EIGENMODES OF SUPERCONDUCTING CAVITIES CALCULATED ON AN APE-100 SUPERCOMPUTER (SIMD)

F. Neugebauer*, DESY Zeuthen, U. van Rienen†, University Rostock

Abstract

The construction of modern accelerators is usually supported by the numerical determination of eigenmodes in the accelerating cavities. Often the rotational symmetry of the cavity is used to simplify the numerical simulation. However, in cases where the cavity lacks rotational symmetry resp. where attached components like couplers have to be taken into account, a fully three dimensional treatment of the maxwell equations is necessary. This requires more computer power than is available on a normal high end workstation. Therefore, in the present approach a parallel SIMD super computer (APE–100) is used to compute the eigenmodes of accelerating cavities. As an example parts of the superconducting TESLA structure are investigated. The geometry input is parsed by MAFIA which transfers the resulting system matrix, incorporating geometry and material information, to the APE–100. The result of the diagonalization procedure is then read back to the MAFIA host where further data analysis and visualization can be done.

1 INTRODUCTION

The construction of modern accelerators is usually supported by the numerical determination of eigenmodes in the accelerating cavities. Often the rotational symmetry of the cavity is used to simplify the numerical simulation. However, in cases where the cavity plus attached rf–components lacks rotational symmetry a fully three dimensional treatment of Maxwell's equations is necessary which requires more computer power than is available on a normal high end workstation. In addition the three dimensional approach allows for the simulation of fabrication errors and surface roughness which are usually not considered to have rotational symmetry.

In the framework of the Finite Integration Technique (FIT) developed by Weiland and coworkers [1] Maxwell's equations in integral representation are transformed to a set of matrix equations. Using rectangular grids the discretization volume is partitioned in two sets of cells which can be considered dual. In the case of determining the eigenmodes of a cavity the grid voltages along neighboring gridpoints are the degrees of freedom of the resulting eigenvalue problem. The system matrix connects grid voltages of a single cell only to grid voltages of adjacent cells. This "next neighbor connection"–property makes the eigenvalue problem especially well suited to be solved on an APE–100 supercomputer for this type of computer is capable of a very fast data exchange between neighboring nodes.

APE–100 supercomputers are mainly used in QCD theory where a profound experience in solving eigenvalue problems [2] does exist. However, the parallel structure of the computer requires the use of special programming tools and a language (TAO) dedicated to the computer topology which is inefficient in programming advanced file IO, string evaluation and in managing pointers.

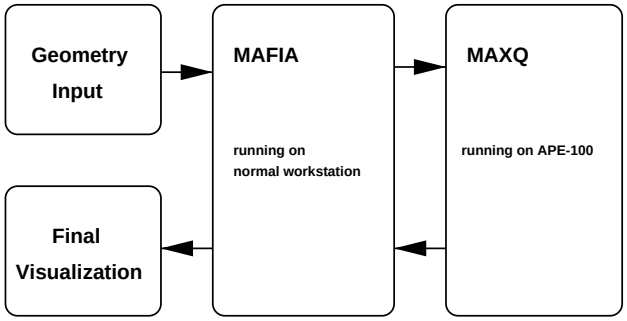

Figure 1: Schematic view of the MAXQ software interface

The parsing of the geometry input, which is mainly a linear task, is left to MAFIA which is running on a usual workstation. The resulting system matrix incorporating geometry and material information is transfered to the APE–100 by means of the MAFIA toolkit (MTK). Then, on the APE–100 supercomputer the numerical expensive task of finding the lowest eigenvalues and corresponding eigenvectors of a large sparse matrix is performed. The result of the diagonalization procedure is then read back to the MAFIA host where further data analysis and visualization can be done.

The approach has been tested first for simple geometries such as sphere and brick in a box which allows for comparison with analytically known results. Next, parts of the TESLA superstructure were studied for a timing comparison between APE–100 and a HP 735.

