

FAST OPTIMIZATION TECHNIQUES APPLIED TO LINEAR AND NON-LINEAR
LATTICES OF ELECTRON STORAGE RINGS AND ACCELERATORS†

S. Peck
Cornell University
Ithaca, N.Y. 14853

Summary

This paper describes a general interactive computer program, ZEN, developed to determine linear focussing parameters for CESR, the Cornell electron storage ring, and application of similar methods to non-linear elements. It is particularly suitable for lattices of electron accelerators, especially those with large numbers of independently variable focussing parameters. ZEN is a descendent of the linear lattice calculating program, MAGIC, Ref. 1, having improved interactive capability and new minimization routine, MINOP Ref. 2.

Introduction

The CESR magnet ring is to be roughly concentric with the Cornell synchrotron and larger in radius by about 1.6 meters. The geometry is fixed by the existing tunnel and by the requirement that the CESR orbit must bulge out away from the synchrotron in one interaction region area to provide room for large experimental apparatus. The basic design philosophy, then, has been to work with fixed component disposition and to optimize machine performance by variation of focussing parameters. For this reason, ZEN has no convenient means for adjusting component disposition. In the beginning a few interactions of basic cell lengths and low beta insertions were required to find a workable lattice over the required operating energy range (1.5 to 10 GeV), and luminosity range ($10^{32}(E/8 \text{ GeV})^2 \text{ cm}^{-2} \text{ sec}^{-1}$). Because of the low superperiodicity of the ring, (2), no attempt was made to design normal cells and match them with transition sections, but rather an entire half ring was treated as a unit which has to meet the optical requirements. Looked at another way, the whole lattice is merely a 49 quadrupole matching section between the two interaction regions. The components in the ring are of course laid out in as regular a fashion as possible.

The Program

Central to the optimization process is F, the 'Figure of Merit', a function of the quadrupole strengths, which have been selected by MINOP on the basis of preceeding function evaluations. The function quantifies the important lattice properties by the weighted sum of squared differences between the momentary and desired value of each term. Selection of the weights is an important component of user-program interaction.

Three preliminary calculations are necessary before finding the sum F itself. The first is performed only once at the start of an entire set of optimizations; in it, the thick lens linear matrices for the sequences of bend and drift components between each pair of quadrupoles are multiplied. Second, for each evaluation of F, the transfer matrix for 1/2 the ring (one superperiod) is found by repeated multiplication of thick-lens quadrupole matrices with the 'bend-drift' matrices found previously. This complete matrix then provides starting values for beta and eta at one interaction region. These are propagated through the same sets of matrices, and a list of beta and eta at each

quadrupole location is created.

Simultaneously, horizontal beta, eta, and their derivatives are traced through individual bend elements in a streamlined fashion to compute the synchrotron integrals. This phase of computation, which uses, wherever possible, invariant terms computed only once per optimization cycle, still consumes about 1/2 of the total CPU time.

Having completed these preliminaries, calculation of F is straightforward. As in MAGIC, desired values for beta and eta at the interaction points may be specified, but initial experience indicated a decreased weight for small errors was desirable. Otherwise, with a large weight on (for example) eta at the interaction region, MINOP would find a minimum with eta = 1.00000 M, and neglect altogether the beam-stay-clear, clearly not an optimum, though a local minimum. However, with less weight for small errors a solution with all desired conditions met and perhaps eta = 1.01, could be found, more in keeping with the physical intents of the optimization. The actual form used to enter the error in vertical interaction region beta is as follows (the coefficients are non-critical) $\text{weight} * (.0001 * \text{error}^2 + .1 * (\text{error over } 1\%)^2 + (\text{error over } 6\%)^2)$ similar forms with .02 as second coefficient are used for horizontal beta and eta, since their exact values are considered somewhat less critical. Attempts to use exponential error terms were unreliable, leading to unstable lattices in some cases.

Using the values of beta and eta already found at each quadrupole, their maximum values in the 'normal' and 'interaction' regions are constrained by adding to the figure of merit the weighted square of the excess (above a target value) of the greatest value found for each. Using the beam size found via the synchrotron integrals, the desired beam-stay-clear tolerances are enforced in the same way. ZEN also allows specification of beta, and/or eta, at individual quadrupole locations, forcing exception to global beta and eta limits as appropriate. These may be specified as upper or lower bounds, or exact values, as desired.

A capability in ZEN peculiar to CESR is making nearly equal the twiss parameters at two interaction regions, since they are not made equal by symmetry. Only operating experience will show if this capability is essential.

The beam size due to synchrotron radiation can be included in the figure of merit, effectively causing MINOP to adjust the relative eta-mismatch and therefore H, the scale of quantum-induced betatron oscillations. In the case where wigglers are used to control beam size ZEN might adjust them in two ways. Provision was originally made for the wiggler excitation to be included with the quadrupoles as a MINOP-controlled variable. When this proved unsatisfactory, ZEN itself was allowed to set the wigglers before each optimization so as to produce either the desired betatron beam size or energy spread. It is usually prudent to include a modest limit on maximum eta in the lattice to avoid a mix of wiggler and eta-mismatch when this is not desired.

Finally, given information about the beam-stay-

clear, interaction region beam size and aspect ratio, RF power, and operating energy, ZEN has provision for the luminosity to be included in the figure of merit. However, initial results using the weighted squared difference between the computed and desired luminosity were disappointing. MINOP generally would stop at a local minimum with an unsatisfactory compromise between luminosity and aperture requirements.

