LHC Project Note 309

January 28, 2003

The Influence of Computer Errors on Dynamic Aperture Results Using SixTrack

M. Hayes, E. McIntosh and F. Schmidt

Keywords: dynamic aperture, SixTrack

Summary

The SixTrack run environment is used to study LHC dynamic aperture. Results are presented from two different computer architectures (Alpha and Intel). The study focused on the differences between the results.

1 Introduction

Studies of the dynamic aperture (DA) of the LHC lattice often use the code SixTrack[1] which is fast and with its improved run environment[2] easy to use. With the current upgrade of computer capacity shifting the production to Intel-based processors away from traditional Alpha-based ones, it was felt necessary to launch a study comparing results from the two architectures. By architecture here it is meant the operating system, compiler and all other nuances as well as the hardware. The two platforms in question were:

- shiftnap: an Alpha-based set of twenty processors running the DEC OSF1 operating system.
- lxplus(lxnap): an Intel Pentium III based set of twenty dual processors running the Redhat 6.2 linux operating system.

A tested version of the LHC sequence was used[3] at injection energy, with only non-linear dipole errors turned on as well as b_3, b_4 and b_5 correction spool pieces. The same LHC definition was used on both architectures to examine the differences introduced by the architectures alone. The latest SixTrack v3.0 was used and its associated run environment. SixTrack is written in double precision FORTRAN77.

To calculate a dynamic aperture 300 seeds are launched split into five phase space angles ($15°$,$30°$,$45°$,$60°$ and $75°$). Each seed consists of 30 particles equally spaced between a series of defined sampling ranges. The standard sampling range is two units amplitude (e.g. 10 to 12 sigma). The dynamic aperture *for a seed* is defined as the smallest amplitude at which the particle

	15°	30°	45°	60°	75°
lxplus	12.3 (45)	11.3 (39)	11.4 (14)	11.9 (35)	12.7 (36)
shiftnap	12.3 (45)	11.4 (45)	11.8 (14)	11.9 (35)	12.8 (36)

Table 1: Dynamic aperture results for the LHC per angle for the shiftnap and lxplus clusters. The seed number causing this DA is given in brackets.

is lost within 100,000 turns. The dynamic aperture *of the machine* is defined as the minimum of all the dynamic apertures of all the seeds. This may be one number for the whole machine or, as in the case of this note, one number per angle.

2 Results

The minimum dynamic apertures for the LHC per angle as found by both systems are given in table 1. As can be seen the results do not vary by much with computer system. The seed number causing this minimum DA is given in brackets next to the value of the aperture. Although not of intrinsic interest in itself, it shows that in two cases the same seed produces a slightly different number and in one case a different seed is responsible for the DA in the two architectures. Although the former may be attributable to rounding effects between the two kinds of processors, the latter hints towards possibility of chaotic effects appearing and the lack of reproducibility of individual seeds.

The effect is startlingly clear in figure 1(a) which shows the difference in DA ($\Delta\sigma$) between the two systems, as defined by:

$$\Delta\sigma_{\mathrm{seed}} = \mathrm{DA}_{\mathrm{seed}}^{\mathrm{lxplus}} - \mathrm{DA}_{\mathrm{seed}}^{\mathrm{shiftnap}} \tag{1}$$

The figure shows a clear 'toothcomb' effect with the difference between teeth of $2/30$ (the initial sampling spacing). The toothcomb seems to build a Gaussian distribution with an RMS of 0.15, but with tails approaching 0.5 units. It is centered around zero. The distribution shows no particular relation to the angle as can be seen from figure 2.

Choosing a particularly bad case (seed 26 at $75°$) where $\Delta\sigma = 0.69$. The DA for both cases individually is 13.8 and 13.1 respectively for shiftnap and lxplus. The automatically generated plots from the run environment are shown in figures 3 and 4. Figure 3 shows at what final distance in phase space two initially adjacent particles finish for a given initial amplitude. An exponential increase in separation of the particles is a firm indication of the presence of chaos. The upper figure is the result from shiftnap and the lower one from lxplus. Both have the same gross features, but differ in details which would be due to small chaotic effects. Apart from an isolated plateau and the odd peak the main chaotic region starts at roughly 13 units of amplitude. Which is consistent with the fact that on both machines the particle is lost above this value. Figure 4 shows the survival time of a particle at a particular amplitude. Again the upper plot is from shiftnap and the lower from lxplus. The dynamic aperture found is indicated by the first dip found. Both sets of plots together show that the dynamic aperture seems to be located in the area of chaotic motion, but it's position found by the program may vary from architecture to architecture, *given the same chaotic area*. It can only jump in units given by the initial sampling of phase space, this in turn causes the toothcomb effect as seen in Fig. 1(a).

