BEAM DYNAMICS SIMULATION FOR THE CTF3 DRIVE-BEAM ACCELERATOR

D. Schulte, CERN, 1211 Geneva, Switzerland

Abstract

A new CLIC Test Facility (CTF3) at CERN will serve to study the drive beam generation for the Compact Linear Collider (CLIC). CTF3 has to accelerate a $3.5\,\mathrm{A}$ electron beam in almost fully-loaded structures. The pulse contains more than 2000 bunches, one in every second RF bucket, and has a length of more than one μs. Different options for the lattice of the drive-beam accelerator are presented, based on FODO-cells and triplets as well as solenoids. The transverse stability is simulated, including the effects of beam jitter, alignment and beam-based correction.

1 INTRODUCTION

In the nominal stage of CTF3, the drive-beam accelerator will have eight klystrons, each feeding two 1 m-long structures. The structures are almost fully loaded, transferring more than $90\,\%$ of their input power to the beam. The average energy gain per structure is $\Delta E \approx 9.1\,\mathrm{MeV}$ [1]. The beam pulse consists of ten short trains of about 210 bunches each. The first train fills odd buckets, the immediately following second train fills even buckets; this pattern is then repeated. An RF-deflector at half the linac frequency is used to separate the trains after acceleration [2]. The initial beam energy is $E_0 \approx 26\,\mathrm{MeV}$, the final beam energy $E_f \approx 170\,\mathrm{MeV}$, the bunch charge $q = 2.33\,\mathrm{nC}$, its length $\sigma_x \approx 1.5\,\mathrm{mm}$ [3] and the transverse normalised emittances are $\epsilon_x^* = \epsilon_y^* = 100\,\mu\mathrm{m}$.

2 STRUCTURE MODEL

The simulations below have been performed using PLACET [4]. The long-range transverse wakefield is represented by the lowest two dipole modes of each cell. These have been calculated neglecting the coupling between cells and the effect of the damping waveguides [5]. The damping of the lowest dipole mode has been found [6] to be in the range $Q = 11$ to $Q = 19$ for perfect loads. In the simulation, the modes are confined to their cells, which allows one to take into account the angle of the beam trajectory in the structure. The loss factors used in the simulation are $50\,\%$ larger than in [5]. This is to account for the effect of higher-order modes. Also, the damping is conservative in the simulation; $Q = 30$ and $Q = 400$ are used for the lowest and the second dipole band. The short-range longitudinal [5] and transverse [7] wakefields have been calculated and are included in the simulation. Almost perfect compensation of the long-range longitudinal wakefields is predicted [1].

Quadrupole wakefields may be important and have been implemented in PLACET. The corresponding modes have not yet been calculated but need to be included in the simulation as soon as they are available.

3 LATTICES

Three different lattices were investigated. One consists of simple FODO-cells, with one structure between each pair of quadrupoles. The other two lattices are based on triplets. In one case (called T1 below), one structure was placed between two triplets; in the other case two structures (T2). The weaker triplet lattice (T1) and the FODO lattice are roughly comparable in length and cost, whereas the strong triplet lattice (T2) is significantly longer and more costly.

In the FODO lattice, the phase advance is $\mu = 102°$ per cell, with a quadrupole spacing of $2\,\mathrm{m}$. In T2 one has $\mu_x = 97°$ and $\mu_y = 110°$, and a distance of $4.2\,\mathrm{m}$ between triplets. The sum of the integrated strengths of the outer two magnets is slightly larger than that of the inner one. With this arrangement, the horizontal and the vertical beta-functions are equal in the accelerating structures, and the energy acceptance of the lattice is markedly improved. For T1 the phase advances are $\mu_x = 84°$ and $\mu_y = 108°$ for a triplet spacing of $3\,\mathrm{m}$. The transverse acceptance is $4.2\,\sigma$ for the FODO lattice, $4.9\,\sigma$ for T2 and $5.8\,\sigma$ for T1.

