

SURFACE FIELD OPTIMIZATION OF ACCELERATING STRUCTURES FOR CLIC USING ACE3P ON REMOTE COMPUTING FACILITY*

Kyrre Ness Sjobak[†] and Erik Adli, Department of Physics, University of Oslo, Norway
 Alexej Grudiev, CERN, Geneva, Switzerland

Abstract

This paper presents a computer program for searching for the optimum shape of an accelerating structure cell by scanning a multidimensional geometry parameter space. For each geometry, RF parameters and peak surface fields are calculated using ACE3P on a remote high-performance computational system. Parameter point selection, mesh generation, result storage and post-analysis are handled by a GUI program running on the user’s workstation. This paper describes the program, AcdOptiGui. AcdOptiGui also includes some capability for automatically selecting scan points based on results from earlier simulations, which enables rapid optimization of a given parameterized geometry. The software has previously been used as a part of the design process for accelerating structures for a 500 GeV CLIC.

INTRODUCTION

The Compact Linear Collider (CLIC) main beam accelerating structures are tapered traveling wave structures, which are composed of a number of cells. In order to find the optimum main beam accelerating structure for CLIC, it is necessary to scan over a large number of accelerating structure geometries, and for each of them estimate its power requirements and effects on the beam. The most practical way of estimating the accelerating mode parameters is to use an analytic model for the accelerating mode power flow along the structure [1].

This analytic model depends on a continuous description of the structure’s Q , R/Q , and group velocity as a function of the position along the structure. Further, the peak surface fields along the structure are also needed in order to evaluate the breakdown constraints. To get these functions, the values for a single accelerating cell with the correct local geometry for the beginning, middle, and end of the structure are pre-calculated. These values are then interpolated along the structure. When scanning over the set of accelerating structure geometries, these interpolation points are themselves found by interpolating from a table of carefully designed cells with different local geometries.

Constructing such a table requires high-gradient optimization of approximately 50 cells, and for each such cell five internal parameters are varied in order to minimize the surface fields. The ACE3P frequency-domain solver Omega3P [2] is ideally suited for this, due to being able to solve fairly complex geometries quickly, and also having the possibility to solve multiple geometries in parallel as it

is running on large batch processing compute clusters. Unfortunately, manually setting up, submitting, and analyzing the large number of small runs required for each cell is a very slow and error-prone process, and thus not practical. A computer program called AcdOpti was therefore written in order to automate this process. This article presents AcdOpti as applied to cell optimization for CLIC.

GEOMETRY OF AN ACCELERATING CELL

Each of the single cells to be optimized has the same waveguide-damped topology as found in the CLIC.G [3] and CLIC_502 [4] structures, and the parameters for this topology are shown in Figure 1. Typically the parameters \mathbf{a} , \mathbf{d} and \mathbf{L} are varied in the overall structure optimization. For each selection of \mathbf{a} , \mathbf{d} and \mathbf{L} five “internal” parameters are varied in order to minimize the maximum surface fields. Due to the field geometry, the surface magnetic field (which determines the pulsed heating ΔT) is dominated by the outer wall geometry, which is described by the parameters \mathbf{eow} and \mathbf{c} . On the other hand, the peak surface electric field and the modified Poynting vector S_c [5] are dominated by the iris geometry, described by the parameters \mathbf{e} and \mathbf{s} . Finally for each point in the outer wall and iris optimization, the cell radius \mathbf{b} is tuned to the frequency 11.9942 GHz. The remaining parameters \mathbf{adw} , \mathbf{idw} , \mathbf{rdw} and \mathbf{rr} are kept fixed in this optimization, while the damping waveguide length \mathbf{ldw} in these simulations is set to be long enough for the main mode field to decay to approximately zero at its end.




Figure 1: Geometry parameters for the single cells, showing iris parameters (left) and outer wall/damping waveguide parameters (right).

This decoupling of the iris- and outer wall properties makes the optimization much easier, as it effectively becomes two independent two-variable optimizations, and the cell radius \mathbf{b} is always determined using the tuning algorithms described below.

* This work supported by the Research Council of Norway.

[†] k.n.sjobak@fys.uio.no

ACDOPTI SOFTWARE

AcOpti is implemented as a Python library which manages simulation input- and output data and communication with the remote computing facility. Access to the data and actions is provided through the modular object-oriented API. The API also makes available multiple algorithms for scanning over geometry parameters, exporting data, and performing common analysis tasks. There is also a GUI, AcOptiGui, which is built on top of this API. The software is available at <https://github.com/kyrsjo/AcdOpti> under the GNU GPL version 3.

