COMPUTATIONAL TOOLS FOR RF STRUCTURE DESIGN

Erk Jensen

Abstract

The Finite Differences Method and the Finite Element Method are the two principally employed numerical methods in modern RF field simulation programs. The basic ideas behind these methods are explained, with regard to available simulation programs. We then go through a list of characteristic parameters of RF structures, explaining how they can be calculated using these tools. With the help of these parameters, we introduce the frequency-domain and the time-domain calculations, leading to impedances and wake-fields, respectively. Subsequently, we present some readily available computer programs, which are in use for RF structure design, stressing their distinctive features and limitations. One final example benchmarks the precision of different codes for calculating the eigenfrequency and Q of a simple cavity resonator.

COMPUTATIONAL TOOLS FOR RF STRUCTURE DESIGN

E. JENSEN

CERN

AB-RF

CH-1211 Geneva 23, Switzerland

E-mail: Erk.Jensen@cern.ch

The Finite Differences Method and the Finite Element Method are the two principally employed numerical methods in modern RF field simulation programs. The basic ideas behind these methods are explained, with regard to available simulation programs. We then go through a list of characteristic parameters of RF structures, explaining how they can be calculated using these tools. With the help of these parameters, we introduce the frequency-domain and the time-domain calculations, leading to impedances and wake-fields, respectively. Subsequently, we present some readily available computer programs, which are in use for RF structure design, stressing their distinctive features and limitations. One final example benchmarks the precision of different codes for calculating the eigenfrequency and Q of a simple cavity resonator.

1 Introduction

Computational tools for RF system design and readily available computer hardware have made remarkable progress over the past decades. This allows us today to design complex RF structures almost entirely on the computer, minimizing the lengthy cut-and-try optimization methods as well as expensive and costly model building for cold measurements. In particular these computer programs also allow us to simulate the interaction with the particle beam, something almost entirely inaccessible to mechanical models. However, the most advanced computational tools also require deep insight into the underlying principles of both RF theory and the numerical tool itself.

The aim of this article is to summarize how the relevant quantities in RF structure design are accessible to numerical computation, to give an idea of how these numerical tools work in principle, and finally to present some of the actual programs available and compare them.

2 RF field calculation methods

The computer codes used to simulate electromagnetic fields numerically start from the idea of discretizing space. For example, one might think of considering the fields only at some discrete values of the three Cartesian co-ordinates x , y

and z . Between these mesh points, some type of interpolation is used to represent the complete solution.

Even though other methods exist (Transmission-Line Method, field matching methods, ...), the two main methods for the simulation of RF fields are the Finite Differences Method (FDM), where the differential equations are directly replaced by difference equations between the field quantities at discrete mesh points, and the Finite Element Method (FEM), where the overall solution is represented as a superposition of simple functions, which are zero everywhere except on localized, finite elements of space.

2.1 Finite Differences Method

The Finite Differences Method (FDM) is based on the idea of transforming the differential equations to a set of difference equations.

The task is to find a solution to Maxwell's equations inside the cavity volume subject to given boundary conditions (in time-domain also initial conditions). Consider a mesh-grid defined inside the volume. Maxwell's equations in differential form are a set of linear differential equations of first order (or – in the form of the Helmholtz equation – of second order). Replacing the differential operators by difference operators results in a set of linear algebraic equations, which relate a field quantity at one mesh point to the field quantities at the directly neighbouring points.

For illustration, consider the (second-order) Laplace equation $\Delta\Phi=0$ in two dimensions and Cartesian co-ordinates. For the potential at some central mesh point, and assuming identical mesh size in both directions, the central difference-quotient would result in the equation

$$\Phi_{m+1,n} + \Phi_{m-1,n} + \Phi_{m,n+1} + \Phi_{m,n-1} - 4\Phi_{m,n} = 0. \quad (1)$$

An equation of this type exists for the potential at every mesh point; at boundary points the boundary conditions (in this simple case a given potential) can be included on the right-hand side. If we further assume for simplicity an equidistant, uniform mesh throughout the entire space, the resulting linear equation would have the form

$$\begin{pmatrix} \ddots & & & & & & & & & & \\ & \ddots & & & & & & & & & \\ & & \ddots & & & & & & & & \\ & & & \ddots & & & & & & & \\ & & & & \ddots & & & & & & \\ & & & & & \ddots & & & & & \\ & & & & & & \ddots & & & & \\ & & & & & & & \ddots & & & \\ & & & & & & & & \ddots & & \\ & & & & & & & & & \ddots & \\ & & & & & & & & & & \ddots \end{pmatrix} \begin{pmatrix} \vdots \\ \Phi_{m,n-1} \\ \vdots \\ \vdots \\ \vdots \\ \Phi_{m-1,n} \\ \Phi_{m,n} \\ \Phi_{m+1,n} \\ \vdots \\ \vdots \\ \vdots \\ \Phi_{m,n+1} \end{pmatrix} = \begin{pmatrix} b \\ o \\ u \\ n \\ d \\ a \\ r \\ i \\ e \\ s \end{pmatrix}, \quad (2)$$

where the ellipsis indicates the non-zero matrix elements. This is a large (number of mesh-points squared), banded, sparsely occupied matrix. Methods of inverting large sparse matrices exist (Cholesky, SOR, ADI, conjugate gradient ...).

In the case of Maxwell's equations in three (or – for time-dependent equations – four) dimensions, the matrix becomes much larger and gets more bands, and for variable mesh sizes and mesh elements, which are not parallel to the co-ordinate axes, the matrix elements become less trivial. Particular care must be taken to properly account for the boundary conditions, in particular if the physical boundary does not coincide with the mesh-grid.

