Examination of the CLIC Drive Beam Pipe Design for Thermal Distortion Caused by Distributed Beam Loss

Colin D. Johnson (CERN, Geneva, Switzerland) and Peter K. Kloeppel (TJNAF, Newport News, VA)

Abstract. Beam transport programs are widely used to estimate the distribution of power deposited in accelerator structures by particle beams, either intentionally as for targets or beam dumps or accidentally owing to beam loss incidents. While this is usually adequate for considerations of radiation safety, it does not reveal the expected temperature rise and its effect on structural integrity. To find this, thermal diffusion must be taken into account, requiring another step in the analysis. The method that has been proposed is to use the output of a transport program, perhaps modified, as input for a finite element analysis program that can solve the thermal diffusion equation.

At CERN, the design of the CLIC beam pipe has been treated in this fashion. The power distribution produced in the walls by a distributed beam loss was found according to the widely-used electron shower code EGS4. The distribution of power density was then used to form the input for the finite element analysis program ANSYS, which was able to find the expected temperature rise and the resulting thermal distortion. As a result of these studies, the beam pipe design can be modified to include features that will counteract such distortion.

Introduction

The Compact LInear Collider (CLIC) study seeks to explore the feasibility of building a two-beam linear collider operating at the accelerating frequency of 30 GHz, with energy and luminosity at the interaction point of 1 TeV (center-of-mass) and 10^{34} $cm^{-2}s^{-1}$, respectively[1]. The two-beam collider is a speculative design in which microwave power for acceleration in the main linac is derived from a drive linac running alongside[2]. The power in a relatively low-energy ($\approx$ 3 GeV) but high-current (20 mA mean) electron drive beam is converted to microwave power that is fed into the main linac to achieve acceleration of a much lower electron or positron beam current to a much higher energy. The drive beam is decelerated in this process in a rather uneven fashion that leads to considerable energy spread (up to a factor 7) within the beam pulse. The conceptual simplicity of this scheme has considerable appeal, albeit with some contra-indications, compared to the conventional linac for which the power is generated by a very large number of klystrons. However, as the total continuous drive beam power required to achieve these high energy and luminosity values is in the region of 63 MW (for each of two drive beams), extremely high power levels must be included in the study.

Just to set the scale: each 3 GeV beam at 63 MW is equivalent to the electrical consumption of a large town of 100,000 inhabitants, or that required to melt 100 kg of copper per second. The required continuous beam current is 100 times that of the Thomas Jefferson facility, which operates at a similar energy. Finally, the microwave power, although at a different frequency, is of the same order as that specified for the ground-based FEL of the US Strategic Defense Initiative[3], which, unlike the linear collider, was never intended to run continuously.

The CLIC drive beam can be considered in three separate stages: the generation of multi-bunch beams of total charge 30μC at a repetition rate of 700 Hz; the acceleration to 3 GeV (i.e. to a total energy of 90 kJ) in superconducting linacs at a frequency of between 200 and 350 MHz; and finally the transport, with deceleration, of the resulting high-power beam through several kilometers of drive linac in which the beam power is converted to microwave power via relatively weak interaction with power-generating transfer structures that, together with the focusing lattice quadrupoles, constitute the drive linac. It is during this deceleration that the large energy spread develops, since the leading part of the beam pulse remains close to the initial 3 GeV whilst the cores of the bunches in the main body of the pulse train lose up to 90% of their starting energy. At all stages in this process, the requirements for high beam power and high power conversion efficiency push the technology beyond present limits; to assess the extent to which these limits can be overcome is the brief of the study group. In this paper we concentrate on the possible consequences of electron losses during the beam's passage through the drive linac.

The method that has been used for the analysis is a combination of EGS4, which can predict the distribution of the power that is deposited in the structure by a missteered beam, and ANSYS, which is used to find the resulting temperature distribution and thermal distortions. Although the programs are not readily combined, we have found meaningful results by the method.

Monte Carlo (MC) programs are widely used for the simulation of all kinds of radiation effects, including shielding penetration and energy deposition. The simulations have the common feature that the physical properties of the shielding or absorbing medium are unaffected in the calculation by the radiation. In reality, the medium can be more or less strongly affected by the radiation, beginning with the thermal effects. It may be distorted by the temperature changes, or it may be melted or even vaporized. These effects are beyond the capabilities of the standard radiation programs.

Finite element (FE) programs have been developed to deal with some temperature effects. Given the distribution of deposited power, these programs can find the temperature distribution, taking into account phase changes (melting or vaporization). They therefore can calculate thermal distortions and even limits imposed by power deposition on structural stability and integrity. The results can describe either equilibrium or the evolution in time of the described effects.

The merger of these two types of programs is clearly desirable for some problems. In the design of accelerator safety systems, for example, we have to determine the time interval in which a beam loss incident will vaporize the shielding, so we can define a rational shutdown system that will react to the beam loss in time to protect personnel. At present, such mergers can be performed only imperfectly. The description of a physical object in an MC program usually breaks it up into a rather small number of homogeneous regions with easily described surfaces as boundaries. The regions should be as large as possible, in order to improve counting statistics. In contrast, FE programs basically attempt to find solutions to partial differential equations, a procedure that is improved by reducing the separations between the nodes that define the geometry of the object. As a result of these conflicting demands, a geometric description that would be regarded as quite large and complicated for MC methods would be very small by FE standards.

Despite this fundamental incompatibility, we have been able to analyze some accelerator structures for their thermal behavior under conditions of beam loss. We have used the familiar MC program EGS4 to find the distribution of the power deposited by the beam in an object that has been modeled for the program, then have used the power distribution as input to the FE program ANSYS to find the thermal response of the same model. The distortion is the result of heating caused by distributed beam loss, such as could be expected during beam setup operations. This shows some of the power of the method, in that it has demonstrated some of the problems that would be encountered with a design without the need for construction and physical tests.

Calculation of the deposited power

The drive linac has a basically simple form, comprising a regular sequence of microwave transfer structures interspersed with strong-focusing quadrupoles. The inner radius of the transfer structures, that also form the vacuum pipe, is in the region of 10 mm, more or less, according to the design, but with a recent trend towards the highest value of 13 mm. Despite the relatively low electrical impedance of the transfer structures, transverse resonant wakefields and to a lesser extent resistive-wall wakes are the chief agents of beam loss. Other losses are due to transverse field inhomogeneity at large radii and straightforward scraping of the beam towards the downstream end of the linac, where, due to deceleration, the real emittance of the cores of the electron bunches increases to fill the acceptance. Quadrupole misalignment could also cause beam loss at high energy, but it turns out that the alignment required to maintain beam trajectories close enough to the machine axis to avoid destructive transverse wakes is much tighter than that needed to ensure free passage of a wake-free beam. Sudden misalignment due to electrical failure in one or more quadrupoles or misalignment due to minor earth tremors could, of course, produce a beam steering error with potentially catastrophic results. A machine protection system that switches off the beam between multi-bunch trains will be needed to look after these problems. The energy density within one train must be reduced below the level that could cause permanent damage to the beam pipe. This implies the choice of an initial high-energy emittance that is sufficiently large to achieve the required low energy density, or the use of beam spoilers, although it is hard to see how these could be introduced all along the linac. A consequence of this requirement is that there is little or no reserve aperture within the beam pipe at the low-energy end, whichever solution is adopted. Taking all these effects into consideration, and bearing in mind that this is an ongoing study, it appears that the microwave transfer structures will need to be aligned in the transverse directions to within a precision of 50 to 100 μm. The quadrupole alignment will have to be much better than that, but they will be supported independently of the beam pipe and so their alignment will not necessarily be directly disturbed by the heat generated from beam losses. This study concentrates on the the possible misalignments of transfer structures that could result from low-level beam losses that may occur more or less anywhere along the drive linac during normal operation, bearing in mind that shutting off the beam may not necessarily be the correct response to minor losses, since below a certain threshold the losses will have to be actively reduced by very fine tuning at full beam power, or simply accommodated by the cooling system.

