CHAPTER 1

EFFICIENT AND ACCURATE BOUNDARY METHODS FOR COMPUTATIONAL OPTICS

Ch. Hafner and J. Smajic

Computational Optics Group, IFH, ETH Zurich
Gloriastrasse 35, ETH Zentrum, ETZ K81
CH-8092, Zurich, Switzerland
E-mail:hafner@ifh.ee.ethz.ch

The Multiple Multipole Program (MMP) is an advanced boundary method for computational electromagnetics and optics. Its closeness to analytical methods allows one to obtain highly accurate and reliable results. After a short introduction of the fundamentals of MMP, additional techniques are introduced that allow one to drastically improve the performance for various applications. This includes 1) a special consideration of ill-conditioned matrices, 2) the so-called connection concept, 3) an advanced eigenvalue solver with a special eigenvalue tracing procedure, and 4) the Parameter Estimation Technique (PET). Special attention is paid to the error estimation and validation of the results. In a second section, the most recent MMP version of the MaX-1 software with its advanced modeling, visualization, and animation features is presented and it is demonstrated how this software is used for handling complicated projects. Thereafter, some typical applications of computational optics with a focus on photonic bandgap computations and photonic crystal structures are presented.

1. The Multiple Multipole Program (MMP) concept for electrodynamics

The Multiple Multipole Program (MMP) concept was developed since 1980 [1] at ETH Zurich as an extension of the traditional multipole

approximation that was used by Mie in 1900 for the computation of electromagnetic waves along circular wires [2] and in 1908 [3] for the scattering at a spherical particle. MMP is therefore very close to analytical solutions and is sometimes called a semi-analytical approach. In fact, MMP approximates the field in any domain by a series expansion

$$Field(P) = \sum_{k=1}^{K} A_k Basis_k(P) + Error(P) \tag{1}$$

as many other matrix methods, for example, the Method of Moments (MoM). The basis functions $Basis_k$ used in MMP are always analytical solutions of the Maxwell equations within at least one domain. The unknown, linear Parameters A_k, i.e., the amplitudes of the basis functions, are determined in such a way that the weighted square errors of the boundary and continuity conditions along all boundaries of all domains are minimized. Therefore, MMP is a true boundary method. In contrast to the Boundary Element Method (BEM), MMP does usually not approximate the boundary by a polygon or by boundary elements, i.e., it works directly on the original boundaries. Therefore, an important source of inaccuracies is omitted. MMP is also closely related to the MoM, although the MoM is traditionally formulated as a domain method similar to the Finite Elements Method (FEM). Since the MoM usually approximates the sources (currents and charges) of the electromagnetic field, it also "works" on the boundaries of the field domains when only perfect conductors and loss-free dielectrics are present. Therefore, MoM codes discretize the boundaries only when no losses are present. In this case, MoM codes usually approximate the boundaries like BEM codes, which causes inaccuracies especially for the near-field computation. The MMP approach allows one to precisely compute the electromagnetic field everywhere, even on the boundaries. This is especially important when the near-field effects must be taken into account, for example, in Scanning Nearfield Optical Microscopy (SNOM).

1.1. *Remarks on completeness and convergence*

In 1948, Vekua provided a strong mathematical basis for multipole approximations of 2D problems [4]: A series expansion of the form

$$F^{approx}(P) = \sum_{k=0}^{K_0}\left[A_{0k}J_k(\kappa r_0)\cos(k\varphi_0) + B_{0k}J_k(\kappa r_0)\sin(k\varphi_0)\right] + \\ + \sum_{i=1}^{I}\left(\sum_{k=0}^{K_i}\left[A_{ik}H_k^{(1)}(\kappa r_i)\cos(k\varphi_i) + B_{ik}H_k^{(1)}(\kappa r_i)\sin(k\varphi_i)\right]\right) \quad (2)$$

may be used to approximate any solution of the 2D Helmholtz equation

$$\left(\Delta + \kappa^2\right)F = 0 \quad (3)$$

(where κ is the 2D wave number) within a domain D with an error that is below an arbitrarily small positive number e everywhere in D – provided that the maximum orders K_i of all multipoles (terms containing the Hankel functions H) and the maximum order K_0 of the Bessel expansion (terms containing the Bessel functions J) are sufficiently big and provided that the linear parameters or amplitudes A_{ik}, B_{ik}, A_{0k}, B_{0k} have appropriate values. This means that we can write

$$\left|F^{approx}(P) - F^{correct}(P)\right| < e\,; \quad \forall P \in D; \quad e > 0 \quad (4)$$

Note that D is a multiply connected domain with I holes (with boundaries ∂D^i) and with an exterior boundary ∂D^0. Inside each of the holes, one must place a set of multipole expansion with its polar coordinate system (r^i, φ^i). When D is simply connected, i.e., no holes are present, the multipole expansions are missing. When the exterior boundary ∂D^0 is missing, the Bessel expansion set with its polar coordinate system (r^0, φ^0) is omitted. Furthermore, one sets all amplitudes $B_{i0}=B_{00}=0$ because $\sin(0)=0$. Note that the location of each field point P inside D must be expressed in $I+1$ different polar coordinate systems (see Fig. 1).

Vekua's proof uses general conforming mapping transformations. Therefore, it holds for very general, multiply connected domains D with Hoelder-continuous boundaries.

Although Vekua's work was translated in 1967 [5], there was a long, confusing debate about the completeness of the series expansion (2) because many authors encountered numerical problems when they used (2) for writing electromagnetic codes. This debate was misleading because most of the authors - who were not aware of Vekua's work - guessed that (2) is incomplete. Finally, methods based on (2) were abandoned in the USA and replaced by MoM and similar codes. In

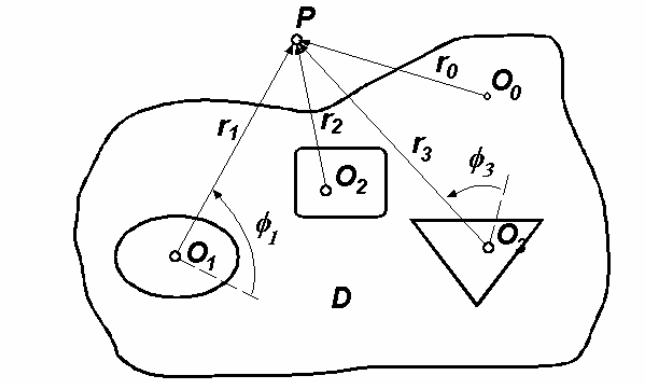

Fig. 1. Domain D with three holes and 3+1 different polar coordinate systems for the Vekua expansion (2). The origins O_1, O_2, O_3 must be inside the three holes. Otherwise, the locations and orientations (dashed lines) of the coordinate systems are free.

Europe, Japan, and other countries, several methods based on (2) were independently developed by several groups under different names [6].

It is interesting to observe that many scientists who develop numerical methods overestimate the value of mathematical proofs of completeness and convergence. This shall be outlined in the following.

The standard procedure of those who design numerical methods based on series expansions, is to first verify that the expansion is complete. If this is guaranteed, they try to prove the convergence, which is even more difficult. If convergence is guaranteed, they consider the convergence rate and assume that the method will be excellent, when a fast convergence, namely when exponential convergence is obtained.

According to Vekua, the standard multipole expansion (2) is complete in the sense that it uses a complete approximation basis for domains with Hoelder-continuous boundaries. The proof of the convergence is more difficult. Fortunately, the situation is similar to the well-known Fourier approximation of a function $f(t)$ that is defined on a limited time interval. It is well known that exponential convergence is obtained when the function $f(t)$ is infinitely many times differentiable. When the nth derivative is discontinuous in a finite number of points, the convergence of the Fourier series is no longer exponential, but still nth order convergence is obtained. Essentially the same holds for the convergence of (2): When all boundaries are infinitely many times differentiable, one

obtains an exponential convergence. This holds for many of the prominent analytic test cases, namely for circular and elliptic boundaries (see Fig.). For technical applications, one often works with boundaries having some corners, where the field may be singular and where the first derivative of the boundary is not continuous. Therefore, the convergence is much worse. From a more detailed analysis, one can see that a rather unbalanced error distribution is obtained along the boundaries when corners are present (see Fig. 2). This is no surprise: When the field is singular in some point, 100% error must be obtained there with any numerical method that uses finite numbers.

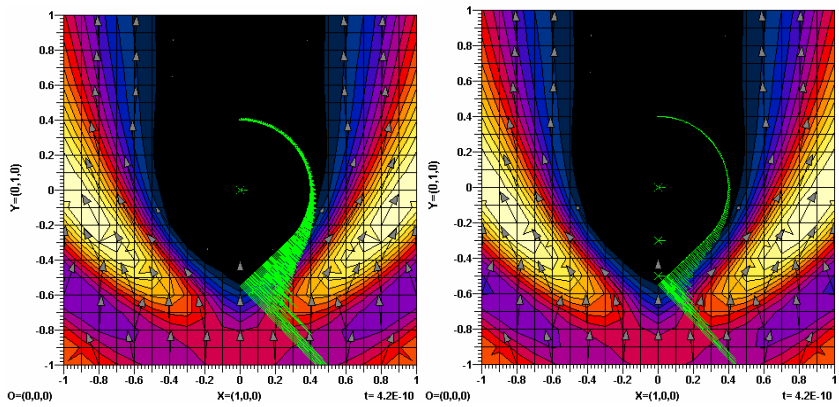

Fig. 2. 2D scattering problem. Plane wave incident in y direction from the bottom on a PEC cylinder with a sharp corner, E field parallel to the z axis Error distribution along the right half of the boundary (green lines) and time average of the Poynting vector field. Left hand side: The scattered field is modeled with a single multipole expansion at x=0, y=0, maximum order 100. The maximum error is 100%, the average error along the boundary is quite big (15%) despite of the high order multipole. Higher orders cause numerical under- and overflow problems. Right hand side: MMP expansion with 3 multipoles (green crosses) of maximum order 10. As a consequence, the number of unknowns is reduced by a factor 3, the matrix size is reduced by the factor 3^2, and the computation time is reduced by the factor 3^3. The maximum error near the corner is still 100%, but the average error is reduced by a factor 2.

We can take advantage of the dependence of the convergence on the geometric shape of the boundary by introducing smoothing procedures. For example, when we round all corners by inserting arbitrarily small circles, the convergence is of order 2 rather than order 1. With a cubic spline approximation, we even can obtain order 3. Theoretically, we can even approximate a boundary that is given in a finite number of points by

a smooth, infinitely many times differentiable boundary that leads to exponential convergence. Although it is nice to apply smoothing procedures for improving the convergence, one can also see that a practical numerical model with very small circles near the corners will lead to almost identical results as the model without smoothened corners. In fact, even when exponential convergence of a series expansion is guaranteed, this does not guarantee that a useful numerical method is obtained, because it may happen that a huge number of basis functions are used for obtaining reasonable results. Obviously, a divergent series expansion that allows one to approximate the solution with, for example, 10 terms having an accuracy of 0.00001% is much better than a series expansion with exponential convergence that requires, for example, 10^{10} terms for 10% accuracy. Therefore, our goal is to find series expansions that allow us to approximate our solution with a given accuracy having a minimum number of terms or – at least – with a reasonably small number of terms.