The FIT algorithm, used for example in MAFIA and in GdfidL, is based on a double grid (see Figure 1), where the mesh points of one grid are centred in the cubical cells formed by the other. It is a feature of Maxwell's equations that the spatial derivatives of the electric field are related to the magnetic field and vice versa. Thus when taking $\vec{E}$, $\vec{D}$ and $\vec{J}$ on the mesh-points of one grid, but $\vec{B}$ and $\vec{H}$ on the edges of the dual grid, one can formulate Maxwell's equations on the double grid to obtain the so-called "Maxwell's Grid-Equations (MGE)." It is convenient to use the integral form of Maxwell's equations here, but to extend the integrals only over the finite areas or volumes of the grid cells (FIT = Finite Integration Theory). This formulation is less prone to discretization errors than a single-grid formulation. A further advantage is that the vector identities $\nabla \times \nabla \varphi \equiv 0$ and $\nabla \cdot \nabla \times \vec{A} \equiv 0$ become properties of the "topological" matrices representing the curl, grad and div operators.

Figure 1: The dual grid used for MAFIA's FIT algorithm.

The method is applicable both in frequency and in time domain. When applied in time domain, one also uses the term "FDTD algorithm" (Finite Differences Time Domain).

2.2 Projection Methods – Finite Element Method

The task is again to find a solution to Maxwell's equations inside the cavity volume subject to given boundary conditions (in time-domain also initial conditions). Let us abbreviate the differential equation as

$$D(\varphi) = 0, \quad (3)$$

where φ represents the (unknown) solution. For illustration, we have set the right-hand side to zero, corresponding to the eigenvalue problem; in a more general case, the sources (currents and charges) appear on the right-hand side. In the case of Maxwell's equation, φ may stand for the vector $\vec{E}$ in the entire volume and $D(\varphi) = 0$ for

$$\nabla \times \left(\frac{1}{\mu} \nabla \times \vec{E} \right) - \omega^2 \epsilon \vec{E} = 0. \quad (4)$$

The projection methods start by assuming that the solution can be represented approximately in the form

$$\varphi = \sum_{n=1}^N a_n \varphi_n, \quad (5)$$

where a is a vector with N (many) dimensions, and the φ_n are known basis functions (trial functions) which satisfy the boundary conditions. This assumed φ may not solve the original equation (3) exactly, so we provide a residual r on the right-hand side,

$$\sum_{n=1}^N a_n D(\varphi_n) = r, \quad (6)$$

and the task will be to minimize r . With the scalar product in function space, denoted $\langle \varphi, \psi \rangle$ (normally an integral over the entire volume), one can now scalar multiply Eq. (6) with N test functions (or weight functions) ψ_m , and solve

$$\langle \psi_m, r \rangle = \sum_{n=1}^N a_n \langle \psi_m, D(\varphi_n) \rangle = 0. \quad (7)$$

The name “projection methods” refers to this scalar multiplication. These projections of the residual r on the test functions will determine the coefficients a_n such that the residual (error) is minimized, assuring that r has zero length in the sub-space spanned by the ψ_m . Since both the basis functions and the test functions are known, the scalar product in Eq. (7) can be evaluated, and we again get a $N \times N$ matrix equation to determine a .

Different choices of weight functions are possible. If the weight functions are chosen to be identical to the basis function, this is called the Galerkin method. If in addition the basis functions are chosen orthogonal, i.e. $\langle \varphi_n, \varphi_m \rangle \propto \delta_{nm}$ (the Kronecker delta), we have spectral methods. Examples of orthogonal functions are the eigenfunctions of the system, trigonometric functions (cf. Fourier series), Bessel functions, Legendre polynomials, Chebyshev polynomials, and others.

The Finite Element Method (FEM) can be classified also as a Galerkin method; each basis function in this case is chosen to be zero everywhere in the volume except for a finite region (the finite element), where the basis functions

are very simple (typically piecewise linear or quadratic polynomials). This makes the evaluation of the integral for the scalar product in Eq. (7) simple. Since the basis functions are non-zero only inside one finite element, the matrix will again be sparse and banded as in Eq. (2), so basically the same methods apply for the inversion of these matrices as for the FDM.

2.3 *Different meshes*

In the FIT example above, we have used a Cartesian global mesh, which simplifies the formulation of the difference equation significantly, since the partial derivatives are taken along mesh lines. However, the contours of arbitrarily shaped RF structures do not generally coincide with mesh lines, and consequently the meshed geometry may be different from the physical geometry to be simulated. This leads to a systematic error, which can be controlled and estimated, and which can be reduced by iterative mesh refinement. However, there remain preferred directions in space, and the meshing (and possibly also the simulation result) will depend on the orientation of the geometry in the mesh, i.e. the mesh can create some intrinsic, numerical anisotropy.

If it is necessary to have a local mesh refinement in a particular area of interest inside the structure, it is typical for the global Cartesian mesh that every added mesh point will add mesh planes which extend through the entire structure, thus extending this refinement also into areas where it is not actually needed.

If the constraint to have mesh lines parallel to the co-ordinate axes is dropped, a completely isotropic mesh may result. When connecting those mesh-points, one can always create a dense packing of tetrahedra (in two dimensions, triangles). The finite difference formulation becomes complicated in this case, but the Finite Element Method often uses tetrahedral meshes.

A local mesh refinement is straightforward with tetrahedral meshes; when we add a mesh point inside an existing tetrahedron, it naturally splits it onto 4 new tetrahedra.

3 **Computing the characteristic parameters of an RF structure**

In the following, we will go through a number of quantities one typically needs a numerical tool to compute. For these quantities, we will mention their mutual relations, explain how they are computed and, if necessary, stress which particular feature of the simulation tool is needed to perform the computation.

3.1 Eigenfrequency

When designing a single-gap RF cavity, the first quantity one typically needs is its resonance frequency f_0 (eigenfrequency), which should correspond to the operation frequency.

