Iterative Solution of Global Electromagnetic Wavefields with Finite Elements

A. Jaun, K. Blomqvist, A. Bondeson, T. Rylander
Department for Electromagnetics, CTH, SE-412 96 Göteborg, Sweden

Abstract. The time-independent Maxwell equations are solved iteratively in 2D geometry for 3D global waves in plasma physics. Krylov space methods, such as the generalized- or the quasi-minimal residuals (GMRES or QMR), are applied together with an incomplete factorization (ILU) preconditioning to a formulation using nodal elements for the electromagnetic scalar and vector potentials. The plasma response is represented as a complex, frequency dependent, dielectric tensor operator and can be used for a variety of applications involving low frequency waves in a tokamak. The iterative approach does not only result in considerable memory savings, but it is also more efficient than a direct solution and paves the way for the parallelization of global wave and stability codes.

PACS codes: 52.35.H, 52.65, 02.70
Keywords: Maxwell, plasma, iterative, GMRES, TFQMR

Corresponding author: (hard to reach by regular mail)
André JAUN, Alfvén Laboratory, KTH, SE-100 44 Stockholm, jaun@fusion.kth.se
Tel: +46 70 7971879, FAX: +46 8 7906574

July 29, 2000

1 Introduction

Applications based on the resolution of the time independent Maxwell equations are numerous, ranging from response calculations in computational electromagnetics to radio-frequency heating and stability models for fusion plasmas.

In computational electromagnetics, integral formulations for the fields in simple dielectric media are discretized in complex geometries using unstructured meshes; they yield sparse linear systems and are solved iteratively to keep the memory consumption low [1, 2, 3]. In plasma physics, the linear response is complicated by the motion of charged particles, which strongly affect the dispersion, the dissipation and even create new types of waves. Studying low frequency perturbations, it has recently been found that the *resonance absorption* from fluid plasma models does not properly account for the power transfers in global wavefields [4]. A more detailed gyrokinetic model is required instead to correctly describe the linear *mode-conversion* process, in which power can be transfered between two waves when they somewhere propagate with the same phase velocity. As a consequence, global stability and ion-cyclotron heating calculations in principle need not only to resolve the global extension of the magneto-hydronamic (MHD) or fast-wavefield, but also the shorter scales associated with the mode-converted drift [5], kinetic-Alfvén [6], ion-hybrid [7, 8] or Berstein waves [9]. This severely challenges not only the physical models for the plasma, but also the memory and execution time required for global 2- and 3-D codes.

Whether only the radial direction [10, 11, 12, 13, 14, 15, 16] or the entire poloidal cross-section is discretized in configuation space [17, 18, 19, 20, 21, 22], global wave codes all solve a large linear system of the form $\mathbf{M} \cdot \mathbf{x} = \mathbf{b}$. This is generally performed with a direct lower-upper (LU) factorization of the matrix in band format, using sometimes a frontal implementation to reduce the central memory requirements [19, 21] or relying on a sophisticated decomposition of the matrix to achieve a limited amount of parallelism [13].

In contrast to direct methods, which are inherently sequential and fill the sparse block tridiagonal matrices during the factorization process, the iterative schemes used in this article are based on matrix-vector multiplications in sparse format. These have been efficiently implemented in the PetSc software library [23] and

result in considerable memory savings and should also allow for an easy parallelisation.

Iterative methods are far from new; when he discovered them, Gauss in person said he would use no other method for linear systems with more than three unknowns [24]. Despite early developments intended for positive definite matrices with eigenvalues distributed along the positive real axis, it is only recently with the advent of Krylov space methods such as the generalized minimum residuals (GMRES) [26] and the transpose-free quasi-minimal residuals (TFQMR) [27] that iterative solutions could be expected also for non-Hermitian matrices and eigenvalues dispersed more generally in the complex plane. An excellent tutorial and review of the subject can be found in a recent book by Saad [25].

In the following sections, we show how a suitable formulation of Maxwell's equations, followed by an incomplete factorization ILU in sparse format, yields a linear system that can be efficiently solved with the GMRES or TFQMR methods for a relatively large variety of global wave problems in plasma physics.