To illustrate this, let us analyze a function f(t) defined on the interval $0<t<1$. When we have, for example,

$$f(t) = \cos(\omega t) \tag{5}$$

we obtain exponential convergence of the Fourier approximation

$$f^0(t) = \sum_{k=0}^{K} A_k \cos(k\,\omega_0\, t) \tag{6}$$

When the frequency ω is equal to $n\omega_0$, we have 100% error as long as $K<n$ and 0% error otherwise. Note that this is even better than exponential convergence. However, this approach is not reasonable for big n. In order to overcome such problems, one can generalize the Fourier series approach by admitting arbitrary frequencies ω_k, instead of multiples of ω_0. This leads immediately to non-orthogonal series expansions that require the solution of a matrix equation for obtaining the linear parameters or amplitudes A_k of the series expansion. When the frequencies ω_k are within a narrow band, one obtains an ill-conditioned matrix, but one can obtain excellent approximations – even when ω of the given function is not within this band – when a matrix solver is used that can handle ill-conditioned matrices. We will encounter similar effects when we consider the condition number of the MMP matrix.

Generalized Fourier series essentially contain two sets of parameters, namely the amplitudes and the frequencies. The former are linear

parameters, whereas the latter are non-linear. The numerical optimization of the non-linear parameters is much more time consuming than the optimization of the linear ones, especially when the number K of terms is high. But with the generalized Fourier approach we can obtain good results even when we only have a rough guess of the frequencies. Furthermore, we can generalize the expansion by admitting arbitrary basis functions

$$f^0(t) = \sum_{k=0}^{K} A_k \, basis_k(p_k, t) \tag{7}$$

where p_k may be a single non-linear parameter or a set of such parameters. This general series expansion - with an arbitrary set of basis functions that are generated with some advanced Genetic Programming (GP) code - was implemented in the GGP code [7]. Although the GGP approximation of an arbitrary function is much more time consuming than the approximation with standard series (Fourier, power series, etc.), it has an important advantage: GGP approximations often provide excellent extrapolations outside the interval where $f(t)$ is given, whereas rather poor extrapolations are obtained with standard series. This plays an important role for advanced methods that speed up numerical procedures. Typical examples are the Parameter Estimation Technique (PET) and the Eigenvalue Estimation Technique (EET) that will be outlined below.

The generalized series expansion (7) is useful for obtaining good extrapolations, but it is also the prototype of the Generalized Multipole Method (GMT) [6]. Here, one can simply write (1) or

$$Field^{approx}(\vec{r}, t) = \sum_{k=1}^{K} A_k \, Basis_k(p_k, \vec{r}, t) \tag{8}$$

Since GMT codes usually are frequency domain codes, we use complex notation and write

$$Field^{approx}(\vec{r}, \omega) = \sum_{k=0}^{K} A_k \, Basis_k(p_k, \vec{r}, \omega) \tag{9}$$

In electrostatics, we simply can express the scalar potential V as:

$$V^{approx}(\vec{r}) = \sum_{k=0}^{K} A_k \, Basis_k(p_k, \vec{r}) \tag{9'}$$

Obviously, Vekua's standard 2D multipole set (2) has essentially the form (9). In addition to the amplitudes, i.e., the linear parameters A and B, one has nonlinear parameters that describe the coordinate transforms required for the different systems of polar coordinates. Theoretically, it is possible to optimize the non-linear parameters, i.e., the locations of the polar coordinate systems, but this is extremely time consuming. Furthermore, it turns out that severe numerical problems may occur even when the multipoles are optimally placed and even when exponential convergence is guaranteed, i.e., the number K of basis functions may be very high for obtaining reasonable accuracy. Furthermore, the condition number of the system matrix may be high.

In order to reduce the number K of linear parameters, we can generalize the series expansion (2) by 1) placing additional multipoles in each hole, 2) placing multipoles outside the exterior boundary, 3) introducing other types of expansions that analytically fulfill the Helmholtz equation or the Maxwell equations in general. Placing additional multipoles is the main idea of the Multiple MultiPole (MMP) approach [6]. The MMP solver of the MaX-1 code [8] also contains additional expansions, namely, plane waves, harmonic expansions, Rayleigh expansions, Gaussian beams, and distributed multipole expansions that help to reduce the number of basis functions in special situations where a pure multipole approach turns out to be inefficient.

It is important to note that the generalization of the Vekua expansion also creates new problems. The main problem is that the condition number of the resulting system matrix can become very high. In order to overcome these problems, several simple rules for the placement of the multipoles were derived. Furthermore, the procedure for obtaining the MMP system matrix was carefully analyzed (see below). This led to the Generalized Point Matching (GPM) technique that is similar to the method of weighted residuals and similar to the Galerkin method. Since the GPM uses direct solvers for the resulting overdetermined system of equations, it reduces problems obtained from ill-conditioned matrices. Consequently, the combination of the MMP expansion with GPM allows one to obtain very accurate results with relatively small matrices with typically less than 2000 columns.

It has been mentioned that Mie used the multipole approach not only for solving 2D problems, but also for 3D scattering. In fact, the MMP approach can easily be extended to arbitrary 3D problems of elecrodynamics. It seems that no proof of completeness has been found

up to now, although the Mie method seems to work very well. However, since MMP uses additional multipole expansions, the MMP set is always complete when sufficiently many multipoles are used. As mentioned before, the problem is neither completeness nor convergence: The problem is to find a set of basis functions that allow one to obtain the desired accuracy with a sufficiently small number of unknown linear parameters. However, for reasons of simplicity, we will focus on 2D applications, i.e., cylindrical structures in the following sections.

1.2. From simple multipole solutions to efficient MMP solutions

In order to illustrate the main MMP concepts, we first consider the simple electrostatic case of the electric field and potential computation of two symmetric, circular, cylindrical wires. Note that the dynamic case was considered by Mie in 1900. The static case is considerably simpler because the 2D Helmholtz equation is replaced by the simpler 2D Laplace equation $\Delta V=0$ that holds in the 2D cross section of the cylindrical structure, where V denotes the scalar potential. The potential must be computed outside the conductors, i.e., in an unbounded domain D with two circular holes (see Fig. 3). Inside the holes (conductors), the scalar potential V is constant. Let us assume that it is equal to +1 in the right wire and –1 in the left wire.

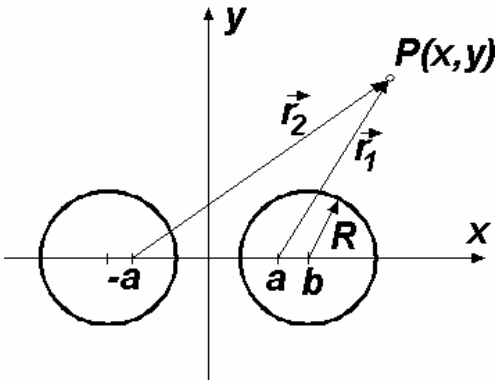

Fig. 3. Cross section of two symmetric, circular wires.

The static counterpart to Vekua's expansion (2) is

$$V^{approx}(P) = A_{01} + \sum_{k=1}^{K_0}\left[A_{0k}\, r_0^{\,k}\cos(k\varphi_0) + B_{0k}\, r_0^{\,k}\sin(k\varphi_0)\right] + \\ + \sum_{i=1}^{I}\left(A_{i0}\ln(r_i) + \sum_{k=1}^{K_i}\left[A_{ik}\, r_i^{-k}\cos(k\varphi_i) + B_{ik}\, r_i^{-k}\sin(k\varphi_i)\right]\right) \tag{10}$$

Note that the radial functions are much simpler, but the form is essentially unchanged. According to an old theorem of Runge, equation (10) is complete in the same sense as equation (2). For our unbounded, symmetric domain D, we can simplify equation (10) considerably because we know that V is anti-symmetric with respect to the y axis and symmetric with respect to the x axis. Therefore, we set one multipole expansion in the right conductor with origin at O_1 (Cartesian coordinates $(a,0)$) and a second multipole at O_2 (Cartesian coordinates $(-a,0)$) and obtain

$$V^{approx}(P) = \sum_{i=1}^{2}\left(A_{i0}\ln(r_i) + \sum_{k=1}^{K_i}\left[A_{ik}\, r_i^{-k}\cos(k\varphi_i)\right]\right) \tag{10'}$$

instead of equation (10). Note that we can omit the sin terms because of the symmetry, when we orient the polar coordinate systems in O_1 and O_2 in such a way that φ_i=0 describes the x axis.

The main problem now is to find a reasonable or even the optimal location of the origins of the multipoles. At first sight, the centers of the conductors seem to be reasonable places for the origins of the multipoles. As long as the distance between the two conductors is not small, one obtains reasonable results with quite small maximum orders (see Fig. 4).

A more detailed analysis shows that the orders required for a given accuracy increase rapidly when the conductors come close together. Therefore, we look for a better solution. Because of the simplicity and because of the symmetry of the problem, we only need to define a single parameter a, i.e., the distance of the multipoles from the y axis. Since the analytical solution of the problem is known, it is easy to find the optimum value of a, i.e., the optimum location of the multipoles. When the centers of the two wires with radius R are at $(b,0)$ and $(-b,0)$, we obtain

$$a^{opt} = \sqrt{b^2 - R^2} \tag{11}$$

When we set the multipoles at $(a^{opt},0)$ and $(-a^{opt},0)$, we obtain the analytical solution

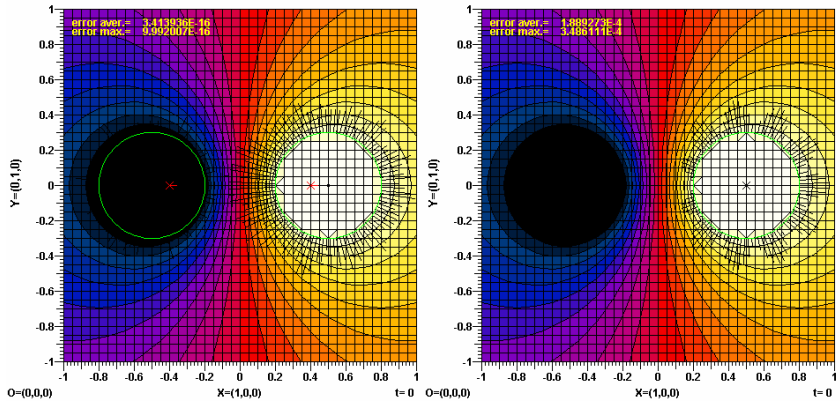

Fig. 4. Potential obtained with a single multipole per conductor. Right hand side: multipole in the center of the conductor, left hand side: zero order multipole at optimal location (origin of the bi-polar coordinate system).

$$V(x,y) = A_{10}\ln(\sqrt{(x-a^{opt})^2+y^2}) + A_{20}\ln(\sqrt{(x+a^{opt})^2+y^2}) \quad (12)$$

already for zero order multipoles, i.e., K_1=K_2=0. Because of the symmetry, we have A=A_{10}=-A_{20}. From the boundary condition at any point, for example, $V(b+R,0)$=1, we finally obtain A. Instead of considering the analytical solution, we may also find the optimal placement from numerical studies (see Fig. 5).

