Sorting the LHC Dipoles using Simulated Annealing

V. Ziemann
CERN-SL, CH-1211 Geneva 23
Switzerland

Abstract

The superconducting dipole magnets in LHC contain many higher multipole components which will vary from magnet to magnet. Ordering the magnets in such a way as to minimize particular beam dynamic quantities constitutes a combinatorial optimization problem that can efficiently be solved using simulated annealing. We investigate sorting the LHC dipoles according to their skew quadrupole and their sextupole contents.

1 INTRODUCTION

The superconducting dipole magnets in LHC [1] will contain many unwanted multipole components due to their design with two magnet channels in a common cryostat and other factors. These multipoles will have systematic and random components of which the latter will vary from magnet to magnet. Before installation in the tunnel the magnets will be thoroughly measured and thus the values of the multipoles will be known for each magnet. This allows sorting the magnets in such a way as to minimize their adverse effect on beam quality.

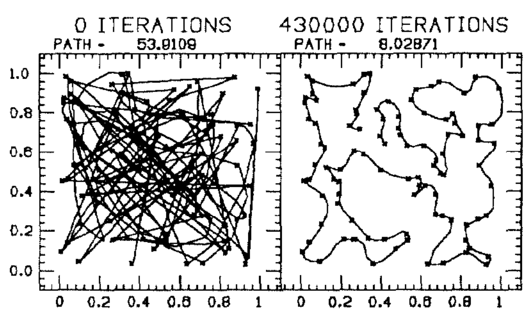

Figure 1: *The Travelling Salesman Problem on the unit square for 100 cities solved by simulated annealing.*

In order to sort a batch of magnets they first have to be delivered and measured. All magnets in a batch must then be stored and finally the magnets will be assembled in the tunnel. Clearly, storing a large batch of magnets will take a lot of space. Therefore, the size of the sorted batch should be as small as possible.

In this paper we consider a single octant of LHC, Version 2 which contains 24 90-degree FODO cells with six 13 m dipoles each. We consider the skew quadrupole component and the sextupole component and will also test whether sorting smaller batches than a full arc yields acceptable results.

2 SORTING BY SIMULATED ANNEALING

The sorting procedure will consist of minimizing an appropriate cost function (discussed below) by reordering the sequence of the magnets. It is an example of combinatorial optimization which can be handled by simulated annealing [2]. The typical example of a problem that can be handled by this algorithm is the travelling salesman problem in which the path length between a large number of points is to be minimized as a function of the sequence in which the points are visited. Figure 1 shows an example of 100 randomly chosen points in the unit square before and after sorting. We see that after about 430 000 iterations the path length is reduced by almost a factor 7 from about 54 to about 8. Clearly the number of iterations and consequently, of the cost function evaluations is very large, necessitating to make them as fast as possible. This excludes the direct use of the dynamic aperture of the entire machine as a cost function, which can only be calculated in a tracking code and is too slow.

Some theory about simulated annealing algorithms is described in ref. 2. We just repeat the gist of the idea here. Starting from an initial arbitrary configuration (sequence of magnets as delivered) we calculate the cost function, then we choose a reordering and recalculate the cost function. If the new cost is smaller, we accept it. If it is bigger, we only accept it with a given probability that is controlled by a temperature-like parameter. This effective temperature is reduced as the optimization goes along, thus reducing the probability to accept configurations that increase the cost function. Reordering the sequence of magnets is done according to two different schemes: (1) choose a section of magnets and translate it to another place, (2) choose a section and reverse it.

3 SORTING THE SKEW QUADRUPOLE COMPONENT

The smallest major multipole component in the dipoles is the skew quadrupolar, which has a sizeable systematic skew component of $K1L = 3.75\,10^{-5}$/m and a fluctuating component which we assume to have a standard deviation of twice that value and is truncated at two standard deviations.

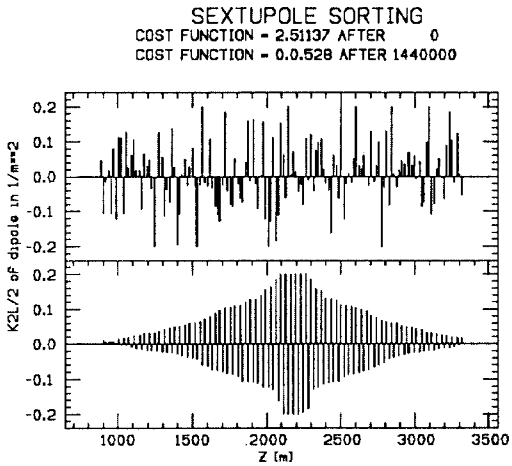

Figure 2: *The longitudinal distribution of the sextupole component of the dipoles before and after sorting.*