Beam damage potential

The potential for damage by an electron beam, in terms of heat deposition leading to strain or melting, depends upon its current and energy. Some typical high-intensity electron beam applications are listed in Table 1, and the 'damage potential' for *continuous beam loss* is roughly indicated by dividing the lost power over the length over which the loss is contained, which for electrons above the critical energy is simply proportional to the beam current. Table 1 serves to indicate the particularly high damage potential of the CLIC drive beam described above. In this table, two drive beams are listed; the second is a variant in which the drive beam is split into 10 beams of lower charge and energy, each supplying power to one-tenth of the main linac. The damage potential is reduced by a factor 5.

Table 1. Comparison of beam power and 'damage potential' for some RF power sources and linacs

Project	Device	Beam energy (MeV)	Time-average current (mA)	Mean power (kW)	Range of electrons in copper (m)	Damage potential–power/range (kW/m)
SLAC NLC	KL-1 Klystron	0.44	125	55.5	2×10^{-4}	2.8×10^{5}
NLC TBA	R.K.	10	22	216	7×10^{-3}	31×10^{4}
TJNAF	Linac	4×10^{3}max	0.2	800	0.28	3.6×10^{3}
CLIC	Drive linac–1	3×10^{3}	21	63×10^{3}	0.17	3.7×10^{5}
CLIC	Drive linac–2	1.5×10^{3}	4.55	6.8×10^{3}	0.085	8.0×10^{4}
TESLA	Main linac	2.5×10^{5}	0.06	16×10^{3}	14.3	1.1×10^{3}

The waveguide considered in the analysis is shown in cross section in Figure 1. It is a four- waveguide or four-slit structure made of copper, with water in its cooling tubes. In order to produce the desired RF field in the cavity, the inner surfaces indicated by the arrows are made to vary sinusoidally in the axial direction, with amplitude 1.883 mm, peak-to-peak. These "corrugations" cannot be handled easily in EGS4, so the model treats the surfaces as flat, at the mean radius from the centerline.

The mesh that was used was chosen as a compromise between the conflicting demands of EGS4 and ANSYS. As can be seen in the figures that show the results, the tubes for cooling water are not distinguished in the mesh. Although both programs would have allowed a more realistic shape, the incremental improvement in accuracy of the results was so slight as not to be worth the effort; in EGS4, the beam is absorbed before it penetrates to the water, and in ANSYS the temperature in the region of the water tubes does not deviate from the base temperature of 300 K. Axially, the waveguide was divided for EGS4 into 10 equal segments, each 10 cm in length. In ANSYS, exploiting the axial symmetry, a single 1 mm axial segment was used.

Two distinct modes of beam loss were considered. In the one, the beam was considered to strike a waveguide ("corrugated") surface at near-grazing incidence. In the other, it struck one of the noses at the apex. The points of impact of the incident particles in the beam were randomly distributed along the length of the entire 1-meter segment. In each case, the angle between the beam and the z axis was 0.001, and the transverse dimensions of the beam were small compared with the size of the mesh. The average energy left in each region, per incident electron, was then found by conventional EGS4 programming. The number of events considered (200000) was large enough to give statistically reliable results. EGS4 showed marked end effects in the first, second, and tenth axial regions, which were therefore discarded. The axial averages of the remaining seven segments were used as the basic input to ANSYS.

Calculation of the thermal effects

A short, ad hoc Fortran program was written to convert the EGS4 output into a form suitable for ANSYS. This program (a) performed the axial averaging, and (b) scaled the output, expressed in MeV per milliliter per incident electron, to units appropriate for ANSYS, watts per cubic meter for beam loss of 10 microamperes per meter. The choice of 10 microamperes per meter for beam loss was only for convenience, and is not significant. Because no nonlinear effects were included in the analysis, the final results are proportional to the current and can be scaled to any desired value.

The ANSYS treatment of each case is in two parts. In the first, the temperature distribution is found. Then, with the temperatures known, the distortions can be calculated. For the present problem, in which only the equilibrium properties are of interest, separating the parts of the calculation in this manner allows the program to run faster than would be possible if the temperature distribution and the resulting distortion had to be found simultaneously. (ANSYS can handle simultaneous calculations when they are necessary, as for example in transient analyses. The price that must be paid is extra execution time.)

Case 1: beam striking corrugated surface

The results of the analysis for the beam impinging on the waveguide surface are shown in figures 2 through 4. Figure 2 shows the temperature distribution, which follows the input power distribution rather closely, as could be expected. The maximum temperature of 318 K is reached right at the projected center of the point of impact of the beam. It falls quite rapidly for points away from the beam, and for most of the region it does not differ appreciably from the temperature of the cooling water, 300 K.

The resulting displacements are very much what would be expected for such a change of temperature. The distortions are highly exaggerated in the figures, as they are typically only about 1.5 μm at maximum. The x and y components of the distortions are axially symmetric (figures 3 and 4, respectively), whereas the z component has reflection symmetry about the axial midplane of the structure. The axial results are not shown, but were included in the ANSYS output as a check.

Case 2: beam striking the nose at its apex

The second case, in which the beam is supposed to strike at the apex of one of the nosepieces, is quite similar to the first. The major difference is that the temperature rise is somewhat smaller, only 12 K, presumably because the impact point is nearer to the cooling water. The results are shown in figures 5 and 6. Figure 5 is the temperature distribution. Figure 6 is the resulting horizontal distortion. Note that midplane symmetry cannot be assumed here as it was for the first case, so the waveguide section must be modeled in full 360°. The displacements are symmetrical about the 45° line; this result is not shown here. The maximum distortion is about 1.0 μm.

Operational strategies

The radial distortions that are predicted by this method, typically a maximum of 1 μm at a current loss of 10 μA per meter, are small but significant in the context of the required beam pipe alignment. The CLIC Transfer Structure is primarily a long, rather slender element that ideally should be made of solid copper. In this case a 1 to 2° temperature difference across the structure would cause it to bend beyond the tolerance limit unless constrained by an external support. Such a temperature differential could result from relatively low-level beam loss—well below the level that should be accommodated during beam tuning and even, perhaps, normal operation. In addition to the provision of a suitably rigid support, the cooling water circuit will be designed to minimize transverse temperature gradients.

Faced with the need to contemplate a multi-megawatt beam handling strategy, one looks for ways of reducing the mean power while preserving the beam qualities for the setting-up phase. One typical scenario starts with machine component alignment employing refined mechanical alignment techniques followed by fine tuning based upon beam observation on low-intensity pilot pulses (of more than one energy so as to set up a dispersion-free quadrupole alignment). Then, at very low repetition rate, the beam intensity would be increased to test the alignment against the wake-dominated effects. Then beam-loading compensation schemes in the acceleration sections would be tested by running at full repetition rate for short periods (burst mode) of say 1 second per minute. Up to this point the beam could be measured for alignment and stability and then dumped on a 1 MW beam dump at the entrance to the drive linac. It is unlikely that one would contemplate a 60 MW beam dump, so the beam at full power (i.e. 700 Hz repetition rate) could only be turned on when the alignment of the drive linac had been assured to the necessary precision. As the drive linac aperture is low compared to the beam emittance, particularly at the low-energy end, this operational strategy would inevitably result in some low-level losses to the transfer structures.