The analytical solution can be interpreted as follows: Outside the conductors we "see" a charge, i.e., zero order multipole or monopole, in each conductor at the virtual positions (a^{opt},0) and (-a^{opt},0). Since the conductors act as mirrors for the electric field, one also calls the monopoles mirror charges or image charges. Unfortunately, the method of mirror charges or image charges does not provide analytical solutions for more complicated configurations. As we know from optical mirrors, a point source can be seen at various positions within a curved mirror or it is even smeared over some area. In 2D electrostatics, we can find the locations of the mirror charges for an arbitrary configuration of two electrodes when we know a conformal transformation that maps the surfaces of the electrodes on two parallel lines. Assume that w=u+iv=C(z) maps the original points z=x+iy in such a way that the surfaces of the two electrodes are mapped on the straight lines at u=-1 and u=+1 (see Fig. 6). When these lines are on the potential –1 and +1

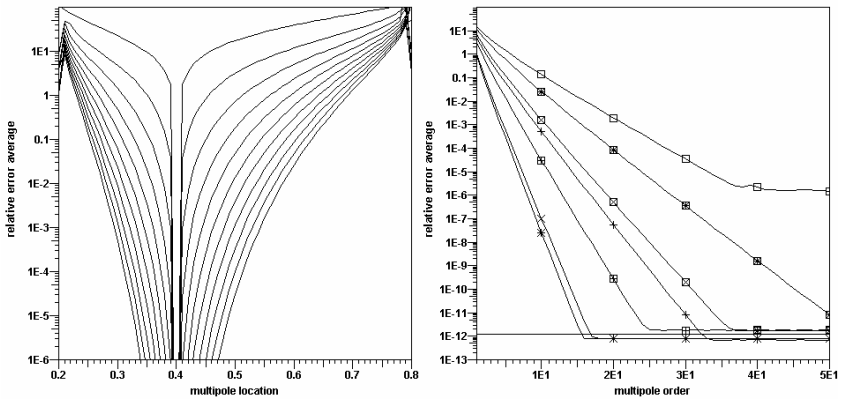

Fig. 5. Average of the relative errors on the boundary for the problem shown in Fig.4. Model with a single multipole inside each conductor at *(a,0)* and *(-a,0)* respectively. Left hand side: Dependence on the multipole location a for the maximum multipole *orders 0, 1, 2,...,10*. Note that the error estimation is inaccurate when the multipole is very close to the boundary, i.e., when a is close to *0.2* or close to *0.8*. All curves exhibit a very deep and sharp minimum at the optimal location *a=0.4*. Right hand side: Convergence for different multipole locations. Curve without marker: optimal placement *a=0.4*; □ marker: *a=0.25*; + marker: *a=0.3*; x marker: *a=0.35*; + and x marker: *a=0.45*; □ and + marker: *a=0.5*; □ and x marker: *a=0.55*; □, +, and x marker: *a=0.6*. Note that one has exponential convergence until machine precision (*14* digits, i.e., 10^{-12}% error) is reached, with one exception: For *a=0.25*, the error no longer decreases when the maximum order is bigger than 35 because the matching point density is not high enough in this model.

respectively, the potential in the w plane is simply $V(u,v)=u$ for $-1<u<1$.

Obviously, we simply have a perfect parallel plate capacitor. We know that the field inside the capacitor is generated by a uniform charge distribution on the electrodes. We can easily extend the field outside the interval $-1<u<1$, by simply saying that we have $V(u,v)=u$ for $-d<u<d$ and that this field is obtained form a fictitious charge distribution along $u=-d$ and $u=+d$. When we approximate the uniform charge distributions along the auxiliary lines $u=-d$ and $u=+d$ by a set of monopoles that are uniformly distributed along the auxiliary lines, we obtain an approximation that is better the bigger the distance d is. In the limit where d approaches infinity, it is even sufficient to use two single monopoles at $u=-\infty$ and $u=\infty$. When we now map the locations of the two monopoles back to the original z plane using the inverse transform $z=C^{-1}(w)$, we obtain the locations of the mirror charges in our simple case

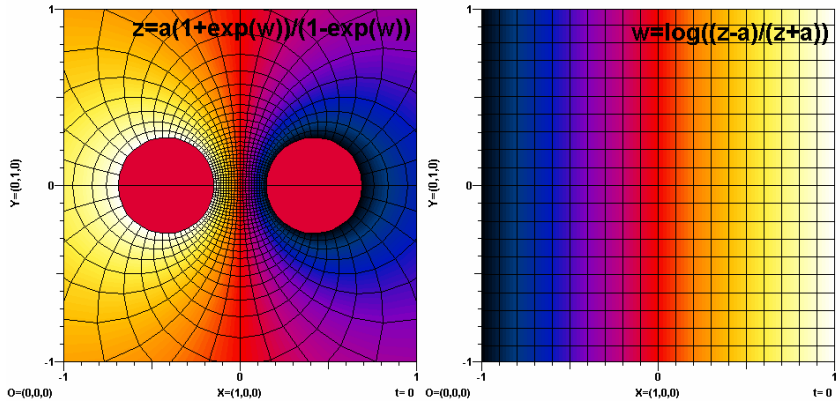

Fig. 6. Conformal mapping of two symmetric circular wires (z plane) on two parallel plates (w plane).

of two circular electrodes. Unfortunately, finding a conformal transformation is not easy in most cases. Furthermore, one often encounters problems with the inverse transform that causes the "smearing" effects mentioned above. However, when we map the auxiliary lines $u=-d$ and $u=+d$ back into the z plane, we obtain auxiliary lines along the surfaces of the electrodes. These auxiliary lines are entirely inside the electrodes. In general, the distance of each auxiliary line from the surface is not constant. Even if we do not know the precise location of the auxiliary line, those who have much experience in conformal mapping can manually draw lines that are a good guess. For a numerical method, this is sufficient. Once the auxiliary lines are defined, one simply can distribute a finite number of auxiliary charges with unknown amplitudes along these lines. This is nothing more than a multiple monopole expansion, i.e., a special case of the MMP approach of the form of equation (10) with $K_i=0$ and I equal to the number of auxiliary charges. Charges are the sources of the electrostatic field. Therefore, this technique is also called Method of Auxiliary Sources (MAS) [9]. Note that the MAS can be extended to dynamic problems as well as to 3D problems although no conformal mapping is known in this case. The MAS is also called Method of Fictitious Sources (MFS) [10].

For computing the I unknown amplitudes of the amplitudes of the auxiliary sources, at least I equations are required. These equations are usually obtained from a Simple Point Matching (SPM) technique, where the boundary conditions in a finite set of matching points are used. For

our simple 2D electrostatic case, we can easily construct one matching point (x_P,y_P) on the boundary for each source point (x_S, y_S) in such a manner that the distance between these two points is minimized. When the transformation $C^{-1}(w)$ is known, one can use it for obtaining an even better set of matching points. In fact, the correlation between the matching points and the locations plays an important role for the accuracy.

From the conformal mapping, we know that it is best to move the auxiliary sources as far away from the surface as possible. If we increase the distance d in the w plane in our simple example, the auxiliary sources in each electrode tend toward the optimal points $(a^{opt},0)$ and $(-a^{opt},0)$. At the same time, the distances between neighboring sources is decreased and tends to zero. As a consequence, the condition number of the MAS matrix tends to infinity. When the condition number is large enough, the positions of the matching points and the quality of the matrix solver become critical. Therefore, one usually moves the auxiliary sources not far away from the surface. Typically, the distance between neighboring sources is in the same order of magnitudes as the distance of the sources from the boundary. However, when one uses a matrix solver that can handle ill-conditioned matrices, one can increase the distances and obtain more accurate results with the same numerical effort. The main road toward a better matrix solver is the Generalized Point Matching (GPM)

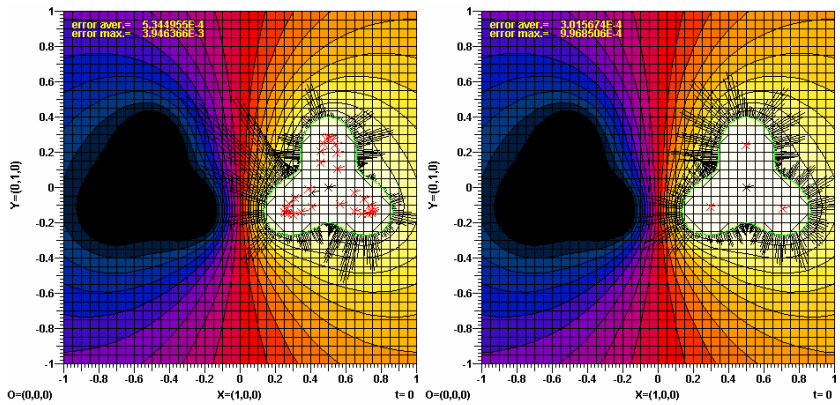

Fig. 7. Two non-circular wires. Error distributions along the surface indicated by black lines. Left hand side: MAS approximation with auxiliary sources relatively far away from the boundary (red crosses). Right hand side: MMP expansion with four multipoles (red and black crosses). 30 unknowns for both computations.

technique that will be outlined below.

When one obtains clusters of monopoles when the distance d in the w plane is sufficiently long, it is reasonable to replace each cluster by a higher order multipole in the center of the cluster. This is similar to what is done in the Fast Multipole Method (FMM) [11]. In the simple example above, we obtain one cluster per wire, i.e., one multipole expansion per wire as in the standard Vekua-Runge approach. When we do not use conformal mapping and "guess" that concentric auxiliary lines are good enough, we finally set the multipoles at $(b,0)$ and $(b,0)$, whereas we find the optimal locations $(a^{opt},0)$ and $(-a^{opt},0)$ when we use conformal mapping for the construction of the auxiliary lines.

For more complicated geometry, we often obtain more than one cluster of auxiliary sources in each electrode. Fig. 7 illustrates this. We then immediately obtain more than one multipole per hole in the domain D, i.e., a Multiple MultiPole (MMP) expansion.