In a closed cavity with given boundary conditions but without any sources, equation (4) above (or equivalently the Helmholtz equation $(\Delta\phi + k^2\phi) = 0$) results in an eigenvalue problem. Each eigensolution corresponds to a resonant mode, with its eigenvalue corresponding to the eigenfrequency, the eigenvector to the field distribution of this mode. The field simulation programs discussed below are able to solve for eigenfrequencies; even the time domain programs (MAFIA, GdfidL and Microwave Studio) have this feature built in.

For travelling-wave accelerating structures, the frequency relevant for efficient acceleration is the synchronous frequency, i.e. the frequency at which the phase velocity of the accelerating mode is equal to the velocity of the beam. For a $v=c$ beam, this frequency is equal to the frequency where the phase

advance $\Delta\phi$ in a single cell of length l is equal to $\frac{\Delta\phi}{2\pi} = \frac{l}{\lambda}$, where λ is the free-space wavelength. In order to calculate such a special eigenfrequency, the simulation program is required to handle periodic boundary conditions, i.e. the field distribution on the downstream boundary of the cell is required to be identical to that on the upstream boundary, but rotated in phase by a given phase angle. Not all simulation programs have this useful feature built in. However, if this feature is missing (e.g. in Superfish), one can still study travelling-wave structures, but only special phase advances (integer multiples of $2\pi/n$), by simulating an assembly of n cells.

3.2 Attenuation

Idealized, lossless cavities have real eigenfrequencies, but in the presence of losses the eigenfrequencies become complex, with the imaginary part of this frequency corresponding to a damping term. From a complex eigenfrequency f_0 , the cavity Q can be calculated from

$$Q = \frac{\text{Re } f_0}{2 \text{Im } f_0}. \quad (8)$$

One often requires that an RF cavity have low losses or a high Q , sufficient to minimize the power needed to produce a necessary accelerating voltage. Maximizing the cavity Q means minimizing the power required for a given stored energy, by making use of the cavity's resonant behaviour. At the same

time, a high Q makes a cavity extremely narrow-band and sensitive to errors, and thus requires high precision in the determination of the resonance frequency (and in the fabrication). The cavity Q is defined as

$$Q = \frac{\omega_0 W}{P}, \quad (9)$$

where W is the energy stored in the cavity, $\omega_0 = 2\pi f_0$, and P is the ohmic loss.

For travelling-wave structures, the Q is defined in the same way, but typically W and P are normalized to the cell length. P again is the ohmic loss, which in this case however is typically only a small fraction of the power input to the cell, since most of the input power travels through. As opposed to a cavity, the Q in this case does not affect the bandwidth of a travelling-wave structure, but rather the attenuation of the travelling wave. With group velocity v_g , which will be defined in Eq. (20) below, the attenuation of the travelling wave is given by

$$\alpha = \frac{\omega}{2Qv_g}, \quad (10)$$

i.e. the forward-travelling power is proportional to $e^{-2\alpha z}$.

3.2.1 Perturbation ansatz for low losses:

It is numerically much easier to find real eigenfrequencies of lossless structures than to find complex eigenfrequencies. A high Q now allows the assumption that the field distribution for the ideal, lossless case is a good approximation of a real one, and that the losses represent only a small perturbation of this ideal solution. The idea of this perturbation ansatz is, starting from the ideal solution, to calculate the tangential magnetic field at the perfectly conducting cavity boundary, equating it to the surface current density, driving this surface current through the surface resistance which is given by

$$R_A = \sqrt{\frac{\omega \mu}{2\sigma}}$$

with wall conductivity σ , and to determine the wall losses by evaluating the integral

$$P = \frac{R_A}{2} \iint_A |H_t|^2 dA.$$

For high- Q RF structures, this perturbation method is very efficient and exact. Most of the simulation programs use this method.

3.3 *R-over-Q*

A third important quantity characterizing an RF structure is the *R-over-Q* ($= R/Q$), which accounts for the influence of the structure geometry; one could say that it describes the feature of the structure that concentrates the electric field where it is actually needed for acceleration. The R/Q is expressed as

$$\frac{R}{Q} = \frac{|V_{acc}|^2}{\omega_0 W}. \quad (11)$$

In Eq. (11), V_{acc} is understood as the accelerating voltage “seen” by the beam, including the effect of its finite speed v , i.e.

$$V_{acc} = \int_{-\infty}^{\infty} E_z e^{j\frac{\omega}{v}z} dz. \quad (12)$$

We have used infinite integration limits in this formulation to indicate that the accelerating field may actually extend beyond the actual structure into the beam-pipe at both ends. However, for a resonance frequency below the cut-off of the beam-pipe, where the fields cannot propagate into it, it is valid to limit the integration to only a little more than the RF structure length.

Note that sometimes R/Q is defined with the finite speed of the beam ignored, and this is subsequently corrected for by introducing the so-called transit-time factor, given by

$$\frac{\left| \int_{-\infty}^{\infty} E_z e^{j\frac{\omega}{v}z} dz \right|}{\left| \int_{-\infty}^{\infty} E_z dz \right|}. \quad (13)$$

The transi- time factor is a number between 0 and 1.

Whereas f_0 and Q are typically directly output from a simulation program, the integral in Eq. (12) requires the possibility of evaluating a path integral over a field component multiplied by a trigonometric function. In the computer codes mentioned below, this functionality either is integrated (Superfish) or can be performed by a sequence of macro commands.

When investigating a single cell of a periodic structure, the integral (12) is extended only over the length l of a single cell. It is common for travelling-wave structures to give R/Q (and R , to be defined below) per unit length (in units of Ω/m) and to denote them with a prime; the definition equivalent to Eq. (11) for travelling wave structures is

$$\frac{R'}{Q} = \frac{\left| \frac{V_{acc}}{l} \right|^2}{\omega_0 \frac{W}{l}}, \quad (14)$$

where V_{acc}/l is called the accelerating gradient and W/l is the stored energy per cell.