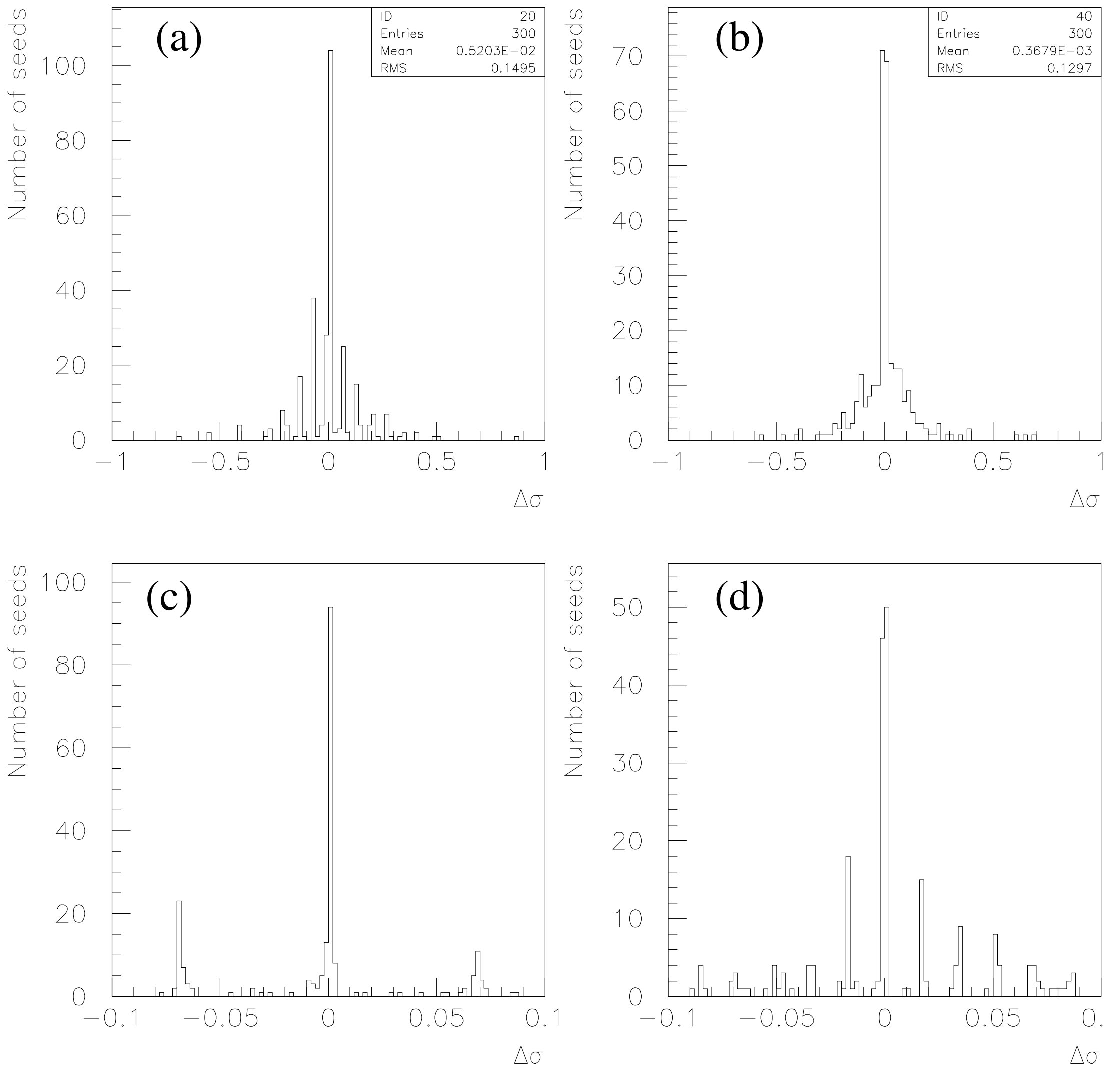

Figure 1: (a) is the difference between the DA found between systems per seed with the standard sampling of 30 particles per two units of amplitude. (b) is the same plot but for four times sampling (i.e. 30 particles per half unit of amplitude). (c) and (d) are enlargements of the center parts of (a) and (b) respectively to enhance the 'toothcomb' detail.