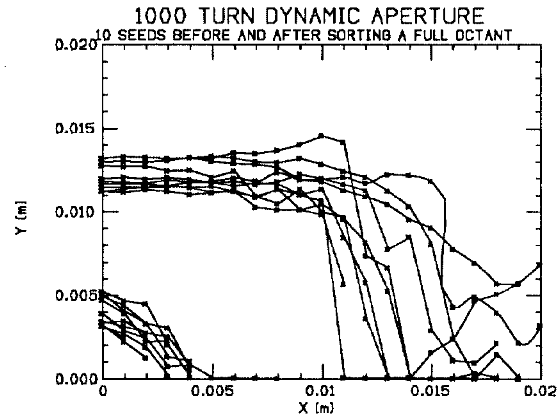

Figure 3: *The Dynamic Aperture before and after sorting for 10 seeds. Here all* 144 *dipoles of one octant are sorted at once.*

In a first attempt we use the sum of squares of the off-diagonal transfer matrix elements through the octant as cost function $C_1 = R_{13}^2 + R_{14}^2 + R_{23}^2 + R_{24}^2$. Since one of the transfer matrix elements (R_{14}) is by far the biggest, the optimization procedure mainly worked on this degree of freedom and ordered the magnets accordingly. Since only one degree of freedom is affected, which depends on the position in the ring where it is calculated, this method is not pursued further.

In a second attempt we use the resonance width, which is independent of position where it is calculated as cost function C_2. Note that only one octant is sorted, and the rest of LHC is replaced by a single uncoupled design transfer matrix. We find that *the coupling width is dominated by the systematic skew quadrupole component* and that sorting almost does not affect the coupling width.

The last observation is confirmed by sorting according to a cost function C_3 that minimizes the local averages defined by

$$C_3 = \sum_{i=1}^{N} (s_i + s_{i+1})^2 + 10 s_1^2 \qquad (1)$$

where s is the sorted quantity (here the skew quadrupole component of magnet i modulo N). The last term causes the dipoles to be sorted such that the magnets with the smallest magnitudes are placed near the interaction regions. Using this cost function we confirm the observation of the previous paragraph that the coupling width is dominated by the systematic skew quadrupole component.

The dominance of the systematic component can be roughly explained by noting that random effects are suppressed with respect to their systematic counterparts by $1/\sqrt{N}$ which for $N = 144$ amounts to about an order of magnitude to a small value of about $5\,10^{-3}$ in the used model. Applying the simulated annealing sorting algorithm to a model with random skew quadrupole components only we find an improvement of the coupling width by two orders of magnitude to about 10^{-5}. The resulting sorted dipole distribution, however, shows no regularity. Despite the improvements achieved by sorting, we feel that *the weak effect of the random component on the coupling width should rather be compensated with skew quadrupoles than by sorting.*

4 SORTING THE SEXTUPOLE COMPONENT

The next major multipole component is the sextupolar component of the dipoles, to which we assign a rms strength of $K2L = 0.1/\text{m}^2$, truncated at two standard deviations.[1] We use a cost function of the type C_3 and check the result in a simple tracking routine that consists of a collection of 145 linear lumped transfer matrices for the sections between the 144 dipoles containing the sextupole kicks. This makes tracking very fast.

Figure 2 shows the sextupole component of the dipoles as a function of longitudinal position through one arc before and after sorting which in this case took 1 440 000 cost function evaluations. After sorting, the sextupole components are matched in magnitude, but opposing signs, as to minimize the sum in eq. 1.

Figure 3 shows the stable region in $x - y$ space in arbitrary units for ten different seeds before and after sorting. The dynamic aperture is determined from short term tracking over 1000 turns. Clearly, *sorting increases the dynamic aperture about three-fold.* Moreover we note that the spread between the seeds before and after sorting is about the same. Note that in this model we consider one arc only and the chromaticity sextupoles are not taken into account.

[1] Originally this strength was accidentically chosen to be 20 times stronger than the expected magnet-to-magnet variation resulting from a rms b_3 of 10^{-4}. Since we are mainly interested in the relative improvement due to sorting this does not affect the conclusions of this paper.