3.4 Shunt impedance

Finally, the product of R/Q and Q is the shunt impedance R . We stress here that one should consider R/Q as a more fundamental quantity than R , even though the naming suggests otherwise.

Maximizing R means maximizing the accelerating voltage while minimizing the input power required to obtain it. Combining equations (9) and (11), this can be expressed as

$$R = \frac{|V_{acc}|^2}{P}. \quad (15)$$

Often, an important design task is to match the cavity to the power source at resonance, i.e. to match to the cavity shunt impedance, typically via a standard 50-Ohm transmission line. The power coupler can be considered as an essential part of this impedance matching circuit. To perform this task, the computational tools are required to account for the power coupler, and to compute input reflection. MAFIA, HFSS, Microwave Studio and GdfidL are suited for this task.

Please note that sometimes a different definition of both R/Q and R is used, which is smaller by a factor of 2 than in Eqs. (11) and (15). We are using here the so-called “Linac definition” of these impedances.

3.5 Beam loading and higher order modes

The above mentioned quantities, f_0 , Q and R/Q , characterize primarily the action of the accelerating mode on the beam (and the cavity as seen from the RF power source). In the presence of a high beam current however, not only the

power source, but also the beam itself will excite this mode, which in the case of acceleration will counteract the accelerating voltage and thus reduce it, a phenomenon referred to as beam-loading. Often it must not be neglected that beam loading influences the matching to the RF power source (except for travelling-wave structures, where the power generated by beam-loading is not reflected, but travelling forward).

This excitation of fields in the RF structure by a high current beam occurs not only at this nominal accelerating mode but also at all higher-order modes (HOMs) that exist in the structure. In general, HOMs excited by the beam result in forces on the beam that can deteriorate it – this is why they need considering. These forces can be both longitudinal (thus potentially perturbing the bunch structure of the beam) and transverse (kicking the beam or parts of it sideways). This coupling of the beam to cavity modes is described by the beam-coupling impedance. It peaks at each mode's resonance, but in general is a complex function of frequency. In many cases (e.g. for single-gap cavities), the coupling impedance at resonance is equal to the shunt impedance (for the driven mode, it is shunted by the power source output impedance). Both longitudinal and transverse beam-coupling impedances should be accessible to the simulation tools.

The transverse coupling impedance is defined similar to the longitudinal one, but with one of the two $|V_{acc}|$ in the nominator of Eq. (11) replaced by an integral along z of the transverse Lorentz force. This seems much more complicated at first sight, but thanks to the Panofsky-Wenzel theorem [1], this problem can be reduced to computation of the transverse dependence of the longitudinal coupling impedance. For example, to determine the transverse impedance of a dipole mode polarized in the x -direction, it is sufficient to calculate the longitudinal impedance as described above, but evaluating the line integral at a transverse offset position x . If the longitudinal impedance on axis is zero (which is true for the dipole mode in structures with mirror symmetry with respect to $x = 0$), the transverse impedance is simply given by

$$R_{\perp} = \frac{c}{\omega_0 x^2} R. \quad (16)$$

The transverse impedance is given in units of Ω/m (or – normalized to length in travelling-wave structures – in Ω/m^2).

To consider HOMs, the simulation tool must be capable of determining the quantities defined above for a number of modes, i.e. a set of f_n , Q_n and R_n/Q_n . For transverse modes, the integral (12) must be evaluated off axis.

3.6 Wake-fields – loss factor

So far, we have been referring to quantities describing the frequency domain behaviour of RF structures. Equivalently, they can be calculated and described in time domain; in this case we refer to longitudinal and transverse wake-fields. Wake-fields are the fields left behind in the structure by a driving charge. They are normally characterized by the so-called wake-potentials, which describe the net forces on a trailing charge, following the driving charge at a distance s through the entire structure [2]. The Fourier transforms of the longitudinal and transverse wake-potential are the longitudinal and transverse beam-coupling impedance, respectively.

One distinguishes the short-range (for small s) and the long-range wake-potentials. The short-range wake-potentials are determined by wide-band and high frequency components of the coupling impedance. Physically, they are dominated by those discontinuities inside the RF structure that are closest to the beam trajectory (typically irises). They are not well accessible to frequency-domain calculation. The high- Q resonances on the other hand, which can be well described in the frequency domain, dominate the long-range wake-potentials. They are dangerous in particular because of their adverse effect on subsequent bunches following at comparatively large distances.

The long-range wake-potentials may be considered as a superposition of the contributions of the individual modes. For longitudinal wake-potentials e.g., the excitation of each of these modes is described by its (modal) loss factor

$$k_{loss} = \frac{|V_{acc}|^2}{4W} = \frac{\omega_0}{4} \frac{R}{Q}. \quad (17)$$

Note that this excitation is independent of the Q of the considered mode, but directly proportional to its R/Q . The energy lost by a charge q into this mode is $k_{loss} q^2$, and the contribution of this mode to the longitudinal wake-potential is

$$W_z(s) = \begin{cases} 0 & \wedge \quad s < 0 \\ -k_{loss} q & \wedge \quad s = 0, \\ -2 k_{loss} q e^{-\frac{\omega_0 s}{2Q c}} \cos\left(\omega_0 \sqrt{1 - \frac{1}{4Q^2}} \frac{s}{c}\right) & \wedge \quad s > 0 \end{cases} \quad (18)$$

where s is again the distance behind the charge exciting the wake, and the minus sign indicates that the wake will be dragging on this driving charge. For transverse wake-fields, the transverse loss factor is defined equivalently; it is sometimes also referred to as kick-factor.