	15°	30°	45°	60°	75°
lxplus	12.3 (45)	11.3 (39)	11.4 (14)	11.9 (35)	11.6 (34)
shiftnap	12.3 (45)	11.4 (39)	11.6 (14)	11.8 (34)	11.6 (34)

Table 2: DA results for the LHC per angle for the shiftnap and lxplus clusters with four times initial sampling. The seed number causing this DA is given in brackets.

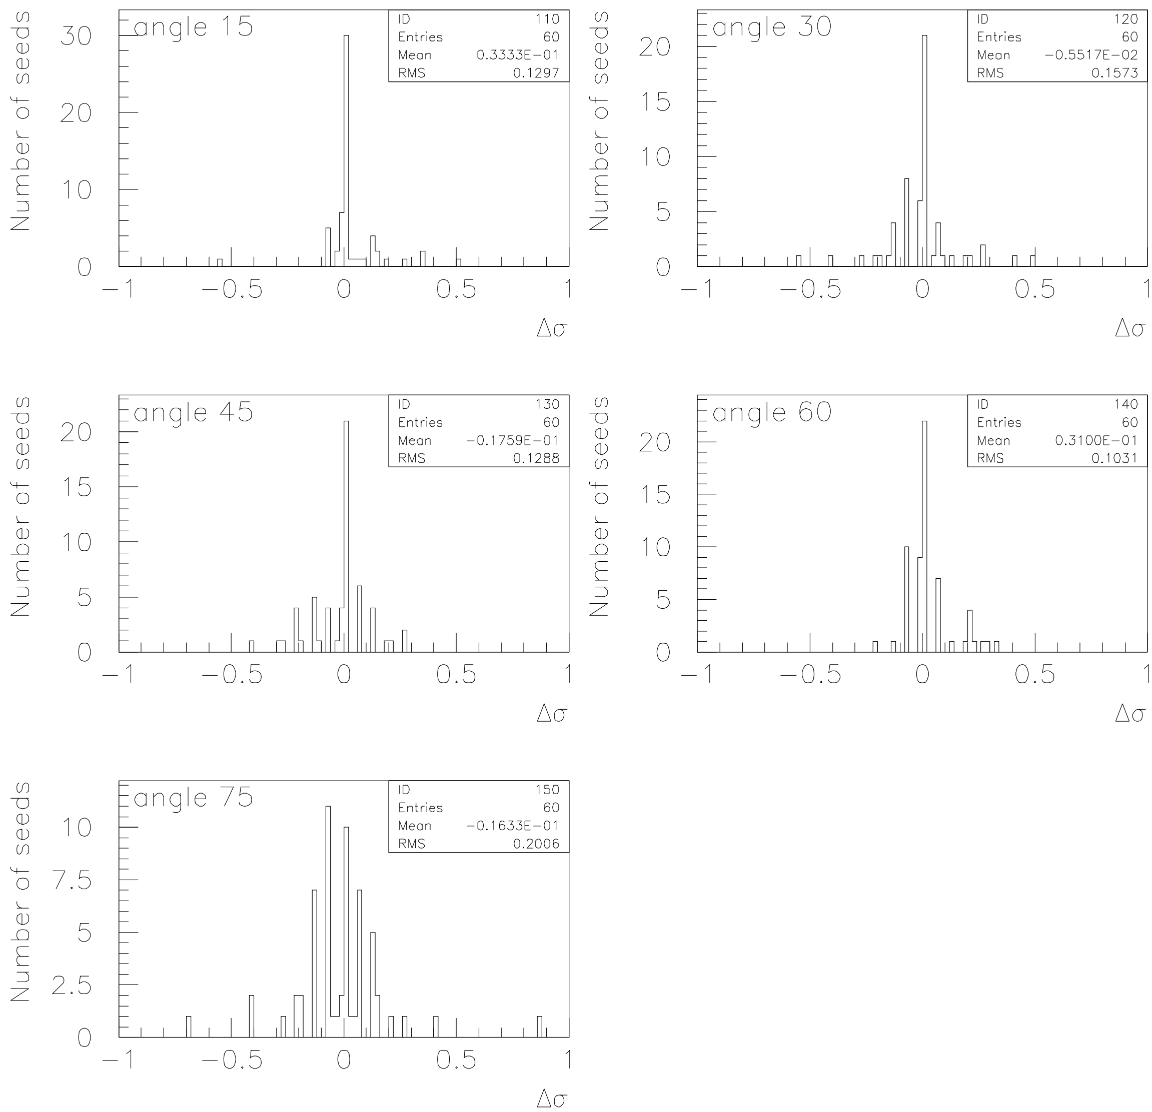

Figure 2: The difference between the DA found between systems per seed with the standard sampling of 30 particles per two beam sigma for each of the five angles studied.