2 FINITE INTEGRATION TECHNIQUE AND THE APE–100 TOPOLOGY

The Finite Integration Technique is based on a discretization of Maxwell's equations using a set of two rectangular grids which can be considered dual to each other [1]. The integral representation of Maxwell's equations is transferred to a discrete version by specifying the integration paths as to be along the edges of the discretization cell. For the case of area integrals the 6 bordering rectangles of the

* Email: fneug@ifh.de
† Email: van.rienen@e-technik1.uni-rostock.de

cell are chosen as the integration area. The degrees of freedom in the discretized version of Maxwell's equations are not the fields themselves, but the grid voltage along neighboring grid points or the flux over a cell surface for example. Therefore the discretized Maxwell's equations remain mathematically equivalent to the continuous case. There is no discretization error and the discretized Maxwell's equations exactly obey the conservation law for charge and current density.

Starting point for the calculation of eigenmodes in superconducting cavities is the matrix equation

$$\left\{ D^{-1/2}\tilde{C}\tilde{D}CD^{-1/2} - D^{1/2}\tilde{S}^T\tilde{S}D^{1/2} \right\} \vec{u} = \omega^2 \vec{u}, \quad (1)$$

where the matrices D, C and S contain material and mesh information and represent the $curlcurl - graddiv = -\nabla^2$ operator. $\vec{u}$ denotes the vector of all grid voltages and ω is the frequency of the wanted eigenmode. In the case of homogeneous material distribution Eq.1 is known as Helmholtz's equation.

An important feature of the matrix equations is its locality. Actually this is due to the fact that the chosen integration space is restricted to the neighboring cells of the selected degree of freedom. As a consequence the resulting system matrix of which the eigensystem has to be determined is sparse with a priori known pattern of entries. A detailed treatment of the theory yields that 13 elements of the system matrix are non zero for each degree of freedom. These elements connect to degrees of freedom belonging to neighboring cells.

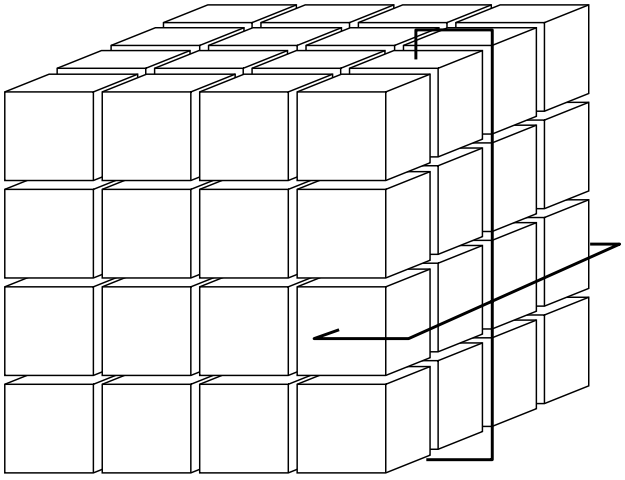

Figure 2: Topology scheme of the APE–100: Each node can access the memory of its six neighbors without latency. The whole cube is subjected to periodic boundary conditions along x-,y- and z-direction resulting in a hyper torus. On each node resides the same amount of grid points and the associated parts of the system matrix (see Eq.1) and the vector of the grid voltages u.

The cells of the discretization volume are distributed to the nodes so that each processor is responsible for its own segment of real space. Then the matrix–vector multiplication is mainly a local operation on each node. Only in cases where the cell lies on the segment boundary data exchange with neighboring nodes will occur. The APE–100 is perfectly suited to such a situation because the SIMD character of the supercomputer accounts for a very fast data transfer with neighboring nodes without latency.

3 RESULTS

The efficiency of the proposed procedure definitely depends on the chosen eigenvalue algorithm. At the moment only a simple variant of a polynomial iteration algortihm is installed. It searches for the lowest eigenvalue in the subspace orthogonal to all eigenvectors already found. No convergence checking has been installed, the interation is simply performed 5000 times. In examples this number has been found sufficiently high to give precise results.

The application of more sophisticated algorithms is in progress. An accelerated conjugate–gradient algorithm for the computation of the lowest eigenvalues is known from Lattice–QCD and tested on the APE–100[2]. Another candidate is the Jacobi–Davidson algorithm.