2 Method

In order to retain the low frequency electrostatic modes that get coupled in a plasma to the global electromagnetic wavefield, the Maxwell equations in the PENN code [21] are written using the Coulomb gauge $\nabla\cdot \vec{A}= 0$ in terms of the electromagnetic potentials $(\vec{A}, \phi)$:

$$\begin{cases} \nabla^2 \vec{A} + \left(\frac{\omega}{c}\right)^2 \overleftrightarrow{\varepsilon} \cdot \vec{A} + i \left(\frac{\omega}{c}\right) \overleftrightarrow{\varepsilon} \cdot \nabla\phi & = -\frac{4\pi}{c} \vec{j}_{ext} \\ \nabla \cdot (\overleftrightarrow{\varepsilon} \cdot \nabla\phi) - i \left(\frac{\omega}{c}\right) \nabla \cdot (\overleftrightarrow{\varepsilon} \cdot \vec{A}) & = 0 \end{cases} \tag{1}$$

Here ω stands for the frequency, c the speed of light and $\overleftrightarrow{\varepsilon}$ is a dielectric tensor operator modeling the plasma response to an external source of solenoidal current $\nabla\cdot \vec{j}_{ext}= 0$. After integrating by parts the corresponding weak variational Galerkin form, the Coulomb gauge needs only to be imposed on the domain boundary to guarantee in effect that it is satisfied everywhere in a variational sense [28]. At first, it certainly seems uneconomical to use four variables $(\vec{A}, \phi)$ where only three components of the electric $\vec{E}$ (or the magnetic field $\vec{B}$) are phys-

ically independent. An advantage can however immediately be exploited for a class of problems dealing with incompressible waves $\vec{B}_{||} = \nabla \times \vec{A}_{\perp} = 0$, where the directions are parallel and perpendicular to the equilibrium magnetic field: using the same program structure, a switch reduces the number of variables to one half, keeping only two potentials $(A_{||}, \phi)$ to compute for example kinetic Alfvén eigenmodes in a low pressure plasma [6].

To speed up the convergence, the linear system is *preconditioned* from the left with an incomplete factorization ILU(0) of the original matrix $\mathbf{M}$. This entails a decomposition of the form $\mathbf{M} = \mathbf{LU} - \mathbf{R}$, so that for a zero level of *fill-in*, $\mathbf{L}$, $\mathbf{U}$ have exactly the same non-zero pattern as the lower and upper parts of $\mathbf{M}$, and $\mathbf{R}$ is a residual [25]. This partial elimination process combines spatially adjascent variables and affects mainly high frequency modes that have very short wavelengths and are limited by the resolution of the mesh.

After preconditioning, an iterative solution is sought using Krylov space methods that are generally applicable for non-Hermitian problems. With GMRES, successive Krylov vectors are constructed iteratively using an Arnoldi process and are stored after explicit Gram-Schmidt ortho-normalization with respect to the entire sequence of search directions that have been generated. Memory consumption rapidly becomes an issue for large problems and only a limited number n of vectors can then be stored before the algorithm has to be restarted with GMRES(n). The loss of precision attributed to the round-off errors after every restart does however dramatically limit the convergence for low $n \leq 20$ and rapidly brings the method to stagnation. It is then advantageous to use the TFQMR algorithm, where the successive Krylov iterates are constructed in a quasi-orthonormal manner using a Lanczos process without having to store the search directions. Due to the finite precision of the computer arithmetic, this method does sometimes also stagnate, but the gain of precision with a much larger number of iterations n is then sufficient to overcome the loss when the algorithm is restarted with TFQMR(n) and convergence can then be demonstrated down to machine precision. The underlying linear algebra can be brought to substantial sophistication and remains a matter of research; the implementation, however, is greatly simplified by using software libraries such as PetSc [23]. Changing a small number

of switches, it is relatively easy to test a variety of preconditioners and solvers, taking full advantage of matrix operations in sparse format.

For global waves in the vacuum, all the eigenvalues $\{\omega_k\}$ of the discretized Maxwell operator are distributed along the real axis; from the finite condition number of a positive definite matrix $\kappa(\mathbf{S}) = \max(\{\omega_k\})/\min(\{\omega_k\})$, it is then possible to give a formal proof of convergence for a variety of iterative methods. The problem is however more delicate in the presence of a plasma, where low frequency ($\sim$ 1-10 kHz) acoustic and drift waves coexist with high frequency ($\sim$ 100-1000 GHz) electromagnetic waves and the dissipation from resitivity or resonant wave-particle interactions moves the eigenvalues far into the complex plane. Only by trying, is it eventually possible to say if an iterative method converges or not for realistic global wave problems.