Since we need quite sophisticated conformal mapping procedures for finding the clusters of auxiliary sources, it is quite demanding to find the optimal locations of the multipoles. Fortunately, one also obtains excellent results with sub-optimal sets of multipoles or with automatically generated sets [6, 8]. Fig. 7 illustrates this. In fact, one has a very high degree of freedom in setting the multipoles. When one starts with a given set of multipoles that is either constructed with some automatic procedure [6] or from simple geometric rules (distances between neighbor multipoles similar to the distances from the boundary), one usually obtains reasonable results. One then can analyze the errors along the boundaries and use this information to improve the accuracy if required. This is simply done by increasing the maximum orders of the multipoles near the areas of large errors or by inserting additional multipoles in these areas, i.e., adaptive MMP modeling is almost trivial.

Although it seems to be obvious that the MMP approach is more general and better than the MAS approach, MMP has an apparent difficulty: MMP does not allow one to associate simply one matching point on the boundary for each multipole. Furthermore, the SPM technique fails for complicated geometry, because the distribution of the matching points may be very critical. We therefore need a more general method that is much more robust and accurate.

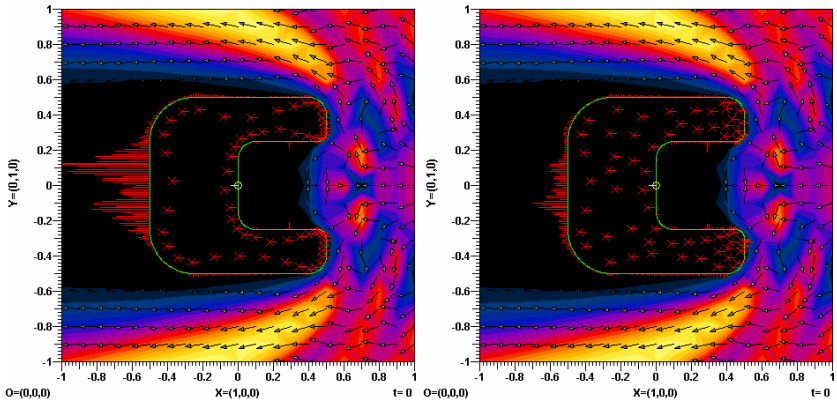

Fig. 8. Time average of the Poynting vector field and error distribution (red lines) along a PEC boundary illuminated by a plane wave. Left hand side: multipole set generated automatically, average relative error 5.3%, condition number of the MMP matrix 4.3E3. Right hand side: "randomized" multipole set, average relative error 2.7%, condition number of the MMP matrix 4.9E7. Note that large errors are obtained in the shadow area, where the field is almost zero. Note also that the model with the higher condition number is more accurate although the corresponding multipole set seems to be less reasonable.

1.3. Generalized Point Matching (GPM) technique

Although the point matching technique may be applied to quite arbitrary problems, we focus on the simple 2D electrostatic Dirichlet problem for reasons of simplicity. Assume that the Laplace equation holds in a multiply connected domain D and that the potential V is known along the boundary ∂D of D. Note that ∂D consists of several pieces when D is multiply connected.

When V is approximated by a series expansion of the form (9'), we can find the unknown linear parameters A_k from the boundary conditions. When we apply the SPM technique, we simply write down the boundary condition(s) in a finite set of matching points P_i :

$$V^{approx}(P_i) = \sum_{k=1}^{K} A_k \, Basis_k(p_k, P_i) = V^{boundary}(P_i); \quad i = 1,2,...,I = K \quad (13)$$

When the boundary is parametrized and s is the parameter, we can also express

$$V^{approx}(s) = \sum_{k=1}^{K} A_k \, Basis_k(p_k, s) = V^{boundary}(s) \tag{13s}$$

and obtain (13) by writing down (13s) in I points P_i located at s_i.
Note that the number of unknowns K is equal to the number of matching points I. Therefore, we obtain a set of $N=I$ equations with $N=K$ unknowns. In matrix notation, we have:

$$\mathrm{MA}^{\mathrm{SPM}} = \mathrm{V} \tag{14}$$

Where M is an N by N matrix with the elements

$$m_{ik} = Basis_k(p_k, s_i) \tag{14m}$$

A is the unknown vector with N elements A_k, and V is a known vector with the elements

$$v_i = V^{boundary}(s_i) \tag{14v}$$

It is well known from the MoM literature, that the SPM technique is not optimal, i.e., the vector $\mathrm{A}^{\mathrm{SPM}}$ obtained is often inaccurate. In order to examone the accuracy, we need a reasonable error definition first. When we would know the correct solution, this would be relatively easy, but this is usually not the case, when we apply a numerical method. However, for the simple Dirichlet problem, we know the scalar potential on the boundary. Furthermore, we know that the extreme values of V are always on the boundary and that V gets more and more smooth when we move away from the boundary. Therefore, the error along the boundary is a reasonable measure that can easily be defined, for example:

$$Error^{boundary} = \sqrt{\int_{boundary} (V^{approx} - V^{boundary})^2 \, ds} \tag{15}$$

Note that the error definition is not unique at all. First of all, we have used the square norm that could be replaced by a more general p norm. Furthermore, we could introduce weights. Weighting is very important for more complicated problems, for example, in dynamics when different fields (electric and magnetic) are involved in the error definition. However, for the simple Dirichlet problem, the definition (15) is appropriate.
As soon as we insert the series expansion of equation (13) in equation (15), we obtain $Error^{boundary}$ as a function of the unknown amplitudes A^k. The square of $Error^{boundary}$ is minimized when

$$\frac{\partial\left(Error^{boundary}\right)^2}{\partial A^j} = \frac{\partial \int\limits_{boundary} (\sum_{k=1}^{K} A_k \, Basis_k(p_k,s) - V^{boundary}(s))^2 \, ds}{\partial A^j} = 0 \quad (16)$$

From this square error minimization, we obtain the matrix equation

$$\mathrm{S}\mathrm{A}^{\mathrm{E2}} = \mathrm{V}^{\mathrm{E2}} \quad (17)$$

Where S is a symmetric matrix with the elements

$$s_{ik} = \int\limits_{boundary} Basis_i(p_i,s) \cdot Basis_k(p_k,s)\, ds \quad (17s)$$

Furthermore V^{E2} is a vector with the elements

$$v_i = \int\limits_{boundary} Basis_i(p_i,s) \cdot V^{boundary}(s)\, ds \quad (17v)$$

and A^{E2} is the vector containing the linear parameters A_k. Note that we could obtain the same matrix equation with the projection technique, when we project the boundary condition (13s) on a set of $I=K$ testing functions $t_i(s)$ and use Galerkin's choice of testing functions, $t_i(s)=Basis_i(p_i,s)$.

In simple cases, we can analytically solve the integrals in (17), but for a general numerical code that shall handle arbitrary boundaries, we must approximate the integrals numerically, for example, with a Riemann sum:

$$s_{ik} \cong \sum_{j=1}^{J} Basis_i(p_i,s_j) \cdot Basis_k(p_k,s_j) \cdot ds_j \quad (18s)$$

$$v_i \cong \sum_{j=1}^{J} Basis_i(p_i,s_j) \cdot V^{boundary}(s_j) \cdot ds_j \quad (18v)$$

where ds_j is essentially the distance between neighboring points. When the integration points s_j are distributed uniformly along the boundary, all lengths ds_j are equal to s^{tot}/J, where s^{tot} is the total length of the boundary. It is well known that this is not optimal for obtaining small errors. When we use adaptive integration routines, the integration points will be at different locations for each matrix element. Although this may be done, it is numerically expensive because the basis functions must be evaluated finally at many positions. It is therefore better to select a sufficiently large, sub-optimal set of integration points for all matrix elements. Such a set may be obtained from experience: Since one knows that the field

varies more rapidly when the radius of curvature of the boundary is small, one will increase the density of integration points in such areas. Furthermore, the set of integration points may be improved adaptively. However, reasonable approximation may only be obtained when the number J of integration points is sufficiently large. First of all, $J \geq K$ must hold. In simple cases, $J=2K$ is sufficient. In complicated cases, higher numbers are required.

Since all matrix elements in (18) are scalar products of vectors, we can write for the numerical approximation of equation (17):

$$S^{approx} A^{E2approx} = V^{E2approx}, S^{approx} = R^T R, V^{E2approx} = R^T V \quad (19)$$

Where R is a rectangular matrix with the elements

$$r_{ik} = Basis_k(p_k, s_i) \cdot \sqrt{ds_i} \quad (19r)$$

R^T is the transposed of R, and V is a vector with the elements

$$v_i = V^{boundary}(s_i) \cdot \sqrt{ds_i} \quad (19v)$$

Note that equation (19) solves the overdetermined system

$$\mathrm{RA}^{\mathrm{over}} - \mathrm{V}^{\mathrm{over}} = \mathrm{E}, \quad \mathrm{V}^{\mathrm{over}} = \mathrm{V}^{\mathrm{E2approx}} \quad (20)$$

where E is the error or residual vector, in the least-squares sense, i.e., it minimizes the square norm of E which contains the mismatching values in the integration points s_i. This allows one to compute and plot the error distribution along the boundary.

When we compare equation (14) with equations (19) and (20), we see that equation (20) is very similar to equation (14). It is obtained from equation (14) essentially by two generalizations: 1) inserting weights (the square roots of the distances between neighbor matching points) and 2) using $I>K$ matching points, i.e., an overdetermined system of equations. Therefore, we call equation (20) Generalized Point Matching (GPM) technique.

GPM is very similar to SPM and needs only minor modifications in the implementation. Despite of this, it provides much better results in general, because it essentially minimizes the square error and is numerically equivalent to the Galerkin method.

When the matrix R is not well conditioned, the direct computation of the GPM equations (20) is even more accurate than computation of the symmetric system (19) obtained from the Galerkin method or from the error minimization technique. The reason for this is that $\mathrm{S}^{\mathrm{approx}}$ is

essentially the square of R and this causes severe numerical problems when the condition number is high. Solving equation (20) with a direct matrix solver (for example, the Givens updating procedure) using single precision is as accurate as solving equation (19) with double precision. Remember that the MAS accuracy is increased when the sources are moved away from the boundary and that the condition number of the MAS matrix is increased as well. The same holds for MMP and similar methods. Therefore, replacing equation (19) by the GPM system ofequations (20) allows one to obtain considerably more accurate results without a significant increase of the memory requirement and computation time. We will consider this in the following section.