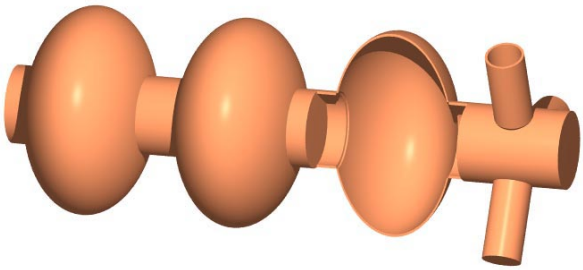

Figure 3: Scetch of the part of the TESLA structure under consideration: Three half cells and a HOM–coupler.

As realistic example a part of the TESLA superstructure consisting of 3 half cells and the beam pipe was discretized on a mesh of 400000 grid points. The first 10 eigenvalues were searched for on a HP 735 in a regular MAFIA session and on the QH4 (512 nodes) running MAXQ. Even with the simple eigenvalue algorithm the APE-100 computed the eigenvalues within only 15 minutes whereas the HP 735 needed about three hours for the same task. The speed-up will be surely increased by the use of better eigenvalue algorithms.

Another important influence on the efficiency comes from the code performance of the matrix by vector multiplication which is the heart of every iterative eigenvalue algorithm for sparse matrices. Measurements show a range from 10 to 20 % of the peak performance depending on the number of grid points per node. Usually, as the number of grid points increases the performance ratio rises. This gives a total computer power of about 2.5 to 5 GFlops on

the QH4.

The total number of grid points which can be used in the discretization is restricted due to memory limitations. Furthermore the used eigensolver needs additional memory for every eigenvector to be found. The memory needed per gridpoint is 52 octets for the system matrix and 12 octets for every eigenvector. This leads to

$$M_{total} = 12 * N * (n + 5) \text{ octets} \qquad (2)$$

for an expression for the total memory used. N is the number of gridpoints and n denotes the number of eigenvectors to be found.

On the QH2 (8x8x4) with 16 MB per node (2) results in:

Number of eigenvectors	Available gridpoints
1	60.000.000
2	51.000.000
5	36.000.000
10	24.000.000
100	3.400.000

Investigations considering the role of the single precision floating point arithmetic of the APE–100 are still necessary. The scalar products are reported to be sensitive to the single precision / double precision problematics. Therefore the scalar products are coded using a software emulation of double precision arithmetics. The same procedure has already been applied in the original MAFIA package to save memory and keep rounding errors at a minimum.

4 CONCLUSIONS

The calculation of eigenmodes on a three–dimensional basis is crucial for the simulation of accelerating cavities. Only a fully three–dimensional treatment of Maxwell's equations can account for effects connected to fabrication errors and surface roughness and most important to effects which arise from devices such as input couplers or HOM–couplers. These coupling devices inevitably break the rotational symmetry of the cavity and cannot be neglected for the determination of eigenmodes of the cavity. However, the lack of computional power on high end workstations normally avoids the inclusion of the three–dimensional effects described above. Therefore in the present paper an approach to a supercomputer solution of the eigenmode problem of superconducting cavities has been made. The reached speed-up by using the APE–100 is at the moment at a factor of about 12, but it is promising that the implementation of better eigenvalue algorithms will increase the efficiency by another factor in the range between 3 to 10.

Not only the computation time for moderate problems is dramatically decreased, moreover the recent approach now opens the possibility to model large scale problems of several 10^6 grid points.

5 ACKNOWLEDGEMENTS

The authors wish to thank Prof. Th. Weiland and his coworkers, especially M. Timm for many helpful discussions on MAFIA and MTK and H-W. Glock and D. Hecht for providing MAFIA command files containing the TESLA geometry. Thanks also to H. Simma for helping implementing the eigensystem solver and for many hints concerning the APE–100. The support given by DESY in providing computing time for the MAXQ project (Maxwell's Equations on Quadrics) is gratefully acknowledged.

6 REFERENCES

[1] Th. Weiland, 'On the Numerical Solution of Maxwell's Equations and Applications in Accelerator Physics', Part. Acc. , **15**, 245–291 (1984)

[2] Th. Kalkreuther, and H. Simma, 'An accelerated conjugate gradient algorithm to compute low-lying eigenvalues – a study for the Dirac operator in SU(2) lattice QCD', DESY 95-137, HUB–IEP–95/10 (1995)