Loss patterns

The patterns of electron losses from the beam halo due to residual alignment errors, the collective and single-particle effects mentioned above, are being extensively simulated using all of the resources at our disposal. We shall soon have experimental verification of these simulations for the particular case of the low-energy dive beam which is a part of the CLIC Test Facility (CTF2) [4]. Meanwhile we can make the following general observations:

i) Because of adiabatic growth of the drive beam as its energy is reduced and the difficulty of arranging a beam transport lattice to accept the large energy spread, the beam will tend to fill the aperture at the far end of the drive linac, with inevitable losses that must be incorporated into the design and operational strategies.

ii) Losses elsewhere along the drive beam will be due initially to quadrupole misalignments and will occur at positions where the beam size is greatest (β_{max}), and these will be restricted to regions of the order of 1 m length. However, the location of β_{max} will depend on the electron energy, and, as the large energy spread develops, these locations may be anywhere along the linac. Electrons in a particular energy range that meet quadrupole misalignments at a combination of betatron phases that conspire to generate large orbit errors may be lost, whilst other electrons in the same bunch may pass.

iii) When losses do occur, a sustained loss of 0.1% will correspond to a current of 20μA, and *pro rata* for larger losses. The results of the EGS/ANSYS simulations can be scaled to these values.

Some conclusions

Minor beam loss (in terms of fraction of beam current) will represent a fairly sizable power loss. One part in a thousand is 60 kW, and this could be deposited more-or-less anywhere along the drive linac, unless special provisions are made to localize losses, and this appears at first sight to be excluded.

Where losses do occur there will be significant distortion of the structures due to thermal stress, and this could, via transfer wake effects, cause further losses and even a run-away loss situation. Some strain-reducing support structures must be envisaged, and possibly the use of a material for the transfer structures having a low thermal coefficient of expansion–copper plated to retain the required electrical properties.

To date in the study, two designs of CLIC transfer structure have been analyzed: a two-waveguide low-impedance structure and more recently a higher-impedance four-waveguide structure. These beam-loss studies have clearly shown that the latter is unsuitable for use as a CLIC transfer structure, although it will be used for the CLIC Test Facility. As the study progresses, the exact choice of transfer structure will be made taking into account several factors, but always including the requirements of the beam loss considerations.

References

[1] H. Braun, R. Corsini, J. P. Delahaye, G. Guignard, C. Johnson, J. Madsen, W. Schnell, L. Thorndahl, J. Wilson, W. Wuensch, B. Zotter, "CLIC–a Compact and Efficient High Energy Linear Collider," CERN-SL-95-26, May 1995.

[2] B. Autin, K. Bane, R. Bossart, R. Corsini, J. P. Delahaye, G. Guignard, C. Johnson, A. Mikhailichenko, J. Nation, L. Thorndahl, I. Wilson, "CLIC Drive Beam Generation: a Feasibility Study," CERN-PS-94-16 and EPAC94, 1994.

[3] R. Warren, "Star Wars and the FEL," personal communication.

[4] The CLIC Study Group, "CTF2 Design Report," CERN-PS-96-14, June 1996.