3 Results

Our task therefore consists in checking a variety of iterative schemes for problems with a variable size and compare the efficiency with a solver based on direct LU factorization of the block-tridiagonal matrix. Physically relevant wavefields are chosen that are smooth enough to be numerically resolved on a succession of refined grids. Bilinear node-based FEM are compared with their bicubic-Hermite counterparts when the wave equations (1) are solved in a vacuum $\overset{\leftrightarrow}{\varepsilon} \equiv 1$ or in a shearless fluid plasma [21]; only cubic FEM are used to integrate the fourth order spatial operators implied by the gyrokinetic plasma model [30]. To keep the figure of merit independent of hardware issues such as indirect addressing, vectorization and parallelization, the performance is measured in total number of real floating point operations (FLOPS) required by a single processor to achieve a solution with a precision better than $\epsilon = 10^{-6}$, where $\|\mathbf{r_k}\|_2 < \epsilon\|\mathbf{r_0}\|_2$ and $\mathbf{r_k} = \mathbf{b} - \mathbf{A}\cdot\mathbf{x_k}$ is the residual after the iteration k. This level of precision is far better than the physical accuracy which can reasonably be expected from MHD or gyrokinetic plasma models. To nevertheless show that higher precisions can be achieved using TFQMR(50), the norm of the residual $\|\mathbf{r_k}\|_2/\|\mathbf{r_0}\|_2$ is plotted in Fig.1 as a function of the iteration number for each of the largest test cases computed in

this paper. The loss of precision after 200 and 400 iterations visible in one of the curves is clearly off-set by the larger gain from the TFQMR algorithm.

The cost of solving the problem iteratively can finally be compared with direct solution methods, where $4nm^3$ complex operations are required for the LU decomposition and $6nm^2$ for the forward-backward substitution. This reflects the block-tridiagonal structure of the matrix with $n = N_s$ blocks of rank $m = N_\theta N_v N_b$, obtained for a concentric numbering of the nodes on a polar mesh with N_s radial surfaces and N_θ poloidal angles. In addition, $N_v = \{2, 4\}$ counts the number of variables in the problem using two $(A_{||}, \phi)$ or four potential $(\vec{A}, \phi)$ and $N_b = \{1, 4\}$ counts the number of linear or cubic FEM basis functions per variable and per node.

3.1 Vacuum

The first test for a nearly Hermitian problem is to compute the four wave-field components $(\vec{A}, \phi)$ of a cavity mode driven by an antenna in a perfectly conducting cylindrical waveguide. Both, linear and cubic FEM are compared when (Eqs.1) are solved with $\overleftrightarrow{\varepsilon} = 1$ on a homogeneous polar grid $N_s = N_\theta \in \{4, 8, 12, 16, 24, 32, 48\}$. A tiny imaginary part is added to the lowest cavity mode frequency $f = f_{TE01}(1 - 0.003i)$ to form a well posed response problem similar to the one used in Fig.5 of Ref.[21]. A variety of preconditioners, including ILU(p) with p=0,1,2, successive over-relaxation (SOR), block-Jacobi and none, have been tried in combination with GMRES and TFQMR solvers.

In general, ILU(0) + GMRES is the most efficient and robust combination when all the Krylov vectors are stored. Restarting with GMRES(n) after a limited number $n = 5 - 20$ of search directions have been completed results in a dramatic loss of convergence. It is then by far more advantageous to use the TFQMR algorithm, trading memory for only a moderate increase in the total number of iterations. Using linear FEM, Fig.2 shows that iterative solutions are indeed competitive with direct LU factorization and for small problems TFQMR is nearly as efficient as GMRES. A preconditioning is clearly desirable and the total number of operations required for the preconditioner and the solver is 2-10 times smaller with an iterative approach than for a direct solution.

Because the band gets wider when using cubic FEM, the gain becomes significatively larger. Fig.3 shows that by the time the resolution reaches 32×32, the Krylov space methods beat direct LU factorization by more than an order of magnitude. In view of the asymptotic dependency in Fig.2, it is however not clear if the favourable scaling in N^3 (instead of N^4 with direct methods) holds for even larger resolutions.

Compared with TFQMR without restart, nearly twice as many iterations are necessary to achieve a precision $\epsilon = 10^{-6}$ when TFQMR(n) is restarted every $n = 50$ steps; at the same time, however, the restarts yield the robustness allowing a convergence down to machine precision in Fig.1. This suggests that relatively large sequences are optimal to maintain both the high efficiency of TFQMR and guarantee a reliable convergence to a chosen level of precision.