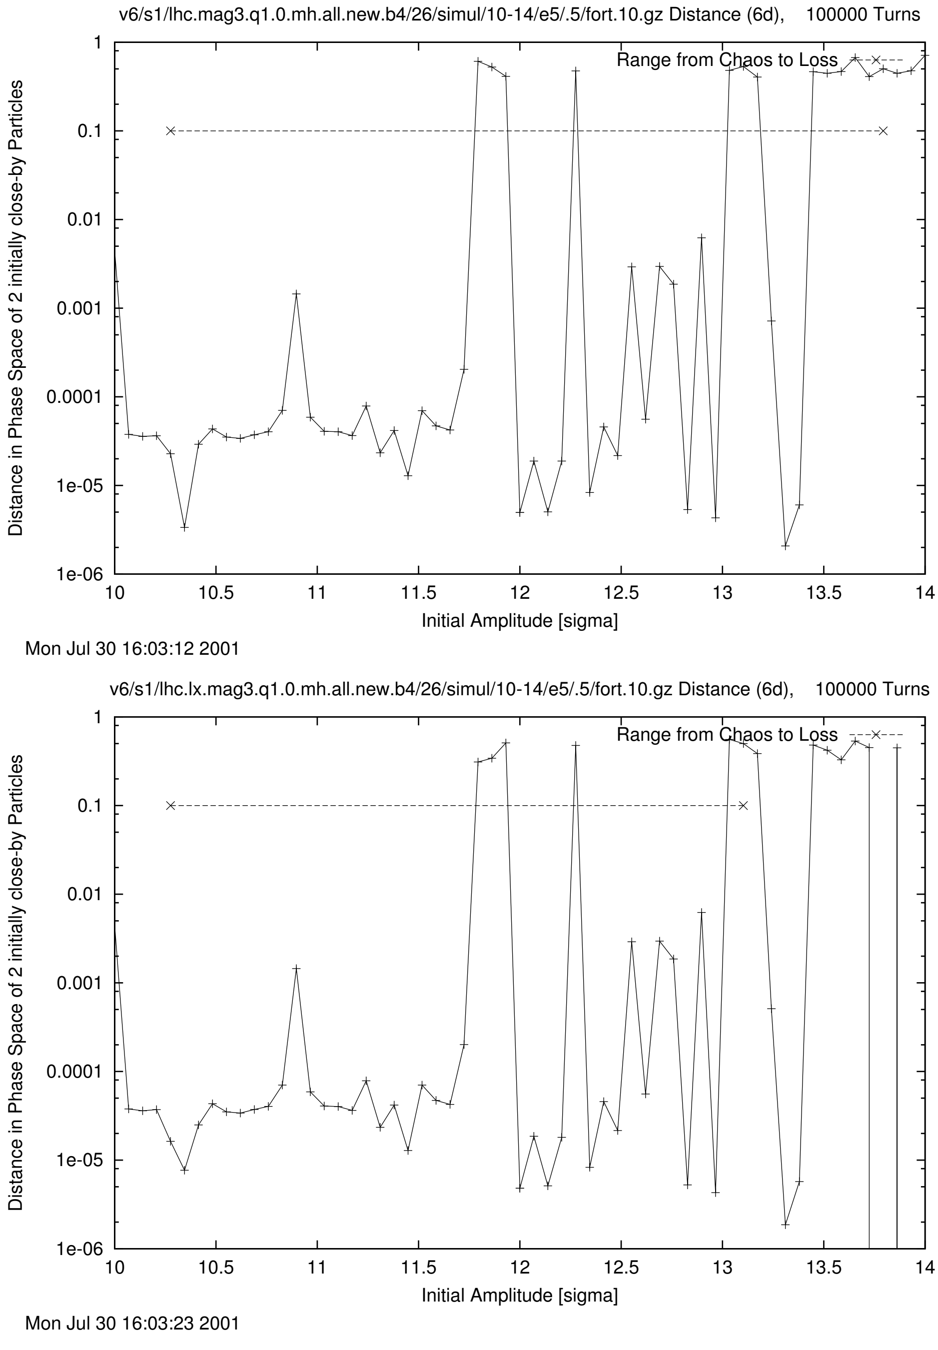

Figure 3: The two figures show the final distance in phase space of two initially adjacent particle, at a particular starting amplitude. Chaotic regions may be defined where the particles end up a distance apart larger than the indicated horizontal dashed line. The top figure is for a particular seed at a particular angle on shiftnap, the bottom is the identical case on lxplus. The DA for both cases are 13.8 and 13.1 respectively.