1.4. The role of ill-conditioned matrices

From conformal mapping and the consideration of the simple case of the two circular wires in the previous sections, it is clear that the accuracy of the MAS results for a fixed number of sources (monopoles) depends on the distance of the sources from the boundary. As we have seen, conformal mapping provides optimal auxiliary lines in some sense. For our simple case, we can use conformal mapping of the entire structure, i.e., of both boundaries and obtain circles as auxiliary lines that are nothing else but coordinate lines of a bipolar coordinate system. Furthermore, conformal mapping provides the optimal distribution of the sources on such a circle, which is non-uniform. Since conformal mapping of the entire structure is useful only in simple cases, one often uses sub-optimal auxiliary lines, for example, at a given distance from the boundary. In our case, such auxiliary lines are concentric circles that would be optimal when only one circular wire would be present or when the two conductors would be co-axial. Figure 9 shows the behavior of the MAS matrix condition number and of the average mismatching error along the boundary for concentric auxiliary lines with different sets of monopoles along them. As one can see, the condition number becomes very large when the distances of the monopoles from the centers of the wires, i.e., the radii of the auxiliary lines become small. This is to be expected. As one can see, the values of the condition number are always below 10^{17}. This is due to the SVDC evaluation of the condition number. When the condition number is above 10^M, where M is the precision of the arithmetic, SVDC cannot accurately compute the condition number and it will return a value in the order of 10^M. For the computations shown

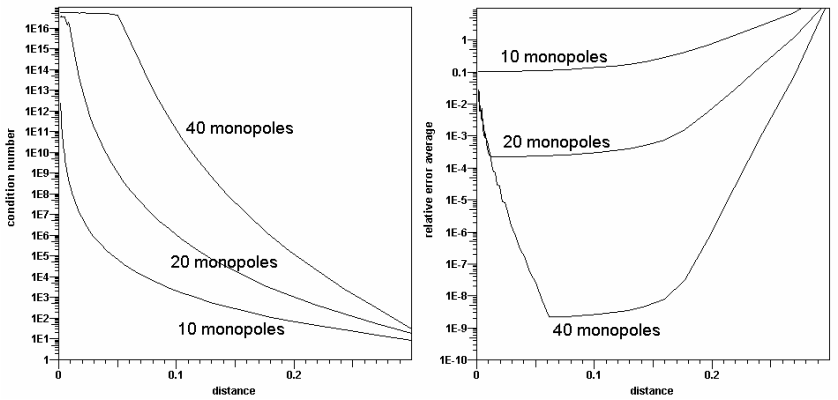

Fig. 9. Dependence of the condition number and of the relative error (in percent) along the boundary for MAS approximations of the problem shown in Fig. 3+4, with different numbers of auxiliary sources (monopoles) uniformly distributed along concentric circles. "distance" indicates the distance of the sources from the center of the wire.

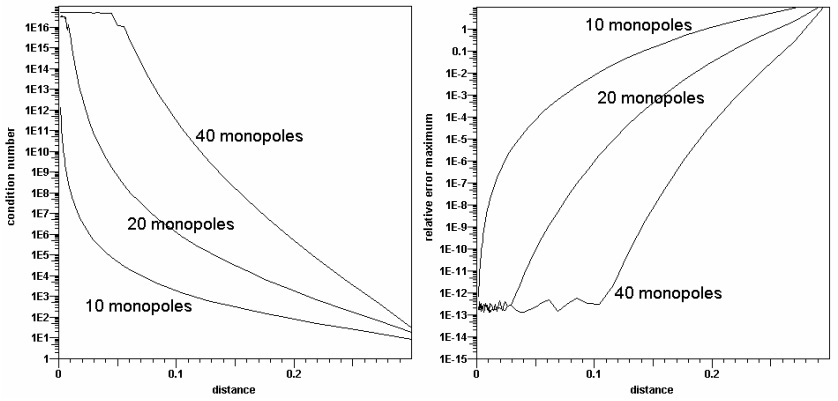

Fig. 10. Same as in Fig. 9, but the sources are moved toward the optimal multipole location at (0.4,0). "distance" indicates the average distance of the sources from the focus point (0.4,0).

here, double precision (M=16) has been used throughout. Essentially the same behavior is observed when we distribute the monopoles along the optimal, bipolar circles (see Figure 10).

When we compare the error distributions of Figures 9 and 10, we see that the error always decreases when the condition number increases, as long as the condition number is below 10^M. In Figure 9, one also observes

some kind of saturation effect for "distance" below 0.18. This effect is due to that the auxiliary circles do not contain the focus point at (0.4,0) when "distance" is too small.
However, it is more important to note that the curves become noisy when the "distance" is so small that the condition number is higher than 10^M. This is because the matrix solver looses accuracy when the condition number is high. As we can see, the standard matrix solver of the MMP implementation of the MaX-1 code is excellent as long as the condition number is below 10M. Even when the condition number is above this limit and cannot be evaluated with SVDC, accurate results with low errors may be obtained. The standard MMP matrix solver is an efficient update procedure based on Givens plane rotations [12]. As we can see, this matrix solver allows us to work with matrices with very large condition numbers and allows us to obtain highly accurate results even when the condition number is above 10^M.
Note that the Givens procedure runs on the overdetermined system (20), but the rectangular matrix R needs not to be stored when the Givens update procedure is applied. In this case, only a triangular matrix must be stored, i.e., the same memory space is required as when one works with the symmetric system of equations (19) obtained from the error minimization or from the Galerkin method. When we solve equation (19), we may apply, for example, the Cholesky algorithm that is a little bit faster than the Givens procedure. As we can see from Figure 11, this algorithm exhibits severe problem as soon as the condition number is higher than $10^{M/2}$. As a consequence, the final error is much higher than when we solve equation (20) with Givens. Furthermore, we have much less freedom in setting the monopoles. For the MAS this is less dramatic than for MMP, because one usually puts the sources not very far away from the boundaries, whereas MMP users often work with multipoles that are far away from the boundary. However, it is very important to see that we may obtain much more accurate results when the condition number is close to 10^M than when we use a method that provides small condition numbers, i.e., minimizing the condition number is not reasonable when one is interested in accurate results. According to our experience, this does not only hold for MMP, MAS, and closely related methods such as the MoM, but also for other methods such as Finite Elements, generalized Fourier series, etc. When symmetric systems of equations are present, this should always be considered as a hint that the symmetric matrix S might be a product of the form $S=R^TR$, where R is a

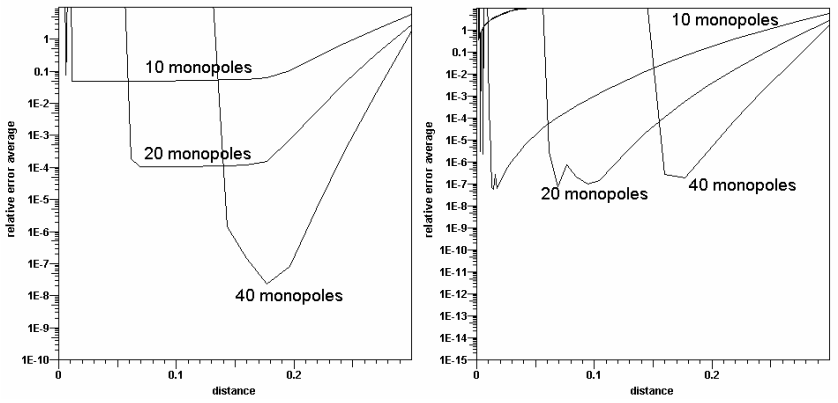

Fig. 11. Relative error average versus distance when the Cholesky matrix solver is applied for solving the system (19) with the symmetric matrix S. Left hand side: same as in Fig. 9, right hand side: same as in Fig. 10.

rectangular matrix. If this is correct, one should try solving the overdetermined system described by R directly.

Unfortunately, there are much fewer numerical algorithms for rectangular matrices than for symmetric ones. Furthermore, direct matrix solvers such as the Givens procedure are very time consuming for big matrices and should be replaced by iterative matrix solvers, whenever possible. A prominent iterative solver for symmetric matrices is the Conjugate Gradient (CG) method. In fact, in the original paper of Hestenes and Stiefel [13] also a CH version for rectangular matrices was presented. For CG, the preconditioning is very important. Unfortunately, almost no information about the preconditioning of rectangular matrices is available. Therefore, we use a simple Jacobi preconditioner that essentially is nothing else but scaling the rectangular matrix R in such a way that the square norm of each column is equal to 1. This leads to the results shown in Figure 12. As one can see, the behavior is considerably better than for the Cholesky solver, but much worse than for the Givens solver. Numerical problems obviously already occur when the condition number is above $10^{M/2}$. Therefore the placement of auxiliary sources and multipoles is much more critical when this solver is applied. One then is usually forced to work with sub-optimal placement for obtaining sufficiently small condition numbers. As a consequence, the errors become larger. When a given problem shall be solved with a certain accuracy, this means that a considerably bigger matrix must be solved

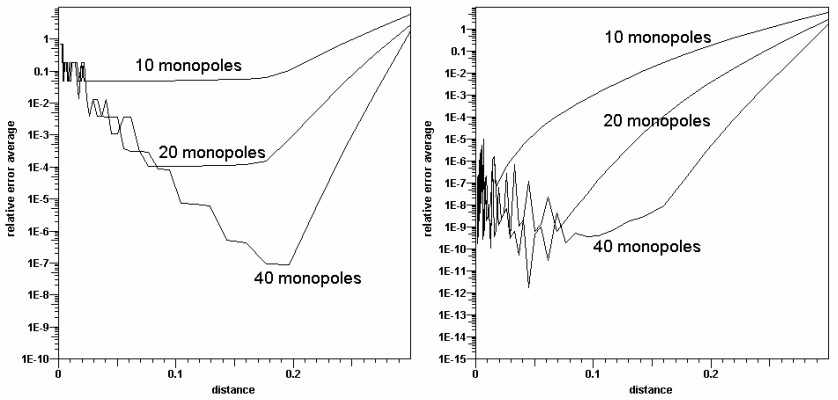

Fig. 12. Relative error average versus distance when the conjugate gradient matrix solver is applied for solving the system (20). Left hand side: same as in Fig. 9, right hand side: same as in Fig. 10.

when the CG solver is used instead of the Givens procedure. Therefore, the CG solver is not always more efficient than the Givens solver. We will see in the next section how to take advantage of the CG solver.
CG without preconditioning, i.e., without column scaling leads to very bad results. Therefore, the column scaling might also be helpful for direct matrix solvers. Unfortunately, efficient column scaling requires the storage of the entire rectangular matrix R. Since R has typically 4-40 times more elements than the triangular matrix stored in the Givens update procedure, scaling requires more memory and computation time. As we can see from Figure 13, scaling also improves the quality of the results obtained with the Givens solver, but here, the effect is not big at all.

1.5. Parameter Estimation Techniques (PET) for speeding up matrix methods

When we compare frequency domain methods with time domain methods, we see that the time domain methods usually use iterative procedures for tracking the time. When the field is known for a certain time t_0, it is computed for the next point $t_1=t_0+\mathrm{d}t$ usually with a single iteration. This procedure only works when the time increment dt is sufficiently small. The known field at t_0 can be considered as a good starting value for the iterative procedure that computes the field at t_1.