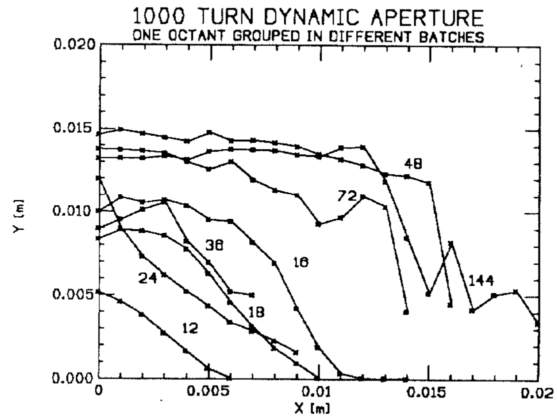

Figure 4: *The Dynamic Aperture with one octant sorted in batches of numbers indicated on the graph.*

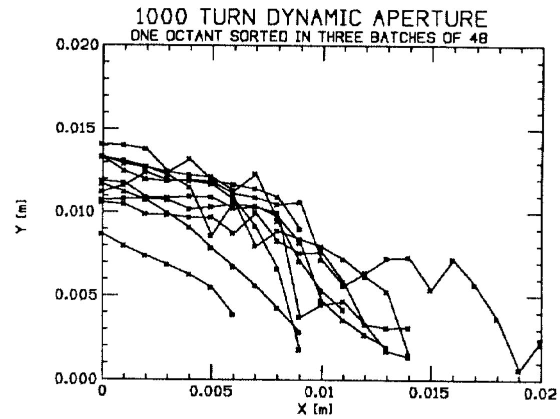

Figure 5: *The Dynamic Aperture for the same seeds used in Fig. 3, but the dipoles are sorted in three batches of* 48.

In a subsequent study we investigate the effect of sorting smaller batches of dipole magnets by splitting the 144 dipoles according to

$$\begin{aligned} 144 &= 2\times 72 = 3\times 48 = 4\times 36 = 6\times 24 \\ &= 8\times 18 = 9\times 16 = 12\times 12\,. \end{aligned} \qquad (2)$$

In fig. 4 we show the dynamic aperture resulting from sorting the same seed using the eight partitions shown in eq. 2. We observe that sorting in batches of 144, 72, or 48 magnets result in similar dynamic apertures. We therefore choose the 48-option to test the same 10 different seeds that lead to fig. 3 and depict the resulting dynamic apertures in fig. 5. Comparing this with fig. 3 we see that indeed the dynamic aperture is remarkably close to sorting the full 144 dipoles, but the spread between the different seed is increased, and one particularly bad seed also appears. Nevertheless, we conclude that *sorting the dipoles in batches of 48 magnets will be sufficient,* thus reducing the space needed to store the dipoles while the rest of the batch is delivered and measured.

The section of beamline that contains 48 dipoles has an unperturbed phase advance of 4π and contains two full betatron oscillations. To what extent cancellations of aberrations take place on that scale and the investigations of why sorting batches of 48 magnets is almost as good as sorting a full arc remains to be investigated in the future.

5 CONCLUSIONS AND OUTLOOK

We investigated sorting the LHC dipoles in one octant according to their skew quadrupole and their sextupole component. Sorting according to the former is entirely dominated by the systematic skew quadrupole component and sorting seems of little use. On the other hand sorting the sextupolar component improves the dynamic aperture of the bare lattice plus sextupole components in dipoles three fold, thus making it an attractive feature. Furthermore sorting only batches of 48 magnets yields almost as good results as sorting a full octant with 144 dipoles, reducing the need for intermediate storage space.

The dramatic threefold increase in dynamic aperture can only be considered qualitative at the present stage of the study, because we neglected the effects of other non-linear elements, such as the chromaticity sextupoles. Furthermore, effects of magnet misalignment and taking the full LHC with all octants into account need to be looked at. Moreover, the two-in-one design of the magnets requires simultaneous sorting of two rings. These effects will certainly reduce the gain in dynamic aperture due to sorting and will have to be investigated in the future.

The presented method is reasonably fast. Sorting an entire octant takes about a few minutes on a fast HP-750 workstation. The program could be cast into a small user-friendly PC version that reads a file with the multipoles, then quickly sorts the dipoles and returns a table with the sorted sequence.

Of course other sorting criteria and cost functions than those used can be envisioned, i.e. forcing the sextupoles which are located 180° in phase apart to be equal as to cancel the second order sextupolar aberrations. Preliminary test, however, show that the resulting dynamic aperture is similar to the one reported in the main part of this paper.

6 ACKNOWLEDGEMENTS

Discussions with W. Scandale and J.-P. Koutchouk are gratefully acknowledged.

7 REFERENCES

1. LHC Study Group, *LHC, The Large Hadron Collider Accelerator Project,* CERN/AC/93-03(LHC), November 1993.

2. W. Press, et.al., *Numerical Recipes,* Cambridge University Press, Cambridge, 1986.