Because of Eq. (17), the modal loss factor is readily accessible to frequency domain calculations (HFSS), while the total loss factor and the wake-potentials are better determined in time domain simulations (MAFIA, GdfidL).

3.6.1 Direct method

When calculating the wake-fields in time domain, one excites the fields in the structure by a charge moving along its axis or parallel to it, and calculates the resulting electro-magnetic fields. For the so-called direct calculation, one has to follow this charge for a large distance through the vacuum chamber both before and after the actual RF structure.

3.6.2 Indirect method

A trick, known as the indirect method [2], which is possible for certain RF structures, significantly simplifies this: if the charge is assumed to travel grazing along the beam pipe outer boundary, it will excite and experience no fields until it enters into the actual gap or – more generally – the RF structure. The obvious advantage is that long beam pipes need not be discretized, which saves computation time substantially; the necessary condition for applying the indirect method is that upstream and downstream beam pipes have the same shape and dimensions, and that there is no obstacle on these special particle trajectories.

The indirect calculation method is based on the fact that the longitudinal wake-potential satisfies the transverse Laplace equation $\Delta_{\perp} W_z(x, y; s) = 0$. Once the longitudinal wake-potential is known at transverse positions corresponding to the boundary of the beam pipe, it is thus determined anywhere inside the cross-section. It follows from the Panofsky-Wenzel theorem that this also determines the transverse wake-potential.

This is particularly simple for round structures, where it is sufficient to determine the longitudinal wake-potential at the radius of the beam pipe. The indirect integration method was implemented in TBCI, a program which today is integrated into MAFIA. Also GdfidL uses the indirect calculation method.

Transverse wake-fields describe the net effect of transverse forces, both electric and magnetic, on trailing particles. As in the frequency domain, the Panofsky-Wenzel theorem can be used to obtain the transverse wake-fields simply from the radial dependence of the longitudinal wake-fields.

At least in principle, wake-field calculation in the frequency domain is possible thanks to the relation (17), which for each mode relates the loss factor to the R/Q . Having determined the modes of a structure, the wake functions can be approximated by a sum of expressions (18). However, the determination of short-range wake-fields is impractical, since the number of modes to be taken into account would become very large.

3.7 Field distribution

Other quantities required from a field simulation program are the electric and magnetic fields themselves, often also derived quantities like the Poynting vector, energy densities or (in particular for precision control) spatial derivatives of the fields. The possibility of graphical output of the field distributions in the structure is extremely useful in the structure design. Visualizing the fields also helps to check the consistency of the computation result and to understand the physics. Particularly interesting are the maximum electric and magnetic fields at the structure surface, since they will determine breakdown limits and surface heating.

3.8 Group velocity

For travelling-wave structures, the group velocity is another important quantity to consider. It was already used in Eq. (10) above. Since the group velocity is defined as the slope of the dispersion curve, one way to calculate it is to calculate the eigenfrequencies at two different phase advances near the synchronous phase advance. With the two eigenfrequencies f_1 and f_2 at phase advances φ_1 and φ_2 in a cell of length l , the group velocity near the frequency $(f_2 - f_1)/2$ is calculated from

$$v_g = 2\pi l \left(\frac{f_2 - f_1}{\varphi_2 - \varphi_1} \right). \quad (19)$$

On the other hand, the determination of the energy stored in a cell of a periodic structure, along with the calculation of the power flow (the integral over a cross-section of the real part of the Poynting vector) allows the direct determination of the energy velocity, which is equal to the group velocity, from

$$v_g = \frac{P_{\text{forward}}}{W/l}. \quad (20)$$

3.9 Time domain versus frequency domain

We have discussed both time domain and frequency domain characteristics of RF structures above. For linear systems, the formulations in time domain and frequency domain are equivalent – one is the Fourier transform of the other. However, for specific problems, one or the other approach is better suited.

For a structure with a single resonance with a very high Q e.g., a time domain calculation would require long computation time, unless one starts off with the steady-state solution. Also if the excitation is known to be only at one frequency (like from the RF power source), a frequency domain calculation is advantageous. On the other hand, the calculation of wake-fields in a structure with no pronounced resonance (or very many modes contributing to it) is more economical directly in the time domain. This is particularly true for short-range wake-fields. If the length of the bunch exciting the wakes is not too short, both time domain and frequency domain calculations require approximately the same computational effort.

For low energy, high current beams, where the excitation of fields by the beam itself (the wake-fields) can influence the particle trajectories, an exact calculation would have to be “self-consistent.” This problem is – strictly speaking – nonlinear and thus cannot be performed in the frequency domain. Only MAFIA’s TS3 module allows today the self-consistent computation of particle trajectories and RF fields excited by them (and influencing those trajectories). A new version of GdfidL, which is currently under test at TU Berlin, can equally handle this difficult case.

However, for most accelerating structures a self-consistent analysis is not necessary. In particular for electron linacs, where the speed of the particles is close enough to c , their trajectory is known beforehand to any practical precision.

4 Specific Computer Codes

A large number of computer codes are available for the calculation of RF fields. Some of the commercially available codes however are not particularly suited for the design of accelerator cavities, since *a)* they do not allow the direct calculation of eigenfrequencies, or *b)* they do not provide a possibility of computing the integral (12) which describes the interaction with a particle beam. So we will restrict ourselves here to a few codes relevant for RF accelerator structure design.

4.1 Superfish

Superfish calculates the eigenfrequencies and fields for round, but otherwise quite arbitrarily shaped cavities, which can be described completely in the two co-ordinates r and z . The output of Superfish includes all quantities relevant for the cavity design directly. Except for trial versions, there is no version of Superfish available today that treats modes with non-zero azimuthal orders. The easy to use PC version is available from the “Los Alamos Accelerator Code Group” (<http://laacg1.lanl.gov/>) free of charge, but users are encouraged to register.