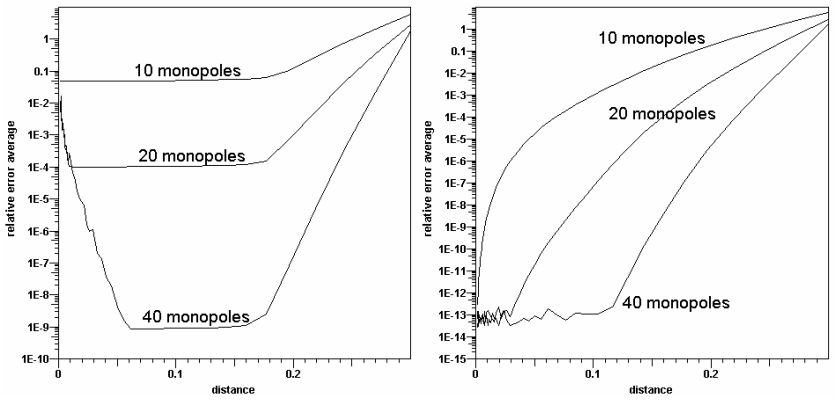

Fig. 13. Relative error average versus distance when the rectangular matrix is scaled before the Givens solver is applied for solving the system (20). Left hand side: same as in Fig. 9, right hand side: same as in Fig. 10.

Frequency domain methods usually consider the frequency as a constant rather than a variable. This is reasonable when the sources operate at a single frequency, but it is not reasonable when one computes the frequency dependency. If we know the field for some frequency f_0, we can assume that it is not much different for another frequency $f_1=f_0+df$, as long as the increment df is small enough. The same holds for all parameters of a series expansion that is used to approximate the field, for example, for the parameters A_k in (1).

In order to take advantage of the fact that $A_k(f_0)$ is a good approximation of $A_k(f_1)$, we need an iterative matrix solver, for example, the CG solver mentioned above. In order to illustrate this, we consider a simple 2D photonic crystal consisting of 8 columns and 8 rows of circular dielectric rods with a missing row of rods in the center as shown in Figure 14. Note that the lower half of the rods is not explicitly shown because of the symmetry of the configuration. The number of CG iterations required depends very much on the frequency. It is between 100 and 500 iterations when the CG iterations are started at $A^{start}{}_k(f_1)=0$. When we start at $A^{start}{}_k(f_1)=A_k(f_0)$ instead, only about 40 iterations are required. This can be considered as a 0 order PET.

As soon as we know the solutions of two frequency points, we may replace the 0 order by a first order extrapolation, i.e., by a linear extrapolation. When we have three frequency points, we may apply second order extrapolation and so on. In fact, high order extrapolations

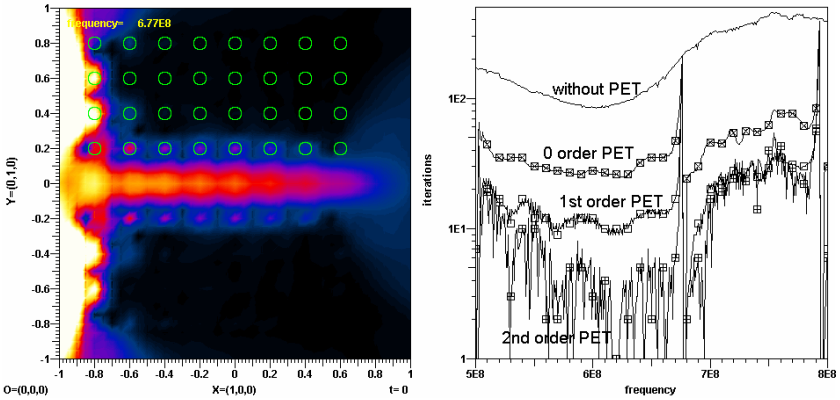

Fig. 14. PET for speeding up the computation of the frequency dependence. Left hand side: Time average of the Poynting vector field of the test example at the critical frequency. Right hand side: Number of iterations required for each frequency step for computations without PET and with 0 up to 2nd order PET.

using standard power series are not helpful. As we can see from Figure 14, 1st order PET is considerably better than 0 order PET, whereas 2nd order PET outperforms 1st order PET only within some frequency range. It is interesting to see that all PET solutions require the same number of iterations as the solution obtained by the brute force approach without PET for some critical frequency. The plot of the Poynting vector field in Figure 14 at this frequency shows that the mode in the defect waveguide in the photonic crystal is evanescent rather than guided above this frequency. There is a higher frequency, where the PET also "fails". This is the upper frequency of the photonic band gap. For frequencies of the guided mode within the band gap, an average of 3 iterations instead of 100 iterations is required, i.e., 2nd order is approximately 30 times faster than the brute force CG solution.

The PET can be further improved by using smart extrapolations that allow one to estimate the starting values of the parameters more accurately. As one can see already from the simple example above, this is not easy because the extrapolation may considerably depend on the frequency. A more detailed analysis shows that the optimal extrapolation also depends on the type of basis functions used in the series expansion and on details of the numerical model [6]. Searching for an optimal extrapolation for a given class of problems is reasonable when such

problems are to be solved frequently, but this search process is time-consuming.
Note that the PET is not restricted to MMP. It may be applied to all kind of frequency domain methods. Furthermore, the PET may also be applied for sequences of similar numerical models, i.e., it is not restricted to the calculation of the frequency dependence. When p is any parameter of the model, for example, the angle of incidence of a scattering problem, we can say that $A^{start}_k(p_1)=A_k(p_0)$ is a reasonable start value that leads to a 0 order PET, and so on. The PET may be further generalized when several model parameters are varied, which is often the case when extended studies are performed, for example, when one wants to know the dependence of the efficiencies of a grating on the frequency and on the angle of incidence. Similarly, the PET is very helpful for model optimizations, where one optimizes several model parameters.
In the following section, we will se a variant of the PET that is helpful for nonlinear eigenvalue problems.

1.6. The MMP eigenvalue solvers

Many standard problems of electro- and magnetostatics, electromagnetic scattering, antenna, etc. are characterized by an inhomogeneous system of equations of the form MX=B, where M is a matrix, X is the vector to be computed, and B is a known right hand side vector that essentially contains the excitation of the field. When one considers resonators and guided waves, one usually uses idealized, energetically closed models without any excitation. As a consequence, a homogeneous matrix equation of the form

$$\mathrm{M}(e)\mathrm{X}(e) = 0 \tag{21}$$

is obtained, where e denotes the eigenvalue. For resonators, e is usually the resonance frequency, for waveguides, it may be the propagation constant. Equation (21) obviously has the trivial solution X=0 that is of no interest. When M is a square matrix, equation (21) has non-trivial solutions when

$$Det(\mathrm{M}(e)) = 0 \tag{22}$$

holds. In general, equation (22) is a transcendent equation that may have several solutions. Therefore, equation (22) must be solved numerically. Note that the numerical solution of this equation can cause several problems. First of all, it is important to accurately and efficiently

compute the determinant of a matrix that may be big and also ill-conditioned, especially near the zeros of the determinant. As a consequence, Det(M(*e*)) may be quite noisy near the zeros. If this happens, the routine that detects the zeros of the determinant will be heavily disturbed. It may even detect several zeros, i.e., several eigenvalues instead of a single one. This is a very difficult problem when one has almost degenerate modes, i.e., two modes with almost identical eigenvalues. This situation is often encountered in the computation of band diagrams of photonic crystals. Therefore, equation (22) is of questionable value.

Since MMP works with rectangular matrices R(*e*), a square matrix equation of the form (21) is obtained by a multiplication with the adjoint matrix R*, i.e., one then has M(*e*)=R*(*e*)R(*e*). We have already seen that it is numerically better to work directly on the overdetermined system, especially when the condition number might cause problems. Therefore, we consider the overdetermined system

$$\mathrm{R}(e)\mathrm{X}(e) = \mathrm{E}(e) \tag{23}$$

where E(*e*) is the residual vector that should be minimized. Note that E(*e*) is unknown. The square norm of the residual vector Res^2 is now a positive function of the eigenvalue *e*. The main question is the following: Is

$$\mathrm{Res}^2(e) = \min. \tag{24}$$

a transcendent equation that defines the eigenvalues? It is hard to answer this question. Numerical experience shows that equation (24) may be used in simple situations, but in complicated situations it may also lead to almost trivial solutions. In order to avoid such almost trivial solutions, one must define the amplitude of the field, which is also a function of the eigenvalue *e*. Once this has been done, one can minimize the following search function

$$F^{search}(e) = \mathrm{Res}^2(e) / \mathrm{Amp}^2(e) \tag{25}$$

instead of equation (24). Note that the definition of the amplitude Amp(*e*) is not unique. Typically, one can use voltages, i.e., integrals over the electric field along a line, currents, integrals over energy densities, the power flux, etc. According to our experience, it is best to define Amp^2 as the integral over the energy density for resonators and similar problems, namely photonic band diagram computations, whereas the definition of Amp^2 as the power flux is best for the computation of

guided waves. For guided waves on transmission lines, the definition of the amplitude as voltages between the conductors or as currents in the conductors seems to be more natural, but these definitions may lead to search functions that exhibit "misleading minims" at high frequencies, i.e., minims that do not correspond to eigenvalues. It is not difficult to detect such minims, but testing the validity of the minims requires additional efforts.

The accurate computation of integrals over some field components or over the energy density may be time consuming. According to our experience, very rough evaluations are sufficient, i.e., the integrals may be approximated by sums over a few test points. Often, even a single test point is sufficient. If a single test point is used, it must be placed at a position where it is guaranteed that the field of the mode to be searched is not zero or even relatively strong.

The standard MMP matrix solver is based on Givens plane rotations. This solver transforms the rectangular matrix R into an upper triangular matrix T. When this algorithm is applied to equation (23), one obtains

$$\mathrm{T}(e)\mathrm{X}(e) = 0 \tag{26}$$

The last element t_{nn} of T contains the residual. A careful analysis shows that the corresponding last parameter, i.e., the last element x_n of the parameter vector X and therefore the last basis function $Basis_n$ of the series expansion used for the approximation of the field plays a prominent role: If $Basis_n$ is chosen in such a way that the scalar product of $Basis_n$ with the field of the mode to be searched is zero (or numerically almost zero), the last equation of (26), i.e.,

$$t_{nn}(e)x_n(e) = 0 \tag{27}$$

holds because x_n is zero. Therefore, the mode cannot be found. This may be useful when one is interested in suppressing some mode. For example, when one has two degenerate or almost degenerate modes with different symmetries, e.g., an even and an odd mode, one can select a basis function $Basis_n$ that has the symmetry of the desired mode. The undesired mode will then be suppressed.