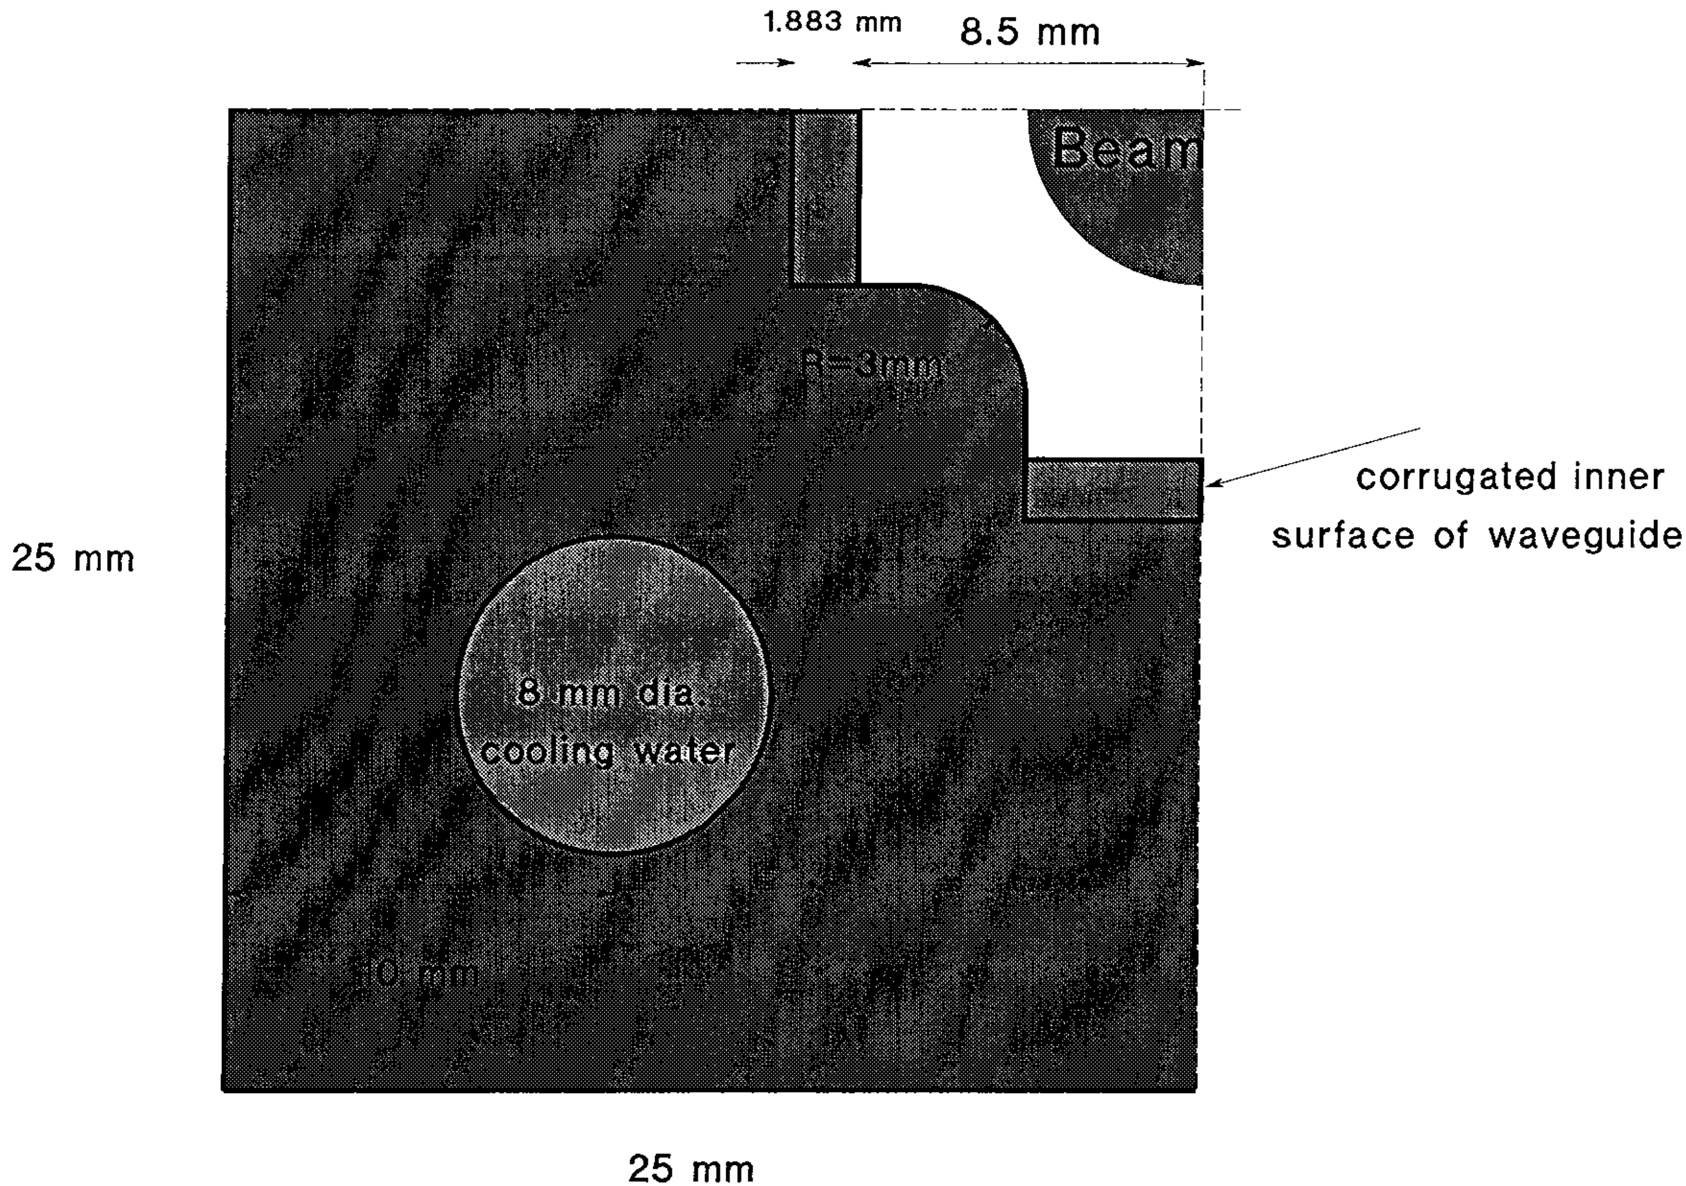

Section of four-waveguide CLIC Transfer Structure

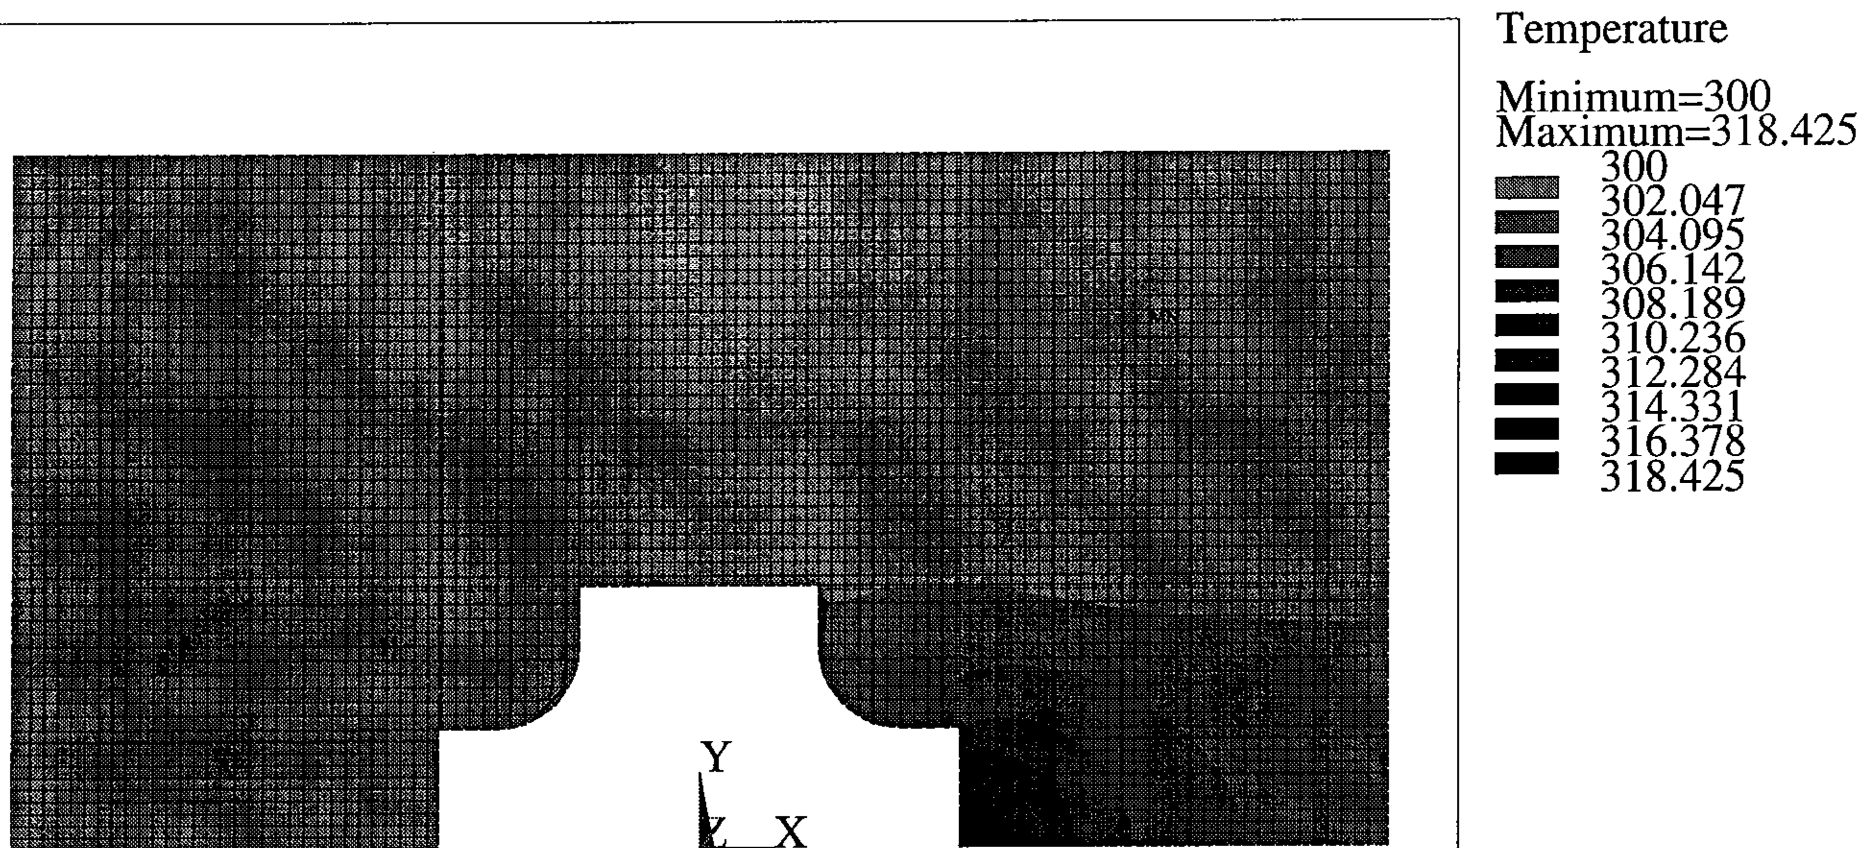

Figure 2. Temperature distribution in CLIC drive waveguide. Case 1.