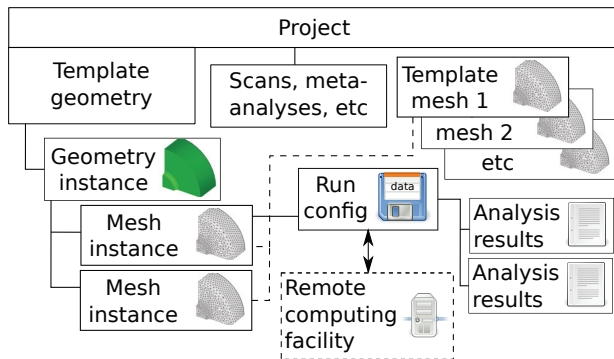


Figure 2: Data structure in an AcOpti project.

Data Structure

The AcOpti software works using a parametric geometry description, such as the one described in the section above. This is described to the program using a CUBIT [6] Aprepo script, where some variables have been marked out to be substituted by AcOpti. This set of variables and their default values plus the template Aprepo scripts forms the “geometry template” shown in Figure 2. By specifying values for the variables and feeding the script to CUBIT through its Python interface, a solid model can be created. The set of variable values, the CUBIT interface, and managing of the solid model data is handled by a “geometry instance”, and a typical optimization project usually contains a few hundred of these.

Meshing works with a similar mechanism: A project contains one or more “mesh templates”, which is a Aprepo script and a list of variables defining such things as mesh density. Such a template can then be used to create “mesh instances”, of which one or more can be attached to any geometry instance. The mesh instance then handles communication with CUBIT in order to create the actual mesh, and then calls a local version of ACE3P’s mesh conversion and post-analysis tool acdtool in order to check the validity of the mesh and convert it to a format which can be used by the solver Omega3P.

Further, for each mesh there may be one or more parameter sets for the simulation software, specifying such things as the phase advance, finite element order, boundary condition IDs, and number of CPUs to request. This is handled by the “run configuration”, which generates the input files for Omega3P and communicates with the remote comput-

ing facility using SSH in order to up-/download data and manage the run. The simulation output data can then be analyzed by one or more analysis modules.

AUTOMATIC TUNING

When searching for the optimum cavity shape by varying one or more parameters, such as when minimizing the surface magnetic field by variation of \mathbf{eow} and \mathbf{c} , it is necessary to compare with trial geometries at the same frequency. Thus every $(\mathbf{eow}, \mathbf{c})$ -point must be tuned, and this is usually done by varying the \mathbf{b} -parameter until the correct frequency has been achieved. AcOpti currently have two strategies for selecting the correct value for \mathbf{b} using as few calculations as possible, and these are described below.

The error estimates are taken as $\pm 1.96 \sqrt{\frac{\text{variance}}{\text{num. points}}}$.

Linear 1D Fit Tuning

This method works by calculating the frequency of two (or more) cells with different values of \mathbf{b} , preferably close to the correct value. The frequency is then fitted as a linear function of \mathbf{b} , and this fit is then used to predict the correct \mathbf{b} for the target frequency. For a sample of 19 such fits from two points with a spread in \mathbf{b} of 0.05 or 0.35 mm, the average frequency error was 0.009 ± 0.02 MHz, and the average error on the radius 4 ± 8 nm. The radius error is estimated by refitting including the first predicted radius, and comparing the old and the updated estimate.

2D Surface Fit Tuning

This method works by selecting calculations that fall within a narrow frequency range (typically ± 0.2 MHz), and then fitting the radius \mathbf{b} as a function of two scan parameters (such as \mathbf{eow} and \mathbf{c}). The resulting plane or quadratic surface is then used to predict the correctly tuned radius \mathbf{b} when creating new scan points. The accuracy of the predictions made by this method depends on how well tuned the initial points are, and also whether the predictions are extrapolating or not. For a set of 12 points tested, the average frequency error was -0.1 ± 0.3 MHz.

GRAPHICAL USER INTERFACE

AcOpti’s GUI, AcOptiGui, is built on top of the AcOpti library API using GTK and Matplotlib, and allows the user to easily manage geometries, runs, and results. It also presents an interface to such things as scans and meta-analyses.