3.2 Fluid resistive plasma model

Increasing the complexity of the physical problem, a two-fluid resistive plasma is inserted into the cavity to compute an Alfvén wave resonance heating scenario. The resistivity is chosen large enough to spread the resonant layer over about a quarter of the plasma radius. The computational cost for solving this problem iteratively with linear FEM is not shown here, but is in fact very similar to Fig.2. This suggests that both the Hermiticity breaking from a large resistive dissipation and the disparity of amplitudes between different wavefield components does not cause any problem.

3.3 Gyrokinetic plasma model

A series of tests is now performed with the gyrokinetic plasma model [30], using the same Alfvén wave scenario as in the previous section. The mode-conversion parameters are chosen so as to obtain a kinetic-Alfvén wavelength large enough to allow for a succession of refinements of the mesh $N = N_s^{plasma} + N_s^{vacuum} = \{15, 20, 30, 40, 60\}$, keeping the partition $N_s^{plasma} = 4N_s^{vacuum} = 4/3N_\theta$ fixed. From the solution plotted in Fig.4, it is clear that a minimum of $N = 10$ cubic FEM are required to resolve the radial oscillations with a little more than two mesh points per wavelength.

Fig.5 shows that an iterative solution is also here very efficient, reducing the number of operations for the largest resolution by more than an order of magnitude using ILU(0) + TFQMR with a precision $\epsilon < 10^{-6}$ and still nearly an order of magnitude using ILU(0) + TFQMR(50) with a precision $\epsilon < 10^{-12}$.

3.4 Example of practical interest

We finish the study with a calculations for the Joint European Torus (JET) discharge 42677, where a record of 12 MW fusion power has been produced in a 50:50 deuterium-tritium plasma. Fig.6 shows the global electric field of the fast magnetosonic wave heating the plasma in the ion-cyclotron range of frequencies. Contrary to the solutions that have been computed so far, realistic calculations involve a non-circular plasma cross-section and regions with short spatial scales. To keep the size of the problem reasonable, it is then important to be able to densify the structured grid both radially and poloidally in the vicinity of the divertor at the bottom. Of course, by making some grid cells smaller, the condition number of a positive definite matrix rises with the frequency of the shortest wavelength mode, so that the linear system gets generally more difficult to solve.

Using an inhomogeneous grid 42×36 with ILU(0) preconditioning, it takes 208 GMRES and a whole 6408 TFQMR iterations to solve the four fast magnetosonic wavefield components (Fig.6,left) with a relative precision $\epsilon < 10^{-6}$. For GMRES, this is less than a third of the operations and the memory of the direct solver (despite the large number of Krylov vectors stored), but TFQMR needs more than a factor five more operations. The convergence can fortunately be improved by keeping a larger number of *fill-in* elements: after an ILU(p), p=0,1,2 preconditioning, GMRES converges in 208, 104, 55 iterations, TFQMR with no restart in 6408, 447, 91 iterations and ILU(2) + TFQMR(50) in 105 iterations.

More work is required to assess the convergence when global wavefields are solved using different plasma models, a variety of frequencies and larger spatial resolutions. Given the long computation times and large memory now required by a single processor, such issues will be studied and discussed along with the parallelization of the PENN code, using a parallel computation of approximate

matrix inverses with an algorithm such as SPAI [31].

4 Conclusions

Global electromagnetic wavefields have been computed iteratively for plasma physics applications using the toroidal PENN code. They combine an incomplete factorization ILU with Krylov space methods (GMRES,TFQMR) and take full advantage of the sparse matrix algebra from the PetSc library to substantially reduce the memory required by the solver. The method also makes it possible to parallelize the computations on a large number of processors, which has not been done so far in global wave codes for plasmas. A variety of problems have been successfully tested using bilinear or bicubic finite elements and a fluid or a gyrokinetic model for the plasma. The total number of operations required to compute an iterative solution is typically an order of magnitude smaller than using direct LU factorization.

Acknowledgements

It is a pleasure to acknowledge useful discussions with K. Appert, J. Vaclavik and Y. Liu. One of us (K. B.) would like to thank the Chalmers Institute of Technology for the hospitality when this work was carried out as a M. Sc. project. This work was supported by the Swedish National Science Foundation and the National Computer Centers in Linköping.