Superfish is a suite of programs, also including Poisson and Pandira for the solution of static fields.

Superfish starts off by defining a triangular mesh inside the cavity. The program “Automesh” then deforms these triangles, so as to better fit the mesh to the given physical boundaries. In the next step, “Fish” or “CFish,” the Helmholtz eigenvalue problem is solved for, which for an accelerating mode of a round cavity can be reduced to the form

$$\nabla^2 H_\varphi = \left(\left(\frac{\omega}{c} \right)^2 - \frac{1}{r^2} \right) H_\varphi = 0. \quad (21)$$

Normally, “Fish” is used, which solves for real eigenvalues, i.e. solutions for closed, lossless cavities with perfectly conducting walls. “CFish” is a complex eigenvalue solver, see Eq. (8).

The following post processor “SFO” allows calculation of a number of relevant parameters, such as Q (using the perturbation ansatz), R/Q , the field distribution and the sensitivity of the solution to small geometric perturbations.

The two co-ordinates used in Superfish can be either r and z of the cylindrical co-ordinates (round cavities, see example above) or Cartesian (arbitrary cross-section waveguides). For the latter case, Superfish allows one to solve for 2-dimensional boundary value problems, i.e. for modes in waveguides of arbitrary cross-section. This is particularly interesting for the investigation of RFQs, and there even exists a utility program “RFQfish” which helps to set up the geometry of the vanes of an RFQ for the determination of the purely transversal electric fields, i.e. without the longitudinal modulation of the vane tips. Also a number of other utility and tuning programs and a very complete documentation are part of the Superfish distribution.

As an example, Figure 2 shows the result of a Superfish run for a single-cell cavity with nose cones to optimize the shunt impedance. Only one half of the

Figure 2: Example of a Superfish output (WSFplot).

cavity is modelled, and the left boundary is assumed to be a symmetry plane. To demonstrate the simplicity, the complete input file listing is given here:

```
Nose cone example
&reg kprob=1, dx=.3, freq=500, icylin=1,
  xdri=0, ydri=23.469, norm=1,
  ezerot=5000000, kmethod=1, beta=1 &
&po x= 0.0, y= 0.0 &
&po x= 0.0, y= 23.469 &
&po x= 2.0, y= 23.469 &
&po x= 15, y= 10.469, nt=5, radius=13 &
&po x= 15, y= 9.1375 &
&po x= 11.5, y= 6.866 &
&po x= 12, y= 5, nt=4, radius=1 &
&po x= 30.0, y= 5.0 &
&po x= 30.0, y= 0.0 &
&po x= 0.0, y= 0.0 &
```

We give here also an excerpt from the output file of SFO, which illustrates the completeness of the solution:

```

Superfish output summary for problem description:
Nose cone example
Problem file: C:\LECTURES\LONG BEACH\S500.AF 7-01-2003 17:46:48
-----
All calculated values below refer to the mesh geometry only.
Field normalization (NORM = 1): EZEROT = 5.00000 MV/m
Frequency = 500.00301 MHz
Particle rest mass energy = 938.271998 MeV
Beta = 1.0000000
Normalization factor for E0 = 7.120 MV/m = 81300.178
Transit-time factor Abs(T+iS) = 0.7022185
Stored energy = 6.4623931 Joules
Using standard room-temperature copper.
Surface resistance = 5.83374 milliohm
Normal-conductor resistivity = 1.72410 microhm-cm
Operating temperature = 20.0000 C
Power dissipation = 458.5893 kw
Q = 44271.3 Shunt impedance = 33.166 MOhm/m
Rs*Q = 258.267 Ohm Z*T*T = 16.355 MOhm/m
r/Q = 110.825 Ohm wake loss parameter = 0.08704 V/pC
Average magnetic field on the outer wall = 21507.1 A/m, 134.922 W/cm^2
Maximum H (at Z,R = 14.125,8.56963) = 29395.9 A/m, 252.053 W/cm^2
Maximum E (at Z,R = 11.0267,5.77084) = 37.5597 MV/m, 1.76356 Kilp.
Ratio of peak fields Bmax/Emax = 0.9835 mT/(MV/m)
Peak-to-average ratio Emax/E0 = 5.2750

```

Superfish will output the numbers for the mesh geometry only, but for the calculation of the losses, it distinguishes between the symmetry planes and the metallic boundaries. This example used roughly 10 000 meshpoints and took a few seconds to run on a Windows PC. Many other data are output by SFO, in particular a complete sensitivity analysis of the resonance frequency on geometric parameters is included.

Summarizing, Superfish has an excellent mesher, the mesh allows one to model curved boundaries quite well, and consequently the results can be very exact. On the other hand, the program is limited to modes with azimuthal index 0 in round cavities (2-d). The output is very comprehensive for RF cavity design. The program runs on Windows, is free of charge and is well maintained by LAACG.

4.2 MAFIA

MAFIA (Solution of MAXwell's equations by the Finite Integration Algorithm) is probably the most widely used software package for accelerator RF structure design. It is based on the FIT algorithm described above. Apart from pre-processor (M, mesh generator) and post-processor (P), it consists of several modules for different tasks: The static solver (S), the eigenmode solver (E), the frequency domain (or Eddy current) solver (W3), the time domain solvers (T2, T3), and the self-consistent time domain solvers including particle dynamics (TS2, TS3). To our knowledge, the latter is a unique feature of MAFIA today.

The latest version 4 of MAFIA includes as well an easy to use graphical user interface and an optimizer module (O), which can be used as a convenient wrapper around the other modules.