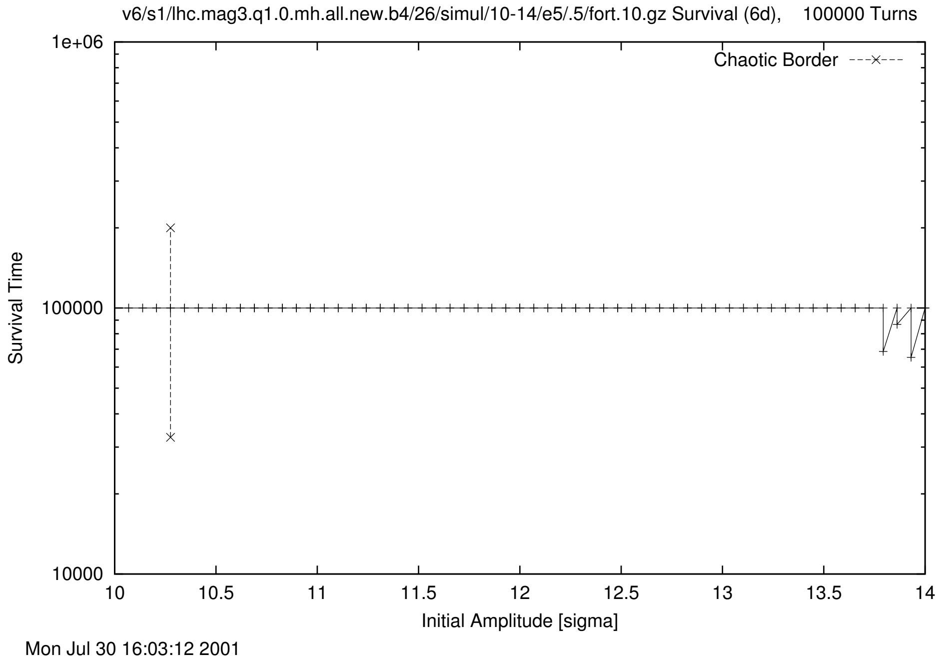

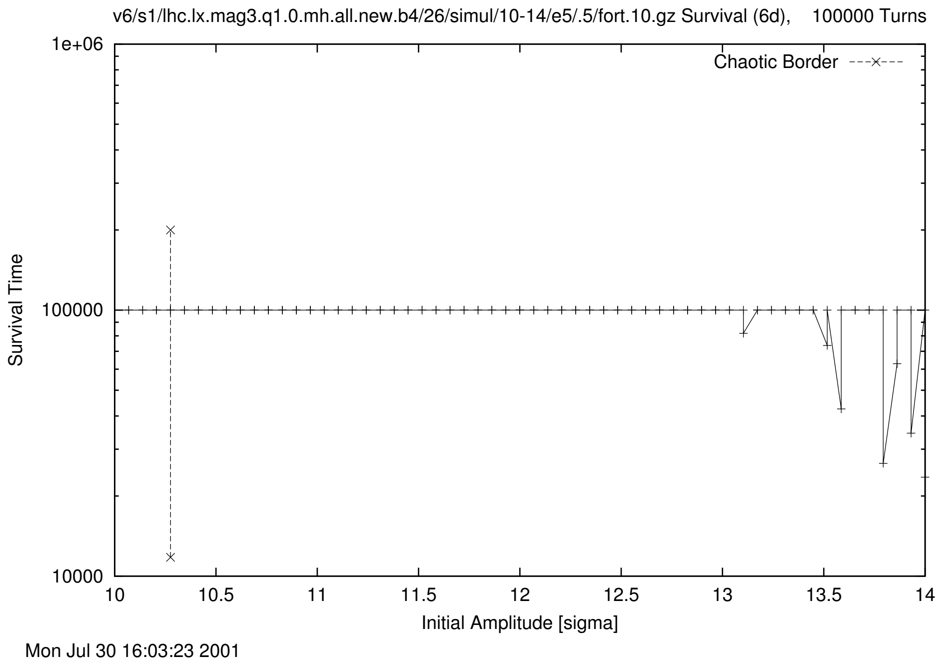

Figure 4: The two figures show the survival times for particles started at various amplitudes. The top figure is for a particular seed at a particular angle on shiftnap, the bottom the identical case on lxplus. The DA for both cases are 13.8 and 13.1 respectively (as can be clearly seen).

Architecture	initial particle x position
Intel	3.7002174852758360000000000000000
Alpha	3.700217485275843820602403866359964
Sun	3.700217485275816287071393162477762
	final particle x position (1000 turns)
Intel	−1.5663574304754480000000000000000
Alpha	−1.566357288619672605278765331604518
Sun	−1.566357230554381807507979829097167

Table 3: The starting and ending horizontal positions of a particle tracked for 1000 turns on the three different architectures.

To confirm this last statement a new run was done for the same 300 seeds but with a four times higher initial sampling of phase space (30 particles over half a unit were launched). The resulting DA for the machine per angle is shown in table 2. The same comments as for the normal sampling apply here. A comparison of the 300 seeds together is shown in 1(b). The same distribution with an RMS of 0.13 can be seen as in the normal sampling of (a). But the toothcomb effect is not visible at this scale. (c) and (d) are a zoom of the horizontal scale of (a) and (b). The toothcomb detail can be seen in (d) now with a spacing of $0.5/30$ equal to the initial sampling spacing. The spread of the differences is caused by two effects:

- The differences in the generated linear accelerator parameters (e.g. β functions) caused by differences in architecture.
- The differences during tracking through non-linear effects.

The former is the cause of the width of the spread and this can therefore not be removed with greater sampling, as is shown in figure 1(a) and (b).

3 Compilers

Using this knowledge and a test run showing a large DA difference between the two architectures (e.g. last graph in Fig. 2) a closer look was taken at the compilers involved. The Alpha version of the compiler performs pure 64 bits processing. The Intel version of the compiler performs 80 bit processing, which can in the hardware be reduced to 64 bit by adding an IEEE flag available on some