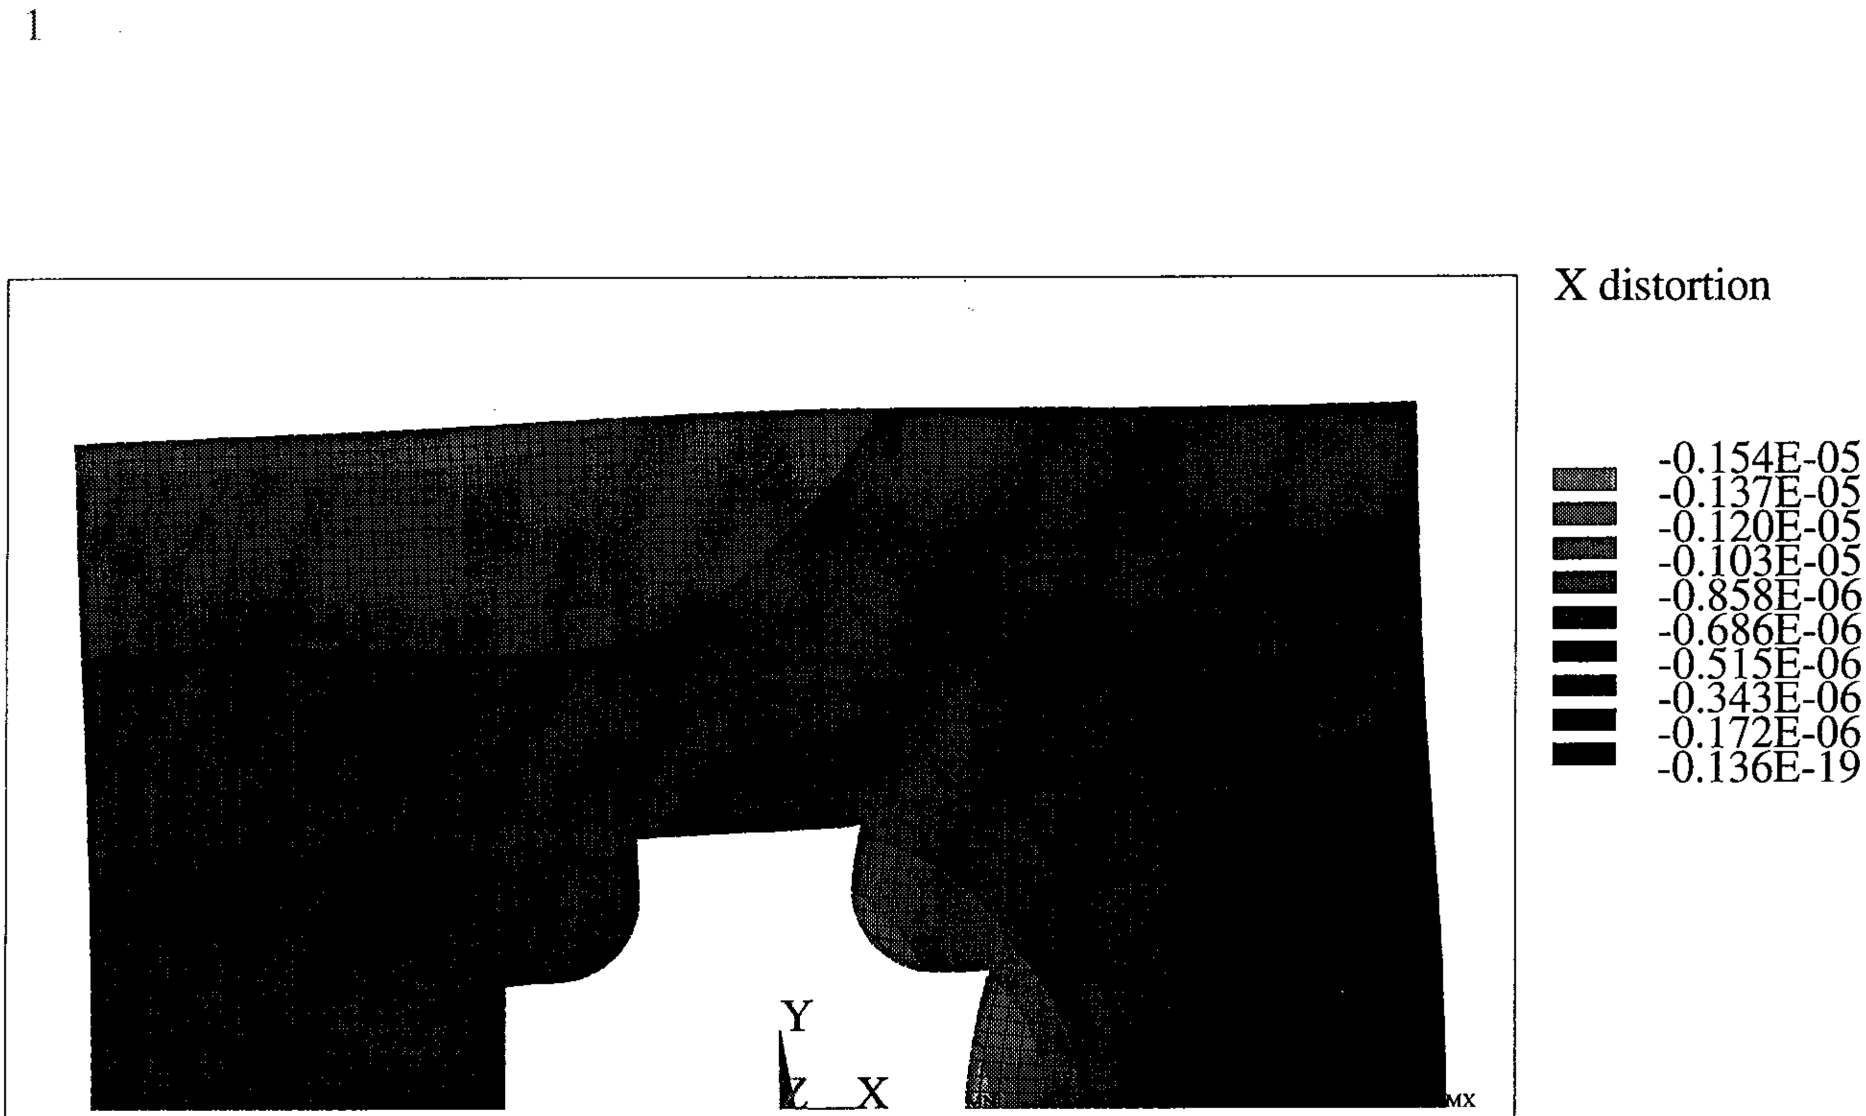

Figure 3. Horizontal (x) distortion in CLIC drive waveguide. Case 1.