Since the beams have to be compressed after the acceleration, the RF-phase cannot be used to optimise the beam transport. It must be chosen to achieve the required compression and to limit the energy spread of the beam before the combiner ring to the latter's energy acceptance. An RF phase $\Phi_{RF} = 6°$ is used in the following.

4 TRANSVERSE BEAM JITTER

No estimate of the transverse jitter of the incoming beam exists. Therefore, only the jitter amplification is calculated. In the simulation, each bunch is cut into slices; the beam is set to an offset of Δx and tracked through the linac. The normalised amplification factor A for a slice is defined as

$$A = \frac{\sigma_{x,0}}{\Delta x} \sqrt{\left(\frac{x_f}{\sigma_{x,f}}\right)^2 + \left(\frac{x'_f}{\sigma_{x',f}}\right)^2}$$

Here, $\sigma_{x,0}$ and $\sigma_{x,f}$ are initial and final beam size, $\sigma_{x',0}$ and $\sigma_{x',f}$ are initial and final beam divergence, Δx is the initial beam offset and x_f and x'_f are the final position and angle of the centre of the slice. For a slice with nominal energy and without wakefield effects, one has $A = 1$. The
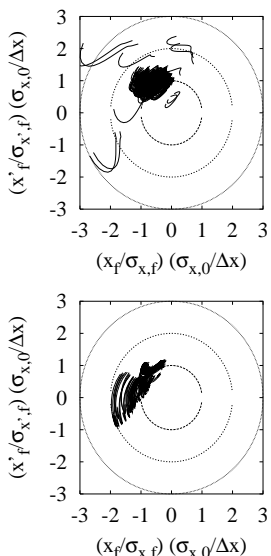
Figure 1: The amplifiaction factor of the beam at the end of the drive-beam accelerator, using the FODO lattice, without a ramp (left) and with a ramp (right). A mono-energetic beam without wakefields should stay on the innermost circle.
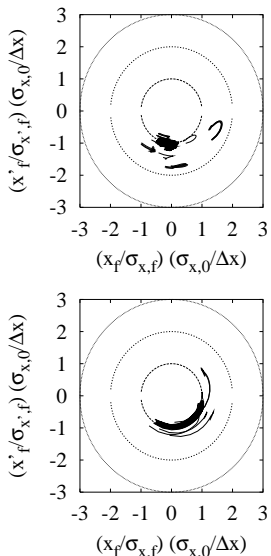

Figure 2: The beam at the end of the drive-beam accelerator in a triplet lattice. On the left-hand side T2, on the right-hand side T1.

maximum amplification factor $\hat{A}$ is the maximum over all slices. The left-hand side of Fig. 1 shows the bunches at the end of the accelerator using the FODO lattice. Different quadrupole strengths were used to find the best phase advance. Some bunches are kicked significantly; the maximum amplification is $\hat{A} = 3.7$. Without knowledge of the acceptance downstream and the size of the beam jitter, it is not possible to decide whether the amplification is acceptable. Within the linac, even a large jitter of $\Delta x = \sigma_x$ does not lead to beam loss.

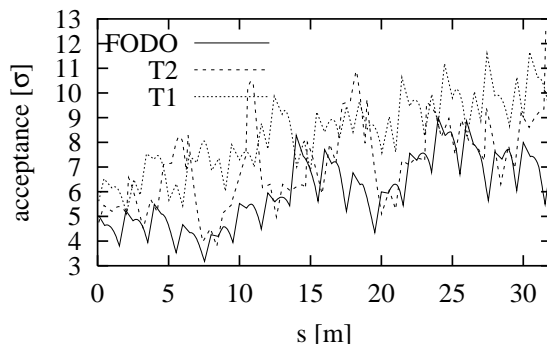The first few bunches in each train are kicked particu-

Figure 3: The minimum acceptance along the linac, with a gradient error. For each lattice, 100 machines have been simulated and their minimum acceptance at each point is plotted.

larly hard. This can be prevented by adding charge ramps. Towards the end of a train that fills even buckets, the bunch charge is slowly decreased from the nominal bunch charge to zero. At the same time one increases the charge in the odd buckets from zero to nominal, to keep the beam current constant. Thus the two consecutive trains practically overlap. On the right-hand side of Fig. 1, one can see that in this case all bunches are well confined, with a maximum amplification of $\hat{A} = 2$.