The interface is split vertically in two parts: The project explorer on the left and panels for interacting with different parts of the program on the right, as seen in Figure 3. The project explorer shows the data hierarchy in a similar manner as shown in Figure 2, and also indicates the status of the different modules with color codes. Below the project explorer there are buttons to batch process a large number of geometries/runs or analyses, and to efficiently interact with the explorer. Different panels are shown when selecting different modules from the project explorer, or when

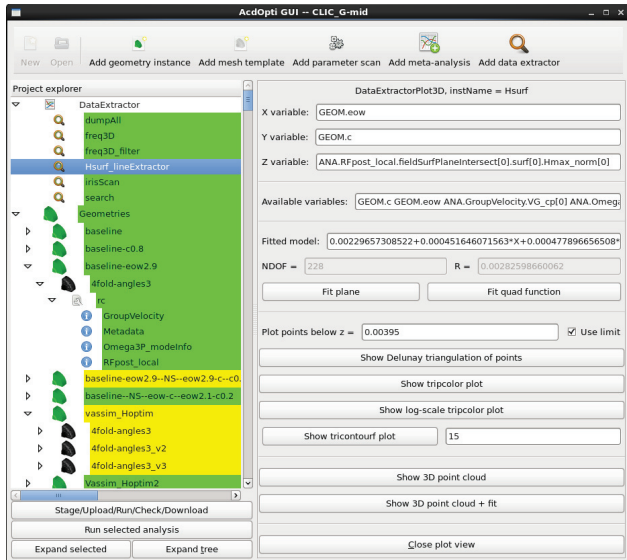


Figure 3: Screenshot of the AcdOptiGui interface, showing parts of the project explorer and a 3D-plotting meta-analysis based on a data-extractor.

doing certain other actions. These panels allow interaction with the modules, and the architecture of AcdOptiGui allows easy implementation of more such panels as they are self-contained classes.

EXAMPLE USAGE

As an example, a peak surface magnetic field optimization of a CLIC.G [3] middle cell was performed. This cell has $L = 8.33159$ mm (120° phase advance per cell at 11.9942 GHz) and iris parameters $a = 2.75$ mm, $d = 1.335$ mm. The outer wall parameters eow and c were first scanned for approximately 20 points in the region of interest, and at each point the frequency at two different radii were calculated. A linear 1D fit was then used at each of these points to predict the correct value for b . These tuned geometries were then solved in order to confirm the tuning, and also used to locate the approximate location of the minimum by plotting the peak surface field as a function of eow and c .

Starting from these tuned geometries, a quadratic surface describing b as a function of eow and c was fitted, and this was used to predict the value of b for subsequent trial geometries. These trials were then placed in the vicinity of the expected minimum until this was clearly defined. The magnetic field as a function of eow and c is shown in Figure 4, and the same plot is easily viewed through the 3D plotting meta-analysis which is shown in Figure 3.

The process is fairly rapid – one trial geometry takes approximately 1-2 minutes to solve on a single node on `hopper.nersc.gov`, and multiple solvers are run in parallel. In addition to this comes mesh generation and file transfer time, which takes on the order of 2-4 minutes per trial geometry depending on the speed of the workstation, internet connection, and geometry complexity. The user will only spend a fraction of this time operating the pro-




Figure 4: Peak surface magnetic field / mean accelerating gradient [mA/V] as function of outer wall geometry. Crosses indicates trial points, where points included in the plot are red, and points that are excluded due to having a normalized surface magnetic field > 3.95 mA/V are blue.

gram – it will typically request intervention (signaled by a beep) for a few minutes once every half hour.

CONCLUSIONS

The AcdOpti program with its user interface AcdOptiGui has been developed in order to support high-gradient optimization of accelerating cells for CLIC, and has already been used to re-optimize the CLIC.502 accelerating structure [4]. The software itself is however built to be a general-purpose tool, and allows optimization of any geometry parameter with respect to any simulation result. The user interface AcdOptiGui provides rapid presentation of results and access to both the data and tools such as automatic tuning. This allows accelerating cells to be optimized quite quickly for a relatively large number of parameters, with minimal amount of manual intervention necessary.

ACKNOWLEDGMENT

This work was enabled by support from ACD at SLAC with K.Ko, A. Candel, Z. Le, K.H. Lee and C.K. Ng. Computing time with ACE3P provided by US DOE at NERSC. Thanks to Walter Wuensch for discussions and support.

REFERENCES

- [1] A. Lunin, V. Yakovlev, and A. Grudiev, “Analytical solutions for transient and steady state beam loading in arbitrary traveling wave accelerating structures,” PRST-AB 14, 2011.
- [2] K.Ko et al, “Advances in Parallel Electromagnetic Codes for Accelerator Science and Development”, LINAC’10, Tskuba
- [3] A.Grudiev and W.Wuensch, “Design of the CLIC main linac accelerating structure for CLIC conceptual design report”, LINAC’10, Tskuba
- [4] K. Sjobak et al, “Design of an accelerating structure for a 500 GeV CLIC using ACE3P”, IPAC’12, New Orleans
- [5] A. Grudiev, S. Calatroni, and W. Wuensch, “New local field quantity describing the high gradient limit of accelerating structures”, PRST-AB 12, 2009.
- [6] CUBIT Geometry and mesh generation toolkit, <http://cubit.sandia.gov/>