For applications where one wants to detect all modes and where the number of modes is high, an algorithm that suppresses some mode may be inconvenient. This especially holds for the band diagram computation of photonic crystals as illustrated in Figure 15.

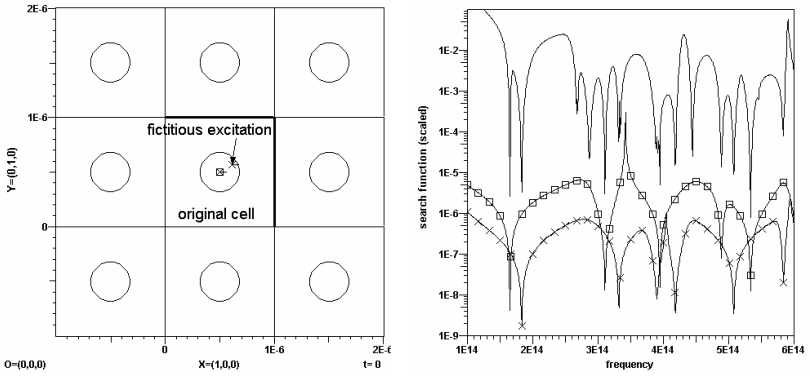

Fig. 15. 2D photonic crystal consisting of dielectric rods. Left hand side: 2D MMP model. Only the two thick lines (periodic boundaries) and the circle in the center are modeled explicitly. The field is modeled by two Bessel expansions and a single multipole expansion in the center. For the MMP-MAS eigenvalue solver, a fictitious monopole excitation is introduced at a random position. Right hand side: The different eigenvalue search functions at the Γ point. 1) MMP-MAS search function (without markers), 2) standard MMP search function for 0 order multipole at "last" position, 3) standard MMP search function for 1st order multipole at "last" position.

When we recognize that the last basis function somehow plays the role of the excitation of the mode, we can easily find an alternative method by introducing a fictitious excitation [14]. This technique has also be used in the context of the MAS [9] and for other methods [15]. We will call it MMP-MAS eigenvalue solver in the following.

Essentially, the MMP-MAS solver mimics the practical measurement of eigenvalues such as resonance frequencies: In order to measure the resonance frequency, one must feed the resonator and one must measure the amplitude somewhere. For doing that, one must open the energetically closed structure. Numerically, we can easily excite any mode, for example, with a fictitious monopole or dipole source and we can test the amplitude of the field in some testing point. Resonances are characterized by peaks of the amplitude observed in the test point P^{test}:

$$\mathrm{Amp}(p^{test}, e) = \max. \quad or \quad 1/\mathrm{Amp}^2(p^{test}, e) = \min. \qquad (28)$$

are therefore alternatives for detecting the eigenvalues. Note that there is some similarity between equation (28) and equation (25), but there is an important difference as well: The residual is missing in equation (28). A careful analysis [14, 16] shows that the residual is a very flat function of e when a fictitious excitation is introduced, but in the close vicinity of the

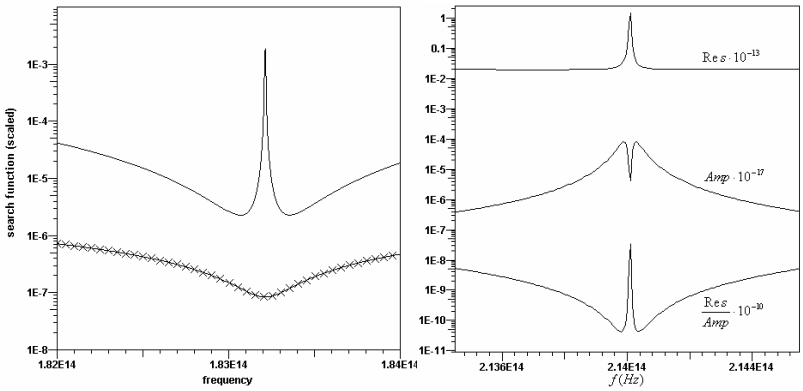

Fig. 16. Double peak and twin minimum phenomenon for a 2D photonic crystal consisting of dielectric rods. Left hand side: Zoom of Fig. 15. The standard MMP eigenvalue search function (x markers) exhibits a single minimum. When a fictitious excitation is introduced, the same search function exhibits the twin minimum phenomenon. Right hand side: The different parts (residual and amplitude) involved in the MMP eigenvalue search function for another point of the band diagram.

correct eigenvalue, it exhibits a sharp peak, whereas Amp(P^{test},e) exhibits a sharp dip as shown in Figure 16. We call this double peak phenomenon.

As a consequence, $1/\mathrm{Amp}^2(P^{test},e)$ and the MMP standard search function (25) exhibit twin minims (see Figure 15). The double peak phenomenon only occurs when fictitious excitations are used and it considerably disturbs the search procedure when a very fine search is performed for obtaining highly accurate results. In order to get rid of this problem, one can introduce a more sophisticated search function that is composed of different powers of Amp and Res, one can implement smart search strategies that consider both Amp and Res, or one can introduce small, fictitious losses: As in the measurement of resonators, the resonance peaks get less sharp and more smooth when losses are present.

Beside the double peak phenomenon, the method of the fictitious excitation exhibits similar problems as the standard MMP eigenvalue search procedure: Both the fictitious excitation and the test point must be appropriately set. Otherwise, some of the eigenvalues will not be detected. In order to make sure that all eigenvalues are found, one can easily introduce several test points and one can also introduce several fictitious excitations. We usually use this technique for the computation of complete band diagrams of photonic crystals [14]. Whereas we use the

standard MMP eigenvalue search method more frequently for the computation of guided waves.

Numerically, the eigenvalue search procedure may be very time-consuming, because the search function $F^{search}(e)$ must be evaluated many times for different estimates of the eigenvalue e. In order to evaluate the search function, a matrix equation of the form of equations (23) must be solved. Therefore, it is important to develop efficient eigenvalue search strategies that can precisely find the minims of $F^{search}(e)$ with a minimal number of evaluations of $F^{search}(e)$. The standard MMP search strategy starts with a rough search on a regular grid within a reasonable search interval. The grid of the rough search should be as coarse as possible (in order to reduce the number of search function evaluations), but fine enough for detecting all eigenvalues. Note that the search grid also depends on the definition of the search function.

As soon as the rough search has detected all minims of $F^{search}(e)$, a fine search routine precisely detects each of the eigenvalue. For the fine search, a smart routine using parabolic and geometric interpolations is used. Beside the search algorithm, stopping criteria are very important of the speed of the fine search. Here, MMP has nice advantages because both the accuracy of the eigenvalue and of the field of the correlated mode can be estimated from the search function. Note that this is not possible, when the eigenvalue search is based on equation (22).

Usually, one is interested in the dependence of the eigenvalue on some model parameters or some variables. For example, when we consider guided waves, we may consider the propagation constant as the eigenvalue and we then want to know its frequency dependence. In the computation of band diagrams, we search for the resonance frequencies depending on the orientation of the wave vector. Thus, we often want to compute diagrams of $e(v)$, where v is either a scalar (real or complex) or a vector. Of course, we can restart the eigenvalue search for a finite set of values v_i , i=1,…I for plotting such diagrams. When I is big, this brute-force method is inefficient. Since we usually have $v_{i+1}=v_i+dv_i$ with a relatively small step dv_i, we may proceed as in the parameter estimation technique, i.e., we may implement an Eigenvalue Estimation Technique (EET) as a variant of the PET. The EET implemented in the MaX-1 code is very efficient. Depending on the desired accuracy and on the step size, one typically requires 4-10 iterations per eigenvalue $e(v_i)$.

More tricky problems occur when one searches for eigenvalues in the complex plane for lossy structures. Note that handling lossy structures as

complex eigenvalue problems is not obvious from the theoretical point of view, because lossy structures are not energetically closed. Furthermore, the search for minims in the complex plane is numerically much more demanding because the search are is two-dimensional rather than one-dimensional, because the search area is not strictly limited, and because the eigenvalue search function may exhibit discontinuities due to cuts in the complex plane. For example, when we consider a loss-free optical fiber, we know that the propagation constant of all guided modes is real within a well-defined, limited interval. When the fiber coating is lossy, each guided mode gets a complex propagation constant, but we observe an infinite number of additional modes that usually exhibit relatively large imaginary parts of the propagation constant [17]. Even if one focuses on the modes with reasonably small imaginary parts of the propagation constant, detecting these modes and tracking their frequency dependence is very demanding. Here, the EET is very useful and allows one to drastically reduce the computation time.
Note that beside losses caused by a finite conductivity, also radiation losses may cause a complex propagation constant. This is typical for waveguides in finite photonic crystals that will be presented later.

1.7. The "connection" concept

It has been mentioned that MMP is a semi-analytic method in the sense that the basis functions or basis fields in the series expansion (9) that approximate the field are analytical solutions at least in a single domain. Typical expansions are plane waves and multipole expansions. These expansions are single domain expansions, because they fulfill the Maxwell equations only in one domain with specific material properties. Outside this domain, they are not evaluated, i.e., set to zero. Because of the linearity of the Maxwell equations, linear superpositions of expansions that fulfill the Maxwell equations will also fulfill the Maxwell equations – provided that the material properties are linear as well. Therefore, we may easily construct more complicated basis functions for equation (9) by defining new basis functions, such as

$$Basis_k(\vec{r},\omega) = \sum_{j=0}^{J} A_{kj}\, Basis_j(p_{kj},\vec{r},\omega) \tag{29}$$

We call this a connection. Note that connections may be nested, i.e., some of the basis functions that define a connection may be other

connections. Furthermore, a connection may contain basis functions that are defined in different domains. Therefore, a connection is usually a multi-domain expansion. A simple example is presented in the following: Assume that a plane wave hits a planar boundary between two different dielectrics. The analytical solution consists of a simple superposition of a reflected and a transmitted plane wave with well defined amplitudes. These plane waves may easily be packed in a connection. When one considers a plane wave that hits a planar boundary with a bump, the connection mentioned above is an excellent basis function that fulfills the boundary condition everywhere, except on the bump. In order to obtain an accurate solution, one can add some multipole expansions in the vicinity of the bump.
More advanced problems where connections play an important role are waveguide discontinuities that usually consist of several waveguide sections (waveguide ports) and the discontinuity region itself. For handling such problems, one can first compute the fields of all modes in all waveguide sections (using the eigenvalue solver outlined above) an pack each mode, i.e., the corresponding basis functions and amplitudes, in a connection [6]. This procedure is extremely helpful for the simulation of arbitrary structures embedded in finite photonic crystals [8, 18] because it allows one to work with very efficient models where only a small discontinuity area must be modeled explicitly. We will consider such problems later in this chapter.