References

[1] C. F. Wang, J. M. Jin, IEEE Trans. Microw. Theory Techn. 46 (1998) 553

[2] T. Rylander, A. Bondeson, Comput. Phys. Commun 125 (2000) 75

[3] W. C. Chew, J.-M. Jin, C.-C. Lu, E. Michielssen, J. M. Song, IEEE Trans. Antennas Propag. 45 (1997) 533

[4] A. Jaun, K. Appert, T. Hellsten, J. Vaclavik, L. Villard, Phys. Plasmas 5 (1998) 3801

[5] A. Jaun, J. Vaclavik, L. Villard, Phys. Plasmas 4 (1997) 1110

[6] A. Jaun, A. Fasoli, W. W. Heidbrink, Phys. Plasmas 5 (1998) 2952

[7] A. Jaun, S. C. Chiu, T. Hellsten, Nucl. Fusion 38 (1998) 153

[8] M. Brambilla, Nucl. Fusion 38 (1998) 1805

[9] A. Jaun, S. C. Chiu, T. Hellsten, in *Radio-Frequancy Power in Plasmas*, (Proc. 12th Topical Conf., Savannah, Georgia, 1-3 April 1997) American Institute of Physics, Woodbury, NY (1997) 281

[10] R. C. Grimm, J. M. Greene, J. L. Johnson, Methods Comput. Phys. 9 (1976) 253

[11] C. Z. Cheng, Phys. Reports 211 (1992) 1

[12] A. Bondeson, G. Vlad, H. Lütjens, in (Proc. IAEA Tech. Com. Meeting on Advances in Simulation and Modeling of Thermonuclear Plasmas, Montreal, Canada), IAEA, Vienna, (1993) 306

[13] D. V. Anderson, A. R. Fry, R. Gruber, A. Roy, Comput. Phys. 3 (1989) 33. See also R. Gruber, J. Rappaz, *Finite Element Methods in Linear Ideal MHD* (1985), and W. A. Cooper, Plasma Phys. Contr. Fusion 34 (1992) 1011

[14] E. F. Jaeger, L. A. Berry, D. B. Batchelor, Nucl. Fusion 3 (1998) 437

[15] G. T. A. Huysmans, J. P. Goedbloed, W. Kerner, Phys. Fluids B 5 (1993) 1545

[16] M. Brambilla, T. Krücken, Nucl. Fusion 28 (1988) 1813; M. Brambilla, Plasma Phys. Contr. Fusion 41 (1999) 1

[17] R. Gruber, F. Troyon, D. Berger *et al.*, Comput. Phys. Commun. 21 (1981) 323

[18] D. Gambier, A. Samain, Nucl. Fusion 25 (1985) 283

[19] L. Villard, K. Appert, R. Gruber, J. Vaclavik, Comput. Phys. Reports 4 (1986) 95

[20] A. D. Turnbull, E. J. Strait, W. W. Heidbrink, M. S. Chu, H. H. Duong, J. M. Greene, L. L. Lao, T. S. Taylor, S. J. Thompson, Phys. Fluids B 5 (1993) 2546

[21] A. Jaun, K. Appert, J. Vaclavik, K. Villard, Comput. Phys. Commun. 92 (1995) 153

[22] L. Degtyarev, A. Martinov, S. Medvedev, F. Troyon, L. Villard, R. Gruber, Comput. Phys. Commun. 103 (1997) 10

[23] S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith *The PetSc homepage*, `http://www.mcs.anl.gov/petsc`, also in *Modern Software Tools in Scientific Computing*, Edited by E. Arge, A. M. Bruaset, H. P. Langtangen, Birkhauser Press (1997) 163

[24] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, New York, 1994

[25] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Compagny, Boston, 1996

[26] Y. Saad, M. H. Schultz, SIAM J. Sci. and Stat. Comput. 14 (1993) 470

[27] R. W. Freund, N. M. Nachtigal, Numerische Mathematik 60 (1991) 315

[28] D. B. Batchelor, E. F. Jaeger and H. Weitzner, in *Theory of Fusion Plasmas* (Proc. Int. Workshop, Varenna, 1988) Editrice Compositori, Bologna (1988) 691

[29] D. R. Lynch, K. D. Paulsen, IEEE Trans. Microwave Theory Tech. 39 (1991) 383