In the design process of a conventional accelerating cavity with MAFIA, one would typically start with a simplified, 2-D geometric model (r - z), and calculate the eigenfrequencies, stored energies and interaction integrals to determine Q 's and coupling impedances. These figures of merit are not directly output from MAFIA, but the relevant operations can be performed in the P-module or programmed in the form of command macros. For a vacuum cavity, the design would involve the E-module to calculate the eigenfrequencies for the lossless case, the P-module for the other parameters (again with the perturbation method for small losses). As opposed to Superfish, MAFIA allows for the r - z geometry to input the azimuthal order (for dipole, quadrupole, ... modes).

The time domain solver of MAFIA evolved from the program TBCI, and it allows the direct or indirect calculation of wake-fields. Also driven solutions for waveguide components or accelerating structures are calculated with the time domain solver; if necessary the output data can be Fourier transformed to get a spectral solution.

MAFIA is supported on different computer platforms. The code is maintained and commercialized by the company CST in Darmstadt, Germany. Detailed information can be obtained from <http://www.cst.de>.

The strength of MAFIA is certainly its versatility. We believe a weakness is its comparatively rigid mesh, which is strictly Cartesian. It can be refined locally, but the refining mesh-lines run through the entire geometry, which results sometimes in quite unevenly distributed meshes. Another important peculiarity of MAFIA is the accounting for the losses at concave, curved boundaries, where the modelled boundary is staircase-like (at least in one of the dimensions of the investigated boundary). In this case, MAFIA tends to over-estimate the Q , and this does not seem to improve with mesh refinement.

4.3 HFSS

HFSS (High Frequency Structure Simulator) is a commercial package from Ansoft, Pittsburgh, PA (<http://www.ansoft.com>), primarily intended to analyze S -parameters and radiation problems. It performs only frequency domain calculations (and is thus strictly limited to linear materials) using the finite element method, solving Eq. (4) under given boundary conditions. Its tetrahedral mesh adapts well to curved surfaces and is co-ordinate independent (isotropic). The implemented adaptive mesh refinement scheme evaluates the residuum (6) to identify the tetrahedra that contribute most to the error, which elegantly allows refining the mesh automatically where it matters most. The program runs under Windows or Solaris operating systems.

The graphical user interface of HFSS is easy to use, but it also allows the use of a powerful, however sometimes counter-intuitive, macro language, including optimizer commands. The RF structure designer's figures of merit, like R/Q , shunt impedance, etc., are however not calculated automatically but require macro programming – a HFSS post-processor macro to calculate the interaction integral (12) along the line “zaxis” could look like:

```
Assign c0 299792458.0
EnterScalar 0;MathFunc "Acos";EnterScalar 4;*      # 2 pi
EnterScalar GetUnitConv;/;EnterScalar c0;/
EnterFrequency;*                                     # omega/c
EnterScalarFunction "Z"; *; Push
MathFunc "Cos"; CmplxR
Exchange
MathFunc "Sin"; CmplxI; +                             # exp(j omega/c z)
Enter "E"; ScalarZ; *; Push; Real                     # Re and Im
EnterLine 'zaxis'; Integrate; Evaluate                # require two
Rollup; Rollup; Imag                                  # separate
EnterLine 'zaxis'; Integrate; Evaluate                # integrations!
Rollup; Pop; Rollup; Rollup; Pop; Rollup; Push; *
Exchange; push; *; +; sqrt                            # ABS
GetTopEntryValue
```

In this example it can be seen that a convenient IPN-type calculator is used for the manipulation of field quantities in the post-processor.

The actual version, HFSS 8.5, allows also the calculation of eigen-frequencies including substantial losses – even modes with Q 's well below 10 can be calculated.

HFSS allows defining periodic boundary conditions (master/slave) as they are helpful for the analysis of periodic structures. However, with the actual version HFSS 8.5, we have encountered some mesher problems in the presence of periodic boundary conditions.

4.4 GdfidL

GdfidL (<http://www.gdfidl.de>) is a program written by W. Bruns of Technische Universität Berlin, Germany. It is in many respects similar to MAFIA. It also uses the dual-grid FDTD approach, but the production version does not handle particle dynamics. The self-consistent particle dynamics code is now under evaluation at TU Berlin. Effort has been put into better meshing near the boundaries, which increases accuracy. This was achieved by introducing diagonal fillings of mesh cells near the boundary. Also the computation efficiency was increased. The mesher of GdfidL is very economical, since the field-free regions of the bounding box are not meshed.

The probably most valuable new feature of GdfidL is its capability to be run on a cluster of Linux-PC's in parallel. This allows highly efficient solution of huge problems, which would otherwise exceed the maximum memory size of 32-bit PCs.

GdfidL has no Graphical User Interface, but a powerful macro language, similar to that of MAFIA. Shown here is an excerpt from a GdfidL command input file which shows the commands to control the calculation of a wake-potential:

```
-fddt
-lcharge
  # The properties of the exciting line-charge.
  charge= 1e-12, sigma= SIGMA
  xposition= 1, yposition= 0
  shigh= 500

-ports
  # Where are absorbing boundary conditions:
  plane= zlow, name= zlow, npml= 40, modes= 0, doit
  plane= zhigh, name= zhigh, npml= 40, modes= 0, doit
  npml= 20
  plane= xhigh, name= xhigha, npml= 40, modes= 0, doit
  plane= yhigh, name= yhigha, npml= 40, modes= 0, doit

-storefieldsat
  what= e, name= a
  firstsaved= 4 * Period / @clight
  lastsaved= 5 * Period / @clight
  distance= 0.2 * Period / @clight
  doit
-fddt
  # Start the time domain computation.
  doit
```

Note in the above example the PML (perfectly matched layer) special boundary conditions which can account for absorbing boundary conditions.