In the triplet lattices, the horizontal plane has a larger jitter amplification than the vertical one. But even the horizontal amplifications are significantly smaller than in the FODO lattice. Figure 2 shows the examples of a pulse without charge ramps, the amplification factors being 1.8 (T2) and 1.5 (T1). With charge ramps, they are reduced to 1.5 and 1.3. If the beam jitters significantly, the triplet lattices are markedly better than the FODO lattice.

5 BEAM-BASED ALIGNMENT

To keep operation as simple as possible, only one-to-one correction is considered. All elements are assumed to be scattered around a straight line following a normal distribution with $\sigma = 200 \, \mu$m. In the FODO lattice, corrector dipoles are located after each quadrupole and beam position monitors (BPM) are placed in front of each quadrupole. In the triplet lattices, the corrector dipoles are positioned after the triplets and the BPMs are positioned in front and after the triplets. The correctors are used to bring the average beam position to zero in the BPMs. For each case, 100 different machines are simulated. The small growths of about $0.5\,\%$ are almost the same for all lattices.

6 GRADIENT AND PHASE ERRORS

The limit on the variation of the bunch energy is $1\,\%$ [8], much smaller than the single-bunch energy spread. In normal operation, the additional dispersive effects are therefore small. Static local energy errors are of more concern and are discussed here.
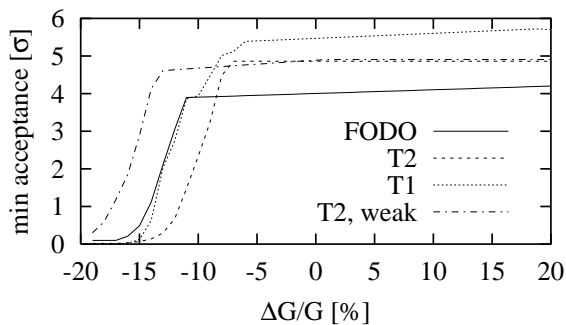
Figure 4: The minimum acceptance of the linac, as a function of the RF-gradient error; 20 machines were simulated for each case.

The initial and final beam energy can be well measured, and from this the average gradient can be derived. A local variation of the gradient is more difficult to detect. It will lead to a quadrupole strength that is not adapted to the beam energy. The worst case is too low a gradient in the first two structures, which are fed by one klystron. In the simulation, 100 different machines with a gradient in the first two structures that is too low by 10% (20%) are corrected with the beam. The emittance growth found after correction is 1% (5%) in the FODO lattice and 0.5% (2%) in T1, which seems to be sufficiently low. In T2, the value for a 10% error is small, 2%, but for an error of 20% it starts to be large: 14%. The transverse acceptance is reduced to $3.8\,\sigma$ ($3.2\,\sigma$) in the FODO lattice, $4.7\,\sigma$ ($3.8\,\sigma$) in T2 and to $5.2\,\sigma$ ($5.1\,\sigma$) in T1. Figure 3 shows the acceptance for a gradient error of 20%. For the FODO lattice, and to a lesser degree also for T2, one starts to worry about beam losses. However, an error of 10% seems acceptable with all lattices. To be able to use the FODO lattice or T2, it necessary to measure the local gradient to better than 10%, to be able to correct the lattice accordingly. For T1, a precision of 20% is sufficient.