1

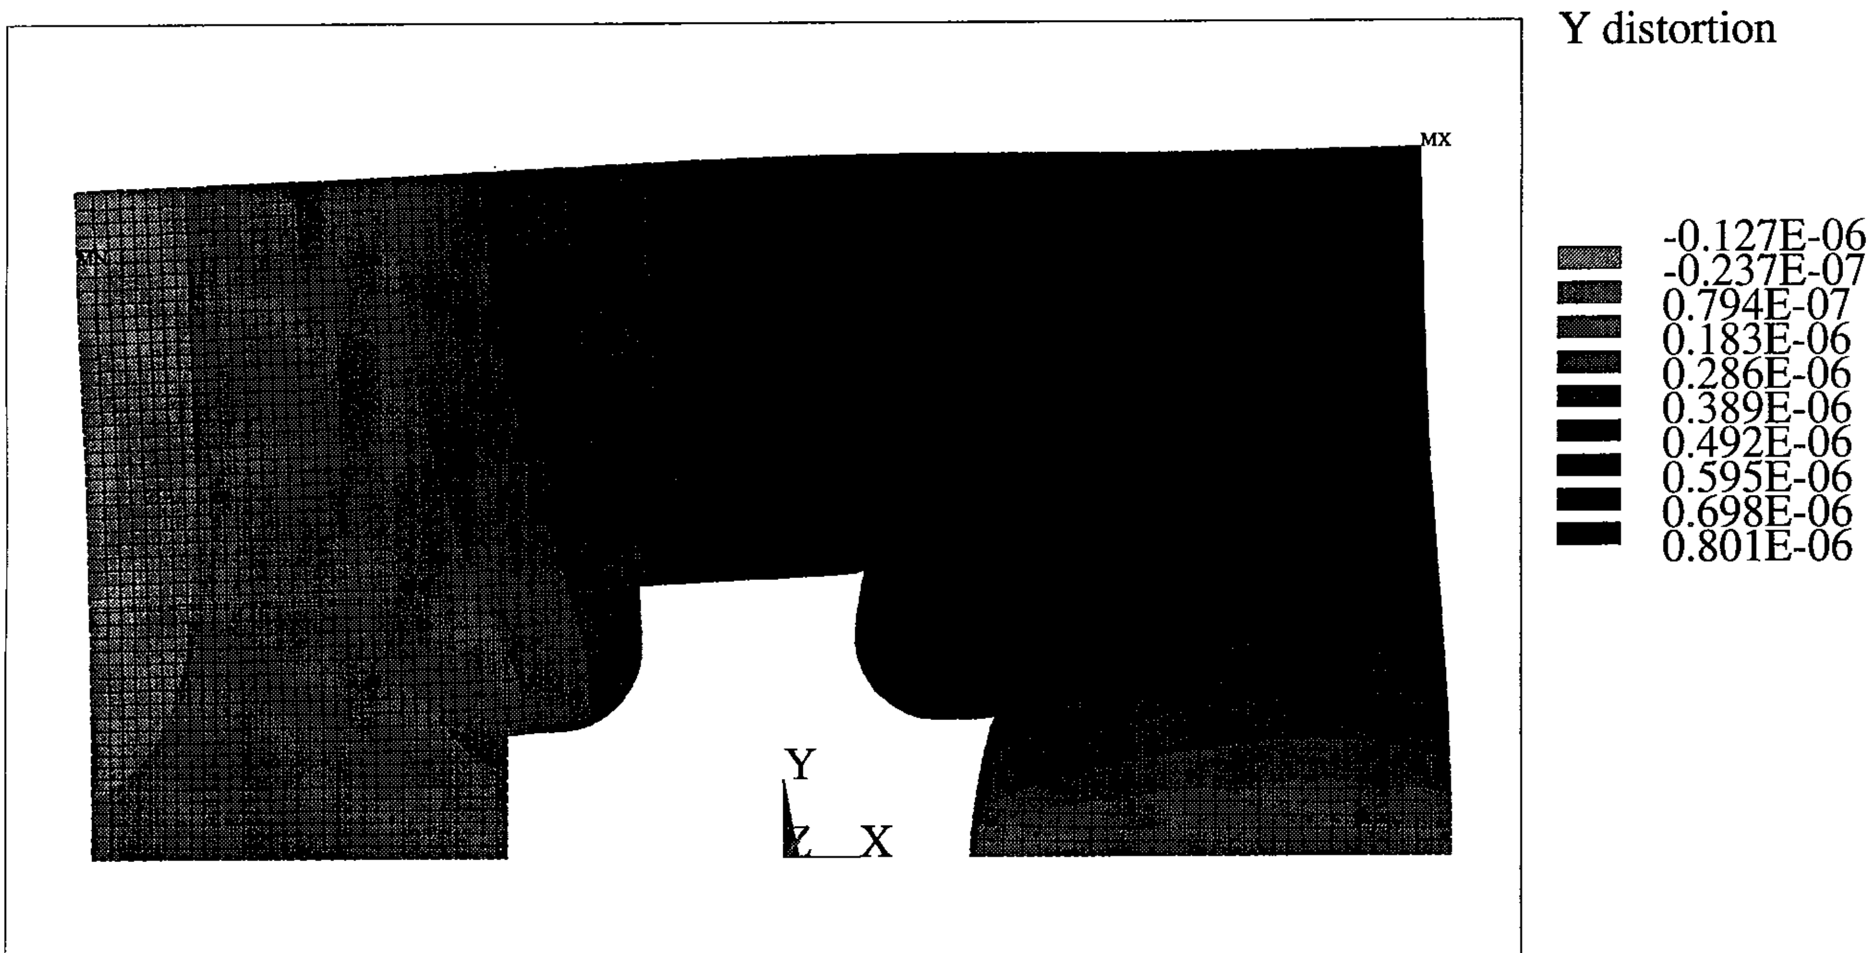

Figure 4. Vertical (y) distortion in CLIC drive waveguide. Case 1.