4.5 CST Microwave Studio

Microwave Studio is a program by CST (like MAFIA), specially designed for the calculation of S -parameters, antennas and radiation problems. It runs in a Windows environment. It has a graphical user interface and uses Microsoft's "Visual Basic" as a macro language; the following lines of code provide an example.

```
'@ define automesh state
  Mesh.Automesh "True"

'@ define frequency range
  Solver.FrequencyRange "2800", "3000"

'@ define solver parameters
  with Solver
    .CalculationType "TD-S"
    .StimulationPort "1"
    .StimulationMode "1"
    .SteadyStateLimit "-30"
    .MeshAdaption False
    .FrequencySamples "1000"
    .StimulationType "Gaussian"
  End with
```

Compared to MAFIA, the approximation of boundaries that do not coincide with mesh planes is improved significantly in Microwave Studio by the so-called "Perfect Boundary Approximation".

The actual version 4.3 of Microwave Studio does not allow wake-field calculations, since it does not provide for moving charges as field sources.

4.6 ANSYS Multiphysics

ANSYS (Canonsburg, PA) offers a commercial suite of finite element programs intended primarily for mechanical structural analysis, including heat conduction and coolant flow. Part of this distribution is Multiphysics, which includes the possibility of simulating RF fields in the "high-frequency electromagnetics" module (<http://www.ansys.com/ansys/multiphysics.htm>). The strengths of this program are the excellent mesher, which can create tetrahedral or hexahedral meshes, and the possibility of integrating heating/cooling and mechanical deformation directly with the RF simulation. Only real eigenfrequencies can be calculated (lossless structures), the cavity Q is determined in the post-processor using the perturbation ansatz.

As in most of the other programs, there exists a macro language, which allows one to write user macros, e.g. to calculate in the post-processor the line integrals necessary to determine R/Q . Also ports can be defined for a driven

solution (in the frequency domain), but only single-mode, simple waveguide geometries, insufficient for the needs of RF structure design.

Ansoft Multiphysics 5.7 does not handle periodic boundary conditions. Perfectly Matched Layer (PML) boundaries are possible. We regret the lack of an adaptive mesh refinement in version 5.7, which could help control the attained precision. This will be improved in version 7.

5 Comparison of RF field simulation programs

The simulation tools presented above each have their strenghts and weaknesses. However, it is interesting to run a case that can be handled by all of them, preferentially one where the exact solution is known. In the following, we will not compare the different functionalities of these programs, which have been mentioned above. The following comparison is limited to the precision obtained in the determination of the resonance frequency and the quality factor.

5.1 Which precision is required

The ultimate goal of a simulation program is the design of an RF structure to a precision that directly allows its fabrication. In this case, it is also possible to predict how mechanical tolerances will eventually influence the performance of the accelerator structure. To this end, the cumulated numerical errors must be below the fabrication tolerances.

A high cavity Q implies a sharp resonance, or a small bandwidth, which on the other hand immediately imposes a requirement on the precision with which the resonance frequency can be predicted. For a Q of 10^4 , e.g., the relative frequency error should be well below 10^{-4} , otherwise there is the risk that the gap voltage will be well below the design value for a given input current.

In choosing the precision required in determining the resonance (or synchronous) frequency, another consideration is comparison with realistic fabrication tolerances. The error of the calculation should be at least in the same order as the error of the fabrication tolerances. It seems that for typical construction techniques, where a precision of 0.01 mm for typical sizes of 100 mm seems realistic, one equally obtains a requirement of 10^{-4} .

As far as the required precision for the calculation of Q is concerned, it must be remembered that an error in Q of 10 % will demand 10 % more input power to obtain the nominal accelerating voltage.

The example chosen for the following benchmark is a spherical resonator; the exact solution for its TM_{110} eigenfrequency is given by

$$f_0 = \frac{\chi c}{2\pi a},$$

where a is the radius of the sphere, c is the speed of light and $\chi = 2.7371\dots$ is a zero of the spherical Bessel-function $\sqrt{\frac{\pi x}{2}} J_{3/2}(x)$. The exact solution for the Q of this mode of the spherical cavity is

$$Q = \sqrt{2a\mu\sigma} \frac{a \int_0^\chi \left(\frac{\sin(x)}{x} - \cos(x) \right)^2 dx}{\chi \left(\frac{\sin(\chi)}{\chi} - \cos(\chi) \right)^2},$$

where σ is the conductivity of the metallic wall of the vacuum cavity. For a 50-cm-radius copper sphere, $\sigma = 58 \text{ MS/m}$, the “exact” numerical result is

$$\begin{aligned} f_0 &= 261.823 \text{ MHz}, \\ Q &= 89899.1 \end{aligned}$$

We summarize the results in Figure 3 indicating the deviation of both the eigenfrequency and the Q from the exact solution. We have included in this comparison some programs that were not mentioned above, notably URMEL (which is identical to MAFIA in r - z Co-ordinates), URMEL-T (using a triangular mesh) and SuperLANS. Furthermore, the number of mesh elements is indicated, which are mesh-points or tetrahedra.

Trying to interpret these results, it becomes apparent that, with the exception of the 2-d programs Superfish and SuperLANS, hardly any program is able to compute the frequency to the desired precision of 10^{-4} . The typical obtained precision is approximately $5 \cdot 10^{-4}$ while the MAFIA results with Cartesian mesh are only around $3 \cdot 10^{-3}$, which is believed to be related to the limited approximation of the spherical geometry.

Figure 3: Results of the benchmark calculating resonance frequency and Q of a spherical cavity.

Regarding the Q calculation, the general results are similar; note however that MAFIA overestimates the Q for this geometry. The errors in the percent range seem acceptable.

References

1. S. Vaganian and H. Henke: "The Panofsky-Wenzel Theorem and General Relations for the Wake Potential," *Particle Accelerators* 48, 239-242 (1995).
2. T. Weiland and R. Wanzenberg: "Wake Fields and Impedances," CERN US PAS 1990:0039-7, DESY-M-91-06 (1991).