compilers[4]. Finally a Sun version of SixTrack was run which also performs 64 bit processing. Results from one tracked particle for 1000 turns in identical set-ups between the three architectures can be seen in Table 3. It should be noted that the Alpha version of SixTrack uses fast maths libraries compared to the Intel version, so these differences are also included in this comparison.

The Sun and Alpha architectures are closest but still different. In terms of initial position they are identical to the 13th decimal place, but after only 1000 turns they differ in the 8th decimal place. The Intel architecture differs also in the 13th decimal place for the initial position, but after 1000 turns it is different in the 7th decimal place.

4 Conclusions

DA is sensitive to the architecture of the machine it is generated on. This is due to the physics associated with chaos and the initial sampling. However this ‘sensitivity’ has a finite spread which is within the normally associated error of 0.5 units as should be given with all DA studies. But it should be noted that the actual value of a DA of a seed may not be reproducible except on the same architecture as it was originally produced. Although compiler flags may be used to minimize the differences, this is not possible in practice due to the multitude of differences between hardware, compilers, libraries and even implementations of the IEEE standards. The smallest differences in the last decimal place are being enhanced by the chaotic nature of the studies concerned.

It should also be noted that this spread of differences is not as much reduced by increased initial sampling as expected due to the fact that the numerical representation of the accelerator also depends on the computer architecture.

5 Acknowledgments

The authors would like to thank the continuing support of the IT division, in particular Tim Smith and Ulrich Fuchs. The authors would also like to thank Werner Herr for useful discussions on the subject of computer precision.

References

[1] F. Schmidt, “SixTrack - User’s Reference Manual”, CERN/SL/94-56 (AP), updated March 2000.

[2] F. Schmidt, “Run Environment for SixTrack”, Beam Physics Note 53.

[3] `/afs/cern.ch/eng/lhc/optics/v6.2`

[4] Specifically IEEE754 specifications for 64 bit floating point arithmetic.