Thus, for example, at RF-limited energies, MINOP might decrease the vertical beta in the interaction region increasing luminosity at the cost of larger aperture in the adjacent quadrupoles, and too-rapid variation of crossing envelope. In the low energy, non RF-limited domain, it might increase the beam current, (and violate stay-clear requirements in the normal lattice).

This impasse was resolved by rendering MINOP blind to 'improvements' in luminosity so gained. Replacing the nominal luminosity with one decreased by a factor equal to $((\beta_V)/(\beta_V \text{ desired}))^{**2}$ when $\beta_V < \beta_V \text{ desired}$, and another equal to $((\text{desired aperture})/(\text{aperture}))^{**2}$ when $\text{aperture} > \text{desired aperture}$, accomplishes this blinding function satisfactorily. MINOP can then improve luminosity in a more useful manner.

ZEN has been successful in finding a wide range of optics for CESR, spanning the range of 1.5 to 10 GeV with luminosity equal to or greater than $10^{32} \cdot (E/8 \text{ GeV})^2$ for $E \leq 8 \text{ GeV}$.

Execution time for ZEN on a PDP-10 (K1) processor may be summarized for a typical run: @ 180 MS per figure of merit evaluation x 50 evaluations per MINOP iteration (for 49 numerically determined derivatives) x about 30 MINOP iterations per subsolution = 370 seconds, or about 45 minutes for a complete set of 10 such solutions. ZEN requires no more than 27,000 words of memory.

Non-Linear Application

Faced with the problems of adjusting the sextupoles of the CESR lattice (some 40 per 1/2 ring) it seemed reasonable to apply the techniques used by MAGIC and ZEN. While only two independently adjustable sets of sextupoles suffice to correct the natural chromaticity of the lattice, and provide a small positive value, making both vertical and horizontal NU values increase somewhat for higher momentum particles, several problems remain, or are created. Three of these problems may be controlled by a figure of merit function, similar to, but simpler than, the F used by ZEN. They are:

1. Variation of beta with momentum, both in the normal bulk of the lattice, and the interaction region.
2. Variation of chromaticity with momentum.
3. Variation of NU with betatron oscillation amplitude.

Of these, (1.) is by far the hardest to control, and is not without dangers of its own. Allowing minop to set very strong sextupole strengths in order to minimize beta variations can cause destructive non-linear resonances to be aggravated. Even when (3.) is controlled. These resonances may still limit maximum stable oscillation amplitudes. This is not to say that improvements in overall stability have not been seen; simply that more factors need be entered into the figure of merit, and equally important, more operating experience is needed.

The non-linear optimization program 'P' uses a thin-

lens sextupole representation, and precomputes all linear transfer matrices (bends, quads, and drifts) into composite linear matrices between sextupoles, analogous to the method used in ZEN between quadrupoles. Then, for two (typically +/- .4%) off-momenta, equivalent matrices (effectively two new linear lattices) are found using linearized sextupoles and tracking sinelike and cosinelike trajectories through them. Additionally, a test trajectory with horizontal betatron amplitude = 10 SIGMA of the nominal beam size is tracked to find its NU-value. Given the NU- and beta-values for these two off-momentum systems a figure of merit is constructed containing the summed squares of the excesses in beta, and errors in NU and interaction beta from the on-momentum values, or other assigned target values.

Limited experience using this program indicates that it does smooth the variations of beta with momentum, while controlling the variation of NU with both momentum and oscillation amplitude. In some cases this has increased the effective usable aperture for the beam; however there have also been cases where overall stability has not improved, or grown worse. It is hoped that more global criteria for beam stability can be included in this optimization scheme in the future.

Observations and Strategies

As in MAGIC the overall optimization proceeds in a (user-specified) number of steps from an initial set of requirements, and an initial lattice, which need not meet them, to a final set of requirements, producing a sequence of lattices, hopefully meeting the successive demands. Both from conversation with one of the MINOP authors and from operating experience, it has been found useful to preserve the Hessian matrix (containing approximations to the partial derivatives of the function F, with respect to all the variables) between these sub-optimizations. This allows successive optimization steps to utilize information about the figure of merit surface found in prior stages.

Two other strategies have proven useful in avoiding stalling of the optimization near a subsidiary minimum or near minimum in F. Before each sub-optimization step, all (except the interaction region quads) quadrupoles or sextupoles are randomly perturbed by a small amount, e.g. .0004 M⁻¹ of a typical .25 quad strength (K). Also the user may specify that starting values for each step be projected part way forward along the path of previous sub-solutions by linear extrapolation. Both methods are routinely used.

In general it has been found possible, by means of variable weights, to incorporate into F all the important design parameters among which engineering compromises must be made. In this way, more of the power of the optimization routine has been brought to bear on the design of accelerator lattices.

I wish to thank R. Littauer, D. Edwards, and R. Talman for significant contributions during the initial phases of this work, and H. Mei, R. Helmke, and S. Ball for their assistance in improving the optimization programming.

Work supported by the National Science Foundation.

1. King, Lee and Lee, SLAC -183 (1975)
2. J. Dennis and H. Mei, An Unconstrained Optimization Algorithm TR75-246 Dept. of Computer Science, Cornell University.