The RF power produced by a klystron has a systematic phase variation during the pulse. One hopes to correct this effect globally, but local variations will remain. To estimate their importance, a linear change in phase of $20°$ over the pulse is assumed for the two structures driven by one klystron. The next pair has an exactly opposite phase variation. The resulting bunch-to-bunch energy spread is 2%, full width, which is not acceptable in the combiner ring; so a better compensation would be needed. In contrast, the emittance growth seems acceptable with about 1.5% averaged over 100 machines for all lattices; the acceptance is hardly decreased. This phase variation does not cause significant transverse effects.

7 ENERGY ACCEPTANCE

During commissioning of the linac, large energy errors may occur. To study the sensitivity to this, 20 machines were simulated for each lattice in the following way: the linac is corrected with a nominal beam; then the incom-
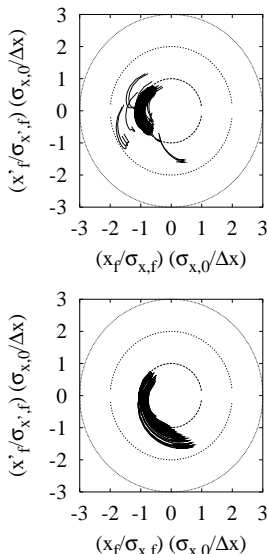

Figure 5: The amplification of beam jitter with and without the charge ramp for the lattice T2 including the injector solenoid.

ing beam is assumed to be accelerated at a different gradient. Figure 4 shows the minimum transverse acceptance of the three lattices as a function of the RF-gradient error. The final energy error is about 1.4 times larger than that of the RF-gradient, since the beam loading does not change. The FODO lattice and T1 have a comparable energy acceptance, whereas that of T2 is slightly smaller. By reducing the focal strength, the energy acceptance can be further improved at the cost of higher transverse wakefield effects. By reducing the strength of T2 to $\mu_x = 83°$ and $\mu_y = 94°$, the energy acceptance becomes larger than that of the FODO lattice. The maximum amplification of an initial jitter increases from 1.8 to 2.6 but is still smaller than the factor 3.7 in the FODO lattice. With each lattice, the linac energy acceptance is largely sufficient during normal operation.

8 SOLENOID

PLACET has been modified to also simulate the effects of solenoids with acceleration. This allows to include the last two structures of the injector which are placed inside a solenoid. Two triplets are used to match the end of the solenoid to the T2 version of the drive-beam accelerator. The field of the solenoid is $0.2\,T$ and its length is chosen such that a horizontal jitter of the nominal beam leads to a final horizontal offset. The end fields of the solenoid are modelled as thin lenses. Neither space charge nor the difference of the particle velocities from the speed of light are taken into account, but the wakefields are considered, in contrast to calculations done with PARMELA [9].

Figure 5 shows the amplification factor. While there is some contribution from the structures in the solenoid, the overall amplification seems still acceptable.

9   CONCLUSION

The simulations show that the lattices considered here can be acceptable; the best is the strong triplet lattice T1. The triplet lattice T2 seems to be a better choice than the FODO lattice. The FODO lattice is less expensive than T2, which is much cheaper than T1. To find the best compromise, more information is needed. For the FODO lattice the ramps have to be studied in more detail. For all lattices, the matching from the injector to the linac and from the linac to the combiner ring needs to be understood.

10   REFERENCES

[1] I. Syratchev. Private communication.

[2] D. Schulte. The Drive-Beam Accelerator of CLIC. *Proceedings of Linac 1998, Chicago, USA* and *CERN/PS 98-042 (LP)* (1998).

[3] L. Rinolfi. Private communication.

[4] D. Schulte. PLACET: A Program to Simulate Drive Beams. *Proceeding of EPAC 2000, Wien, Austria* and *CERN-PS-2000-028 (AE)* (2000).

[5] L. Thorndahl. In: The CLIC RF Power Source. *CERN 99-06* (1999).

[6] E. Jensen, A. Millich and L. Thorndahl. Private communication.

[7] A. Millich. Private communication.

[8] R. Corsini. Private communication.

[9] F. Zhou. To be published as a CTF3-Note.