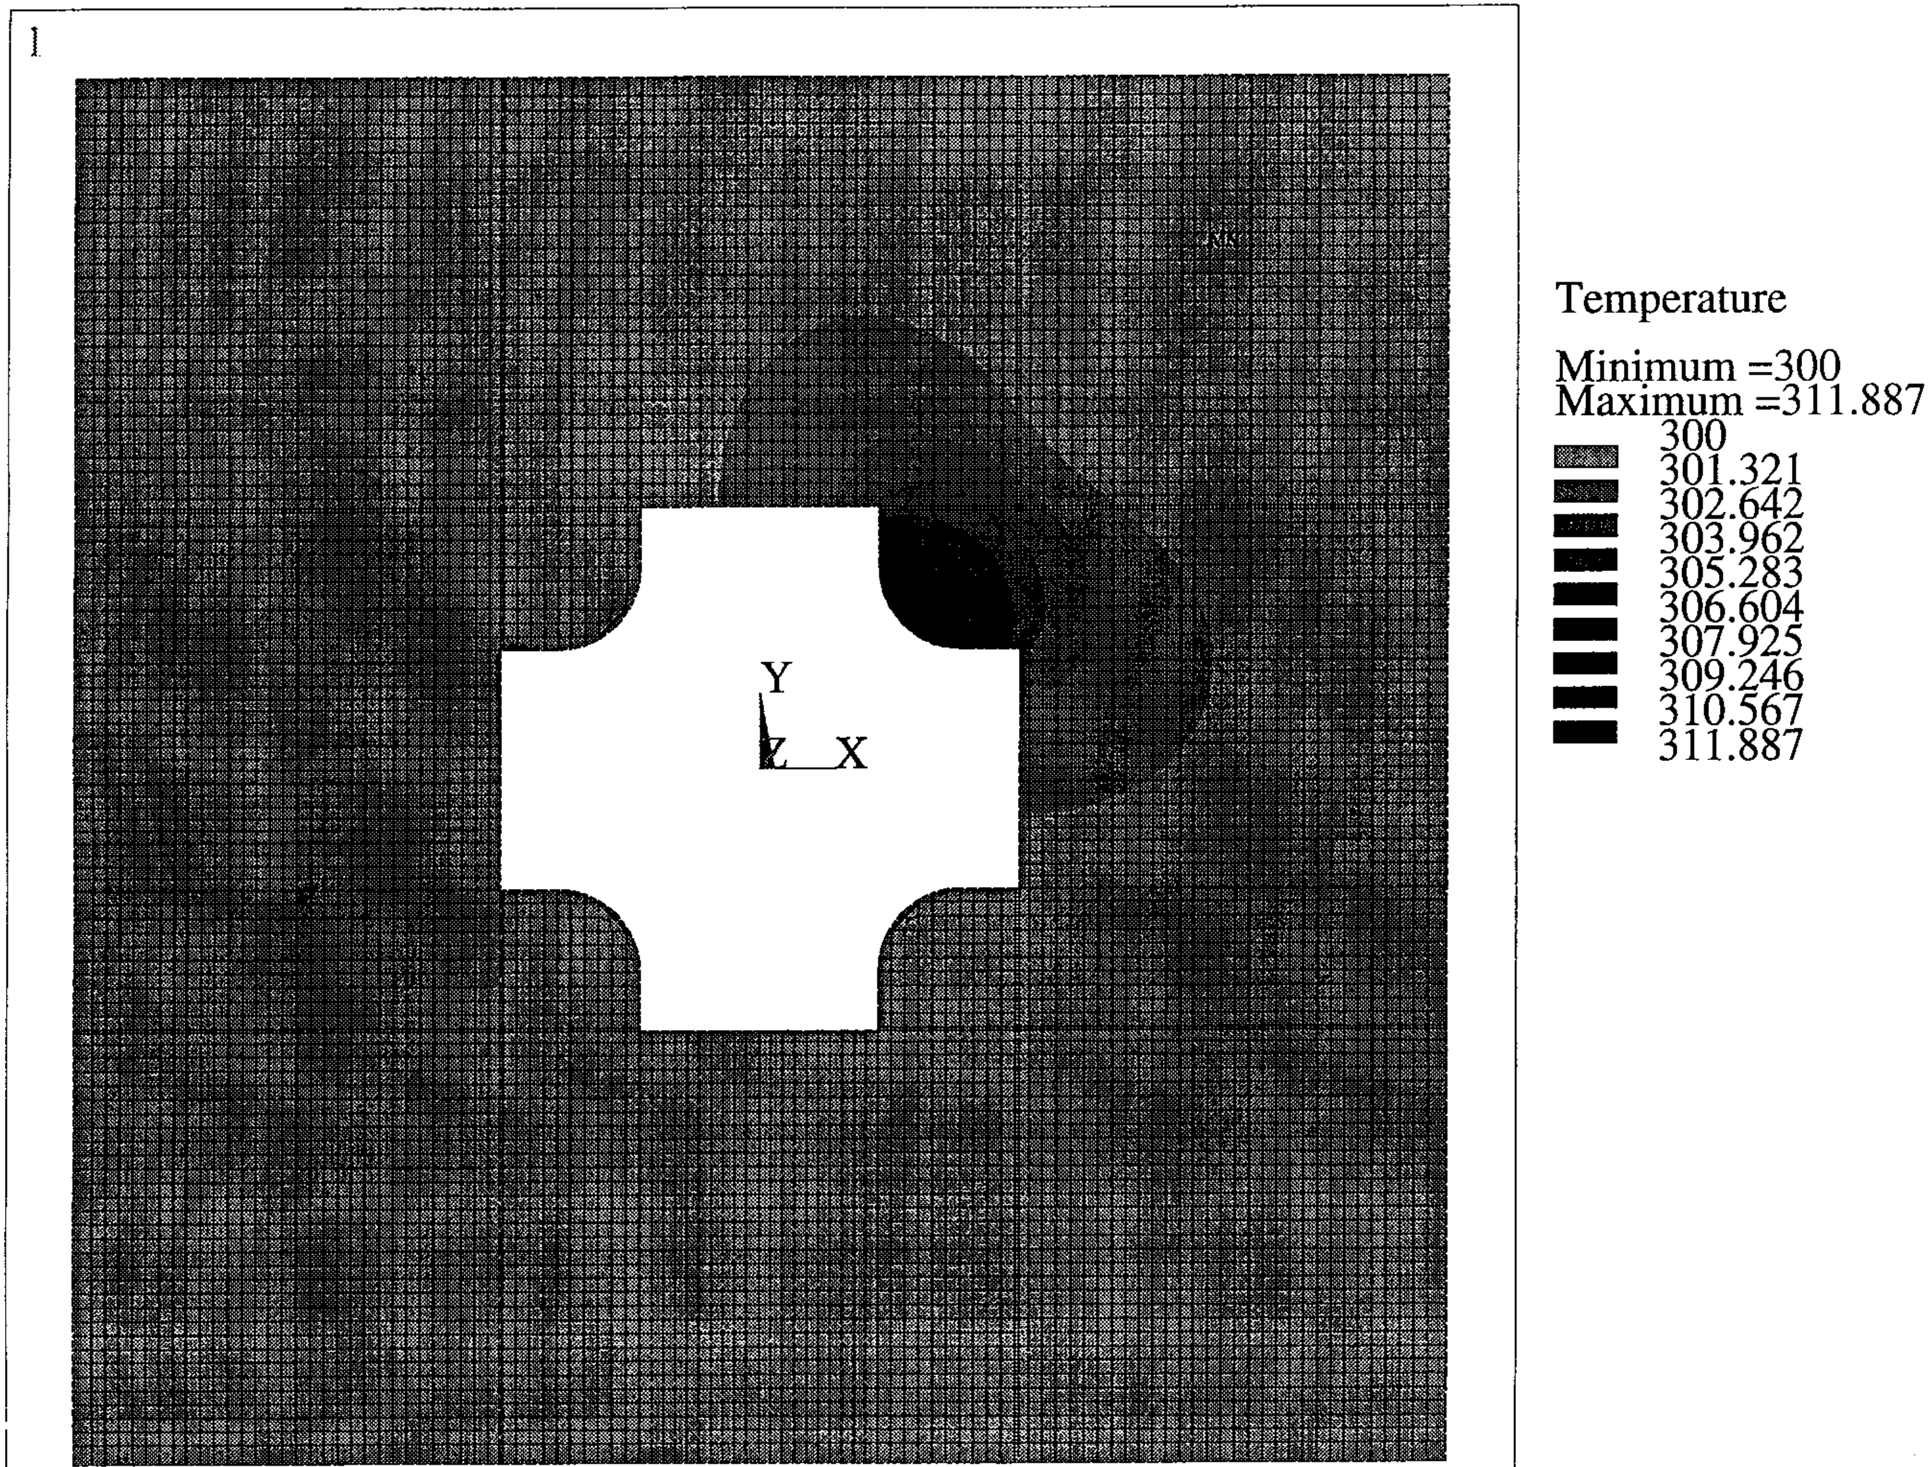

Figure 5. Temperature distribution in CLIC drive waveguide. Case 2.