[30] S. Brunner, J. Vaclavik, Phys. Fluids B 5 (1993) 1695

[31] M. J. Grote, T. Huckle, SIAM J. Sci. Comput. 18 (1997) 838

(E-mail address of A. Jaun: jaun@fusion.kth.se)

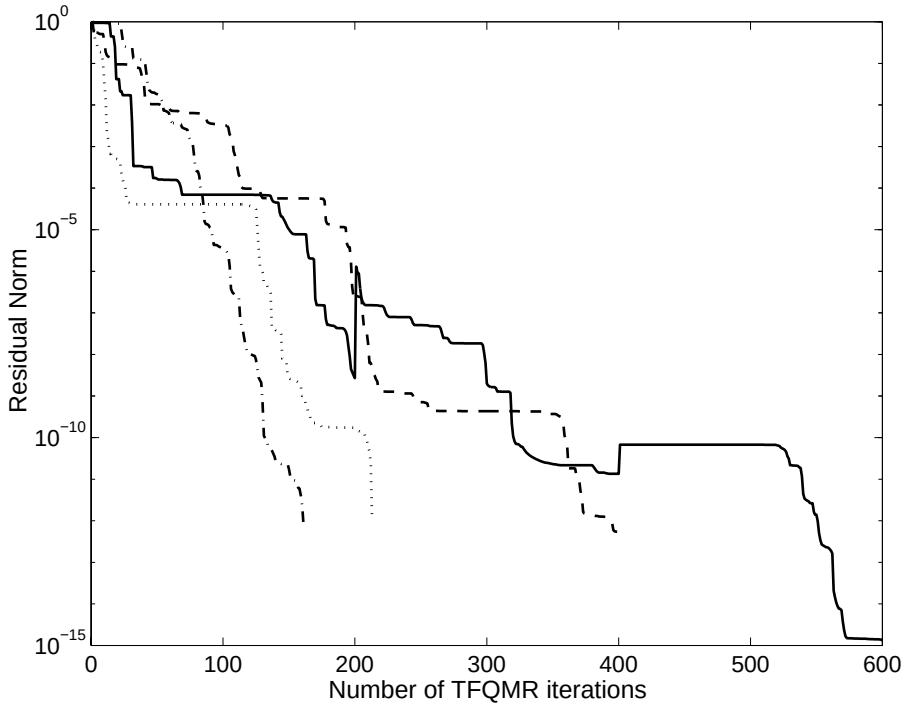

Figure 1: The norm of the residual $\|\mathbf{r_k}\|_2/\|\mathbf{r_0}\|_2$ is plotted as a function of the iteration number for each of the largest test cases computed in this paper, using the TFQMR(n) algorithm restarted after $n = 50$ iterations: linear (dots) and cubic (dashes) FEM in a cylindrical cavity, Alfvén mode-conversion (line) and radio-frequency heating (dash-dots) scenarios both with cubic FEM.

