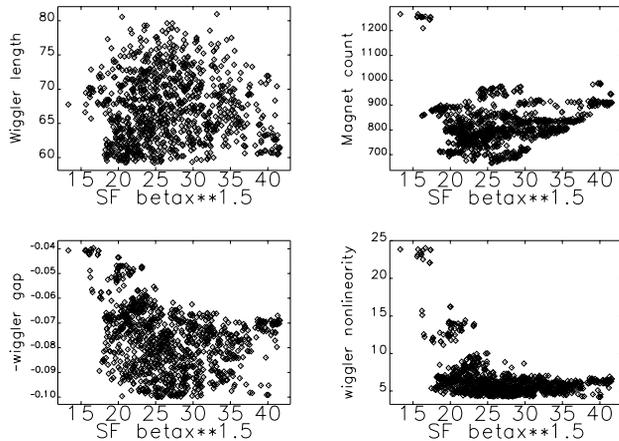


Figure 2: Pair-wise objective plots for best solutions. Many objectives appear uncorrelated.

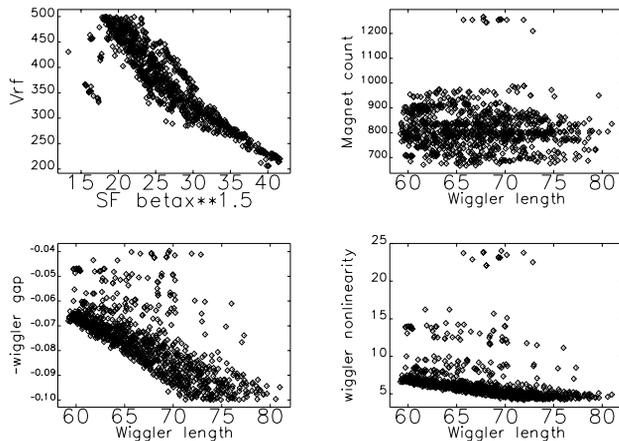


Figure 3: Pair-wise objective plots for best solutions. See text for comment on rf voltage.

In these figures, the plots that are discernible as lines may be interpreted as a trade-off between the two objectives, some are directly due to constraints. A change in constraint (e.g., bunch length) can completely change the appearance of the objective plots (not shown.) The objective can also be plotted against the variables for further insight.

Finally, when a solution is selected from the formula model, the real lattice can be generated and characterized.

The rf voltage objective (range of 200-500 MV) turned out to be surprisingly large compared to the energy loss per turn (about 8 MeV) because of the relatively high momentum compaction of a FODO cell. In any case, we can see

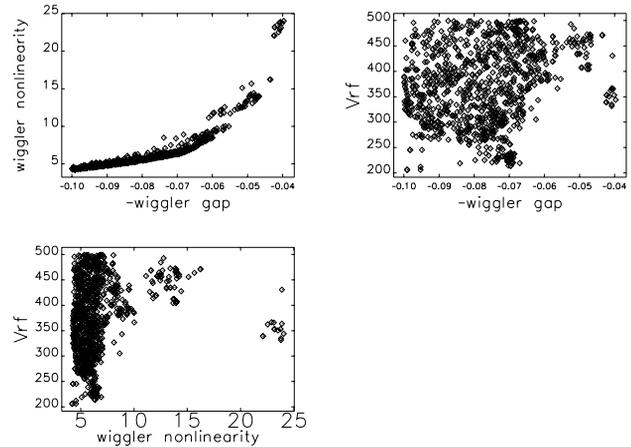


Figure 4: Pair-wise objective plots for best solutions. The correlation between wiggler gap and wiggler nonlinearity is due to the constraint from the wiggler model and the wiggler field solution near the upper range limit of 2 T.

the trade-off in Figure 3 between rf voltage and sextupole nonlinearity, which are connected through the number-of-FODO-cells variable. For a reasonable value of rf voltage, say, 50 MV, a short bunch length can be achieved only with a much shorter FODO cell and at the cost of higher nonlinearities. The seriousness of the problem may be reduced in other types of damping rings.

CONCLUSION

Software tools for the algorithm and full-lattice generation have been implemented. We intend to apply this analysis to other damping ring types. The computation time is an issue. The full lattice model is only feasible in runs that require 100-200 iterations and 100 in population. For large problem dimensions, only a formula model seems practical. Also for large dimensions, some tweaking of the optimizer setting (iteration number versus population number) may be required. We speculate that a reduced set of variables and objectives may be useful to analyze actual design decisions.

REFERENCES

- [1] ILC Damping Ring Collaboration, <http://www.desy.de/awolski/ILCDR/Lattices.htm>.
- [2] P. Emma and T. Raubenheimer, "Systematic Approach to Damping Ring Design," *Phys. Rev. ST Accel. Beams* 4 (2001) 021001.
- [3] K. Deb et al., "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Trans. Evolut. Comput.* 6 (2002) 182.
- [4] I. V. Bazarov and H. S. Padamsee, "Multivariate Optimization of ILC Parameters," these proceedings.
- [5] I. V. Bazarov and C. K. Sinclair, "Multivariate Optimization of a High Brightness DC Gun Photoinjector," *Phys. Rev. ST Accel. Beams* 8 (2005) 034202, <http://link.aps.org/abstract/PRSTAB/v8/e034202>.