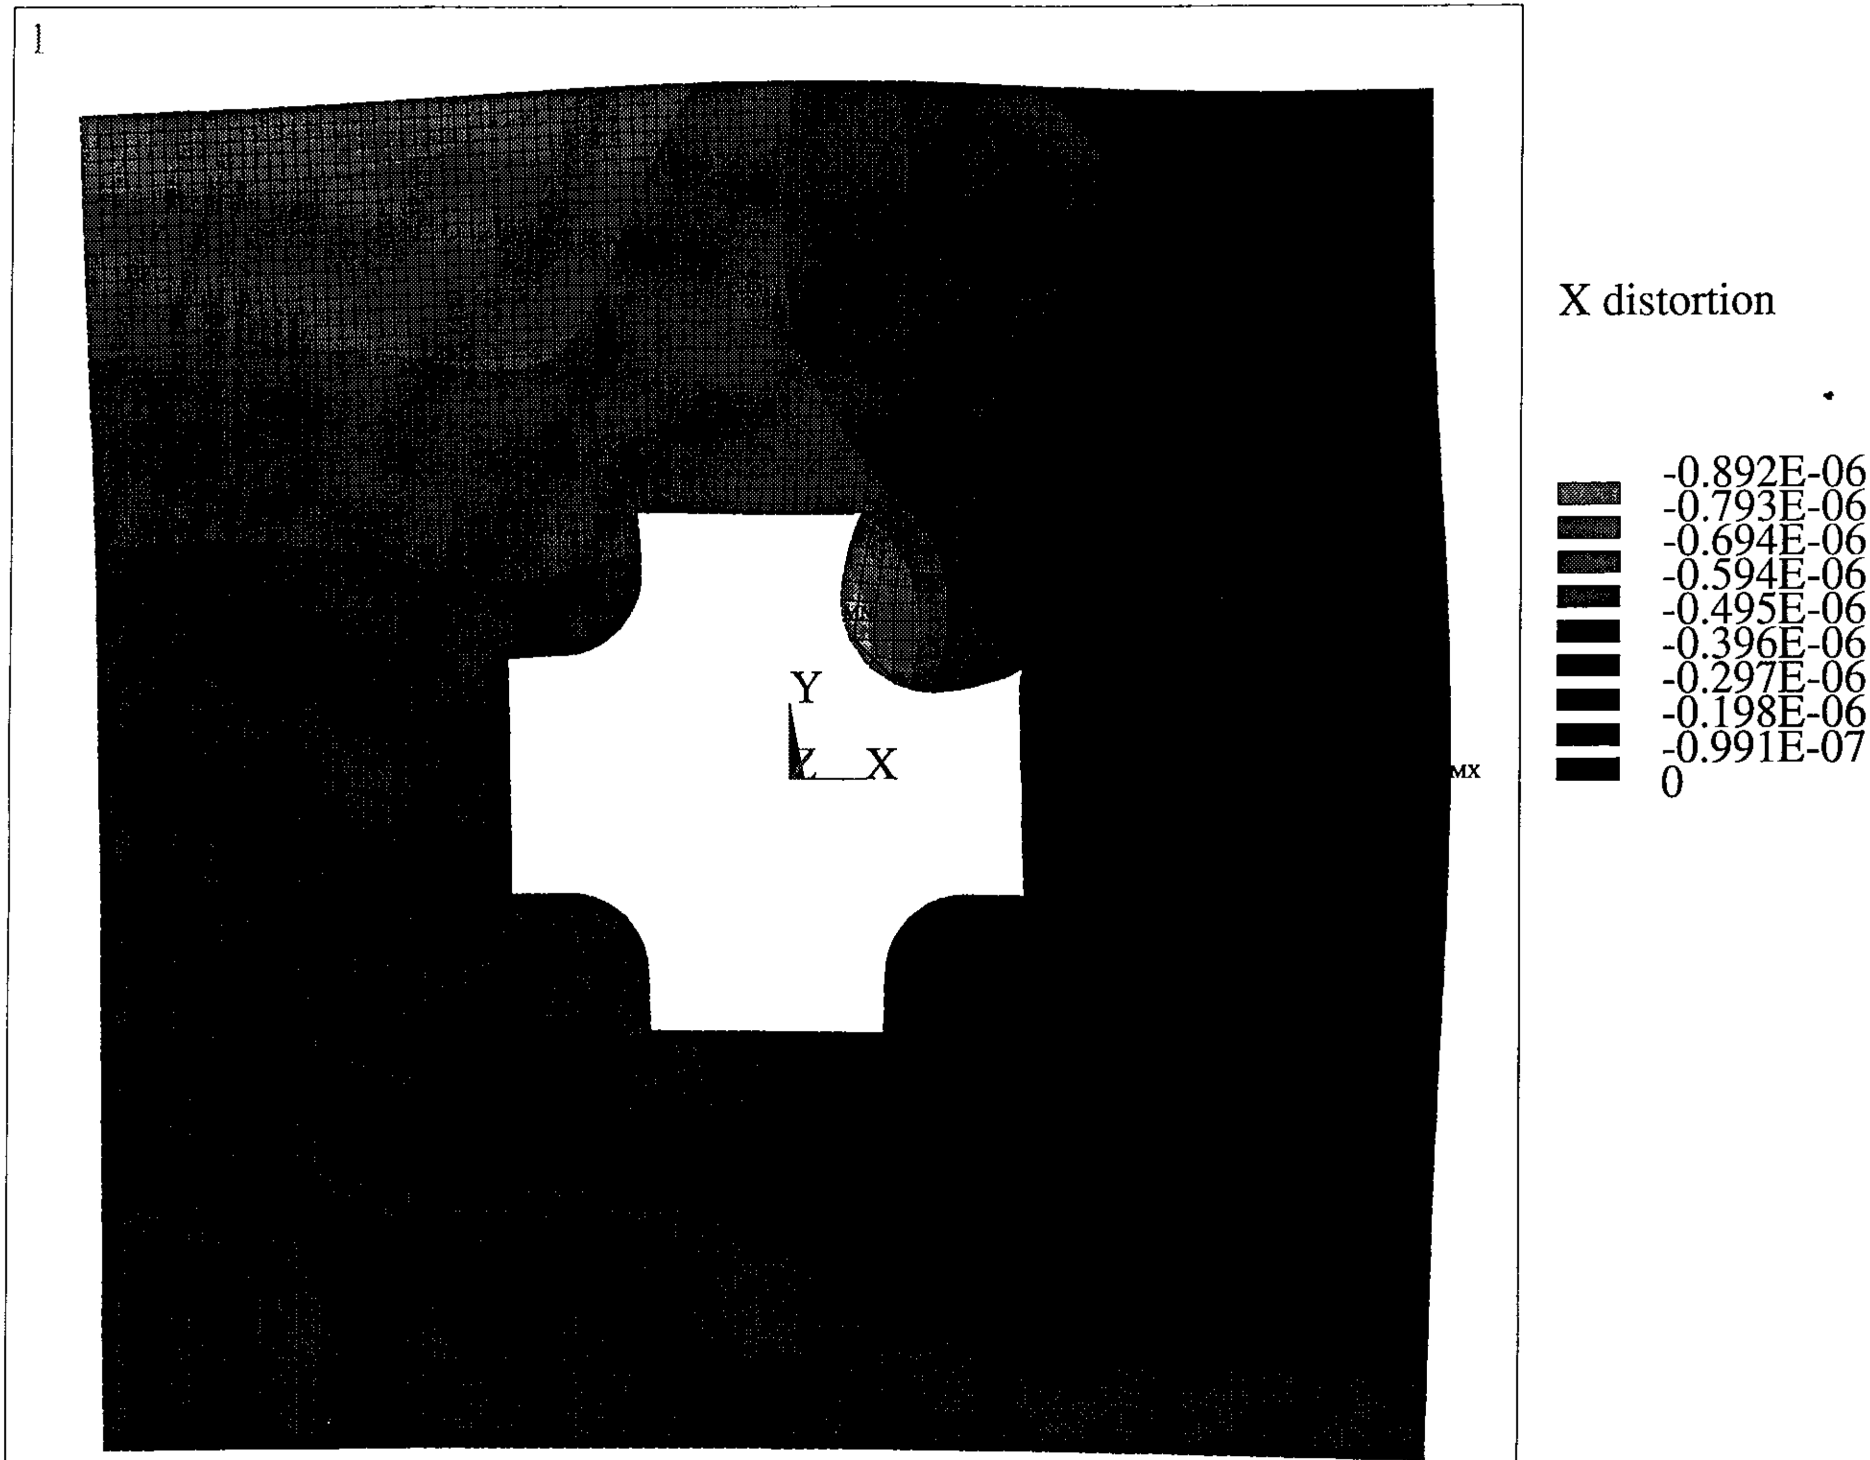

Figure 6. Horizontal (x) distortion in CLIC drive waveguide. Case 2.