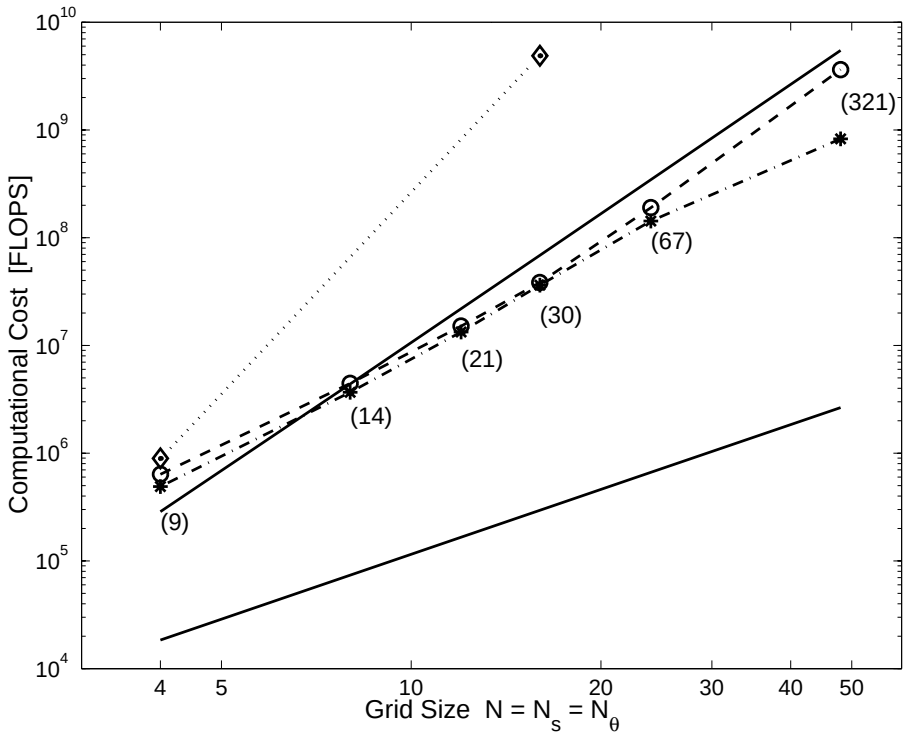

Figure 2: Computational cost for a vacuum wavefield problem computed with bilinear finite elements. The total floating point operations (FLOPS) is plotted as a function of the mesh size $N = N_s = N_\theta$ in a comparison between direct LU factorization (upper continuous line), ILU(0) + GMRES without restart (circles), TFQMR iterations with ILU(0) (stars) and without preconditioning (diamonds). The number of TFQMR iterations is shown (in parenthesis) with the cost of a single matrix-vector multiplication in sparse format below (lower continuous line).

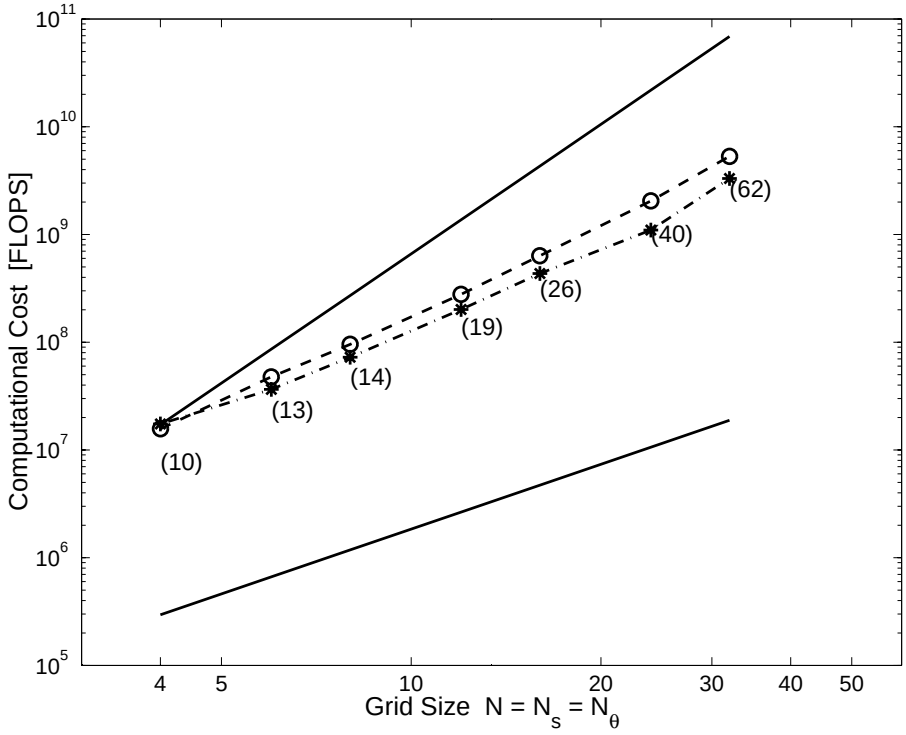

Figure 3: Same notation as Fig.2 using bicubic-Hermite finite elements.

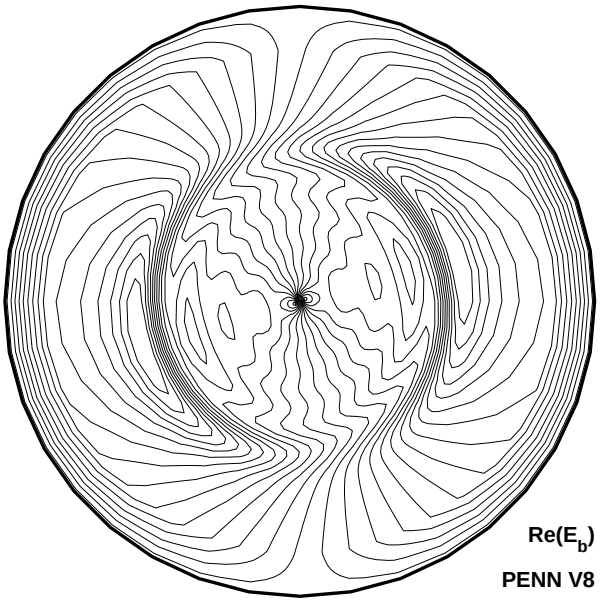

Figure 4: Alfvén wave mode-conversion scenario computed with a gyrokinetic plasma model for $(\vec{A}, \phi)$. The contour lines suggest how an evanescent field driven in the vacuum region tunnels into the plasma to a mode-conversion layer. A kinetic-Alfvén wave escapes and gets Landau damped as it propagates further into the plasma.

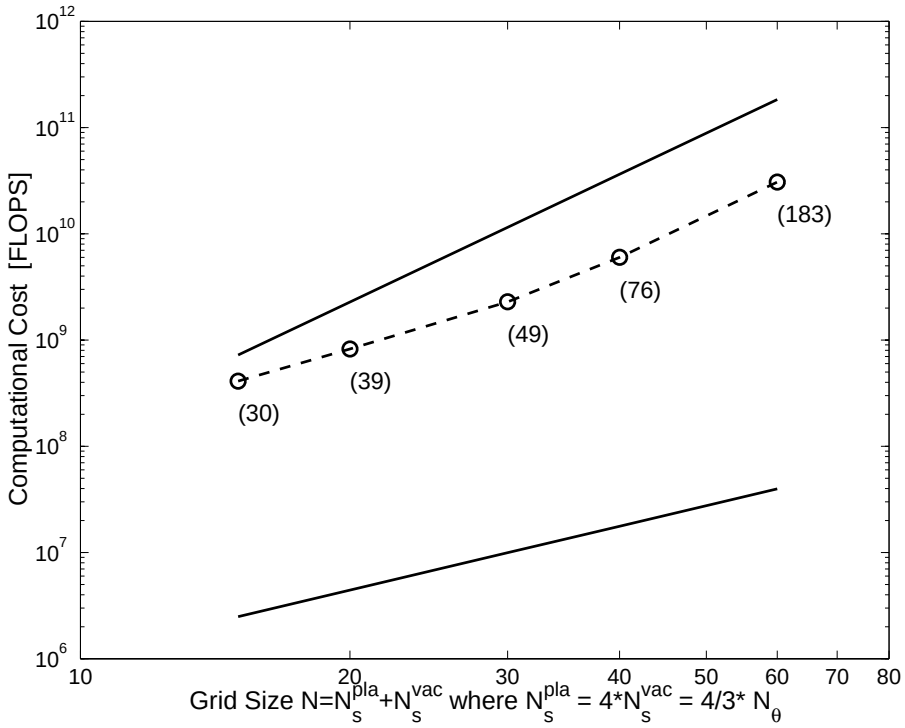

Figure 5: Computational cost to solve the Alfvén mode-conversion scenario of Fig.4 using bicubic-Hermite finite elements. The total floating point operations (FLOPS) for standard LU factorization (upper line) is compared with an iterative calculation combining ILU(0) + TFQMR (circles). The number of iterations is given in parenthesis with the cost of a single matrix-vector multiplication in sparse format below (lower line).

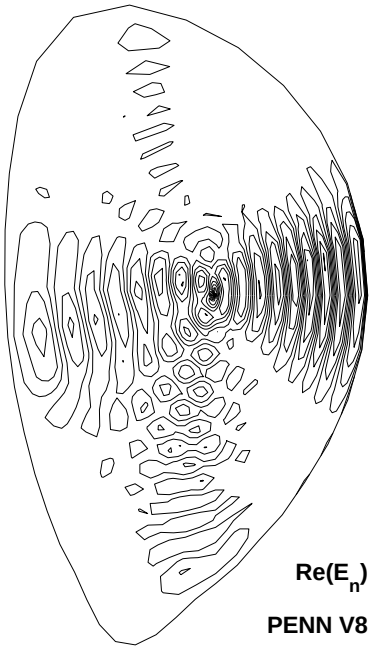

Figure 6: Radio-frequency heating calculation in a JET dicharge (42677) that produced 12 MW of fusion power. A fast magnetosonic wave launched from an antenna and propagates to the ion-cyclotron resonance of a minority species where it is absorbed.