1.8. Validation problems

It is well-known from traditional domain methods that the results sometimes look reasonable but turn out to be rather inaccurate or completely wrong. Therefore, many scientists do not trust the simulations and demand experimental validation of the results. When numerical simulations shall be used to abbreviate the design process, to optimize a structure, or to analyze new concepts of structures that cannot be fabricated yet, such codes are not very helpful. Therefore, there is a strong need for robust codes with reliable results that maybe validated even when neither a comparison with experimental results nor with other numerical methods is possible.
Because of its semi-analytic nature, MMP allows one to obtain highly accurate results. In the close nearfield area and even on the boundaries, the field may be computed accurately. From the analysis of the boundary

conditions that are matched numerically, one obtains an excellent error estimation because the electromagnetic field is usually most complicated near the boundaries. Therefore, the maximum errors typically occur on the boundaries. The error distribution along the boundaries and their characteristic average and maximum values allow one to easily validate the results. The MMP implementation of MaX-1 allows one to plot the absolute and relative error distributions along the boundaries, not only in the matching points but everywhere else (see Figures 2, 4, 7, 8, etc.).
The error distributions along the boundaries may not only be used for getting confidence in the results, but also for improving the numerical model. Usually, it is easy to reduce the error in some region by adding multipole expansions in this region. Therefore, one can systematically reduce the maximum error with adaptive multipole setting procedures.
When the error distribution is known, one can also use it for obtaining more efficient models with acceptable accuracy by reducing the orders of the multipoles in those areas where the accuracy is relatively high.
In many applications, the near-field computation plays a minor role because one is mainly interested in derived quantities such as the absorbed power, the directivity of an antenna, the S parameters of a filter, and so on. Such quantities usually contain integrals over the field that may be accurate enough even when the field along the boundaries is very inaccurate. For estimating the accuracy of such quantities, one can compare results of different numerical methods. Here, the fast convergence of usual MMP models is very helpful. MaX-1 contains several tools for the comparison of solutions obtained with different models. Thus, one usually can internally validate the MMP results very well, without any comparison with measurements or with other codes.

1.9. Comparison with MoM

Although the MoM is originally designed as a domain method that is very close to the method of Finite Elements (FE), it is also closely related to the MMP method in the loss-free case. When one compares MoM and FE, one sees that the main difference is the following: in the FE one usually approximates the electromagnetic field or the potential, whereas one starts with a series expansion approximation of the sources (charges, currents) of the field in the MoM. In the loss-free case, the sources are located on the boundaries. Therefore, a discretization of the boundaries is required (In terms of the MoM, the boundaries of the field domains are

the domains of the field sources. In order to avoid confusions, we will not use this terminology here.). This avoids the problems associated with the truncation of the infinite space (absorbing boundaries) of standard domain methods.

When one considers the MoM as a "source expansion method", one immediately can see its vicinity to the MAS. Since the MAS is a special case of the MMP method, it is sufficient when we compare the MoM with the MAS. The main difference between the MAS and the MoM for loss-free problems is that the MoM approximates the sources where they really are, i.e., on the boundaries, whereas the MAS introduces auxiliary sources at some distance from the boundary. We have already seen that moving the auxiliary sources away from the boundary allows one to drastically increase the accuracy of the results without increasing the memory and computation time. The only price that has to be paid is the higher condition number of the matrix. Putting point sources (point charges in electrostatics) directly on the boundaries has very bad consequences. First of all, the field of these sources on the boundaries has singularities that cause severe numerical problems. As a consequence, the field evaluation on the boundaries and in the close nearfield becomes completely wrong.

Standard MoM implementations avoid or reduce such problems by more sophisticated approximations of the source along the boundaries. For reasons of simplicity, we only consider the simple electrostatic case here. On a sufficiently smooth metallic electrode without corners, the surface charge density is also a smooth function without singularities. Therefore, it can be approximated by a piecewise constant function, by a piecewise linear function, or by some other sophisticated functions. Instead of this, one can subdivide the boundary into several boundary elements (subdomains in the MoM terminology) and assume that the charge distribution on each element is constant, linear, or a more sophisticated function. In order to obtain the corresponding electrostatic potential, one must solve the Coulomb integrals, i.e., one must integrate the charge distribution over each element. This integration is numerically time-consuming when the element has a complicated shape or when it is curved. Therefore, one usually approximates the physical boundary by a piecewise flat boundary, in the 2D case, by a simple polygon. This approximation introduces modeling errors that cannot be analyzed easily. Furthermore, when simple basis functions are used on each element, the resulting electric field is discontinuous or it may even have singularities

on the boundary. Therefore, more sophisticated basis functions are desirable, but for such functions one does not have analytical solutions of the Coulomb integrals and the numerical integration is time consuming. Since one usually has many boundary elements, the price to pay is numerically high.

When we consider the field of an auxiliary source, we can notice that this field may also be obtained from a smooth charge distribution along the entire boundary. In the MoM terminology, one therefore has an "entire domain" basis function that is known to be much better than the standard "subdomain" basis functions. Furthermore, the Coulomb integral over this "entire domain" basis function is known and must not be solved. What is unknown, is the charge distribution along the boundary. In fact, it is not necessary to know the charge distribution for the implementation of the field solver. Therefore, the MAS (and MMP) is an elegant and efficient technique that removes the major MoM problems. Once more, the only prize we pay is the higher condition number.

Since the field distribution along the boundaries (where the boundary conditions must be imposed) is not smooth is standard MoM codes, the Simple Point Matching (SPM) technique is not very reliable. Therefore, a variant of the projection technique is often used in the MoM. For rather obscure reasons, a special product is introduced that only coincides with the traditional scalar product for real functions. This product is then used to "project" the boundary conditions on some testing functions defined along the boundaries. Both the basis functions (charge distributions) and the testing functions along the boundaries may be defined everywhere or in some boundary elements only. The SPM is obtained, when Dirac testing functions are used. This allows one to analytically solve the integrals contained in the products mentioned above, but it is known to be the worst choice of testing functions. The optimal choice of testing functions is considered to be Galerkin's choice, where basis and testing functions are the same. We have already seen that this leads to an error minimization in the least squares sense and that a numerically better way is the GPM that avoids time-consuming numerical integration as well as problems with high condition numbers.

From these considerations, one can easily understand that the MAS-MMP approach is more efficient, more accurate, and easier to implement than standard MoM, provided that accurate results or near-field computations are desired. When only farfield approximations of a moderate accuracy are desired, the situation may be different for the

following reason: The MAS-MMP codes usually have a very rapid convergence when the surface is smooth enough, but they have not been tuned for obtaining quick and approvimate results for complicated geometry that exhibit field singularities near sharp corners etc. Typical examples are airplanes, microwave circuits, etc. For such situations, one can design specialized MoM codes based on advanced concepts, such as non-free-space Green's functions, that are much better suited.

A simple case, where the standard MMP-MAS approach is not efficient, is a thin but long wire. For such an object, the auxiliary sources or multipoles are best placed along the center of the wire. In order to obtain accurate results, the distances between neighboring sources should not be much larger than the radius of the wire. Therefore, one obtains a huge number of source when the wire is long – even when it is short compared with the wavelength. One then observes that the amplitudes of the sources vary slowly along the wire. In order to overcome this problem, one can introduce a new type of basis functions, so-called distributed multipoles [6]. When an array of multipoles is distributed along a straight line, we call it line multipole. When it is distributed along a circle, we call it ring multipole. Line multipoles are a generalization of the thin-wire approximation that is frequently used in MoM codes. The standard thin-wire expansion is a zero order line multipole. The more general line multipoles may be used not only for thin wires, but also for non-circular wires and for relatively thick wires. In fact, each Green's function that may be used in a MoM code is a solution of the Maxwell equations. Therefore, it can be implemented and used in the MMP code as well. As soon as such a basis function is available in MMP, the main difference to its use in a MoM code will be the following: In the MoM code, it will be used in its "natural or physical" place, whereas it can be used in the MMP code at some arbitrary place - just as a new basis function that is helpful for the field approximation. For example, a MoM current patch is used on the surface of a conductor, where the currents really are, whereas the MMP user will move the same current patch inside the conductor in order to avoid the field discontinuity problems on the boundary. Therefore, MoM must not be considered as a concurrent of MMP. It is a useful source of possible basis functions.

Note that ring multipoles are not only useful for toroidal objects but also for general axisymmetric problems. Such problems can often be handled with ring multipoles with similar computation times as 2D problems.

2. The MaX-1 implementation of MMP

The most advanced MMP code is contained in the MaX-1 software package [8, 18]. This package contains helpful tools for the advanced modeling, visualization, animation, and validation. The main concepts are outlined in the following section.

2.1. Advanced modeling features

Since MMP is a semi-analytical field solver that is very close to analytical solutions, it is important that modeling tools are available that allow one to describe the geometry and other data as accurately as possible because the accuracy of the input data limits the accuracy of the results. For example, when the locations of the matching points and the tangential directions of the boundaries in the matching points are defined with a certain accuracy, one cannot obtain the near field with a higher precision. Therefore, a very precise input is desirable. For this reason, MaX-1 does not only permit real and integer data input, but also formula input. In order to avoid time-consuming formula evaluations, it uses a simplified formula interpreter that is general enough for standard tasks.
The formula input can be used, for example, for the definition of frequency dependent material properties, for the analytical definition of boundaries, and for the definition of curved grid lines.
For 2D modeling, standard boundaries in most of the technical applications are polygons, sometimes with circular parts. It is well-known that the electromagnetic field usually exhibits singularities near corners. In order to avoid the difficulties associated with singularities, one best inserts reasonably small arcs at the corners of the polygons. Since manual insertion of arcs is tedious and a source of inaccuracies, this is done automatically in MaX-1. A polygon with appropriate arcs in each corner is called a C-polygon.
The error analysis shows that the highest errors of models with C-polygonal boundaries always are observed near the junctions of the arcs with the straight polygonal lines. From the theoretical point of view, this is to be expected, because the second order derivatives of the boundary function in these points are discontinuous. An improvement may therefore be obtained by a cubic spline approximation of the C-polygons. This approximation is also implemented in MaX-1. It is considered to be sufficiently accurate for most of the practical applications.

When the 2D boundaries and additional model parameters (material properties, frequency, geometric symmetries, etc.) are defined, an appropriate set of basis functions – mainly 2D multipoles – must be defined. For inexperienced MMP users, this is the most difficult task. Therefore, MaX-1 contains several automatic pole setting procedures.
The construction of geometric 2D objects is relatively simple because a computer screen is well suited for that. 3D modeling is much more tricky. Since an analytical definition of the boundaries of 3D Objects is desired for 3D MMP models, standard input procedures of CAD packages are not appropriate. MaX-1 contains several transformations that allow one to construct parts of complicated objects from the definition of boundaries and 2d multipoles defined in the xy plane. In order to illustrate the procedure, we consider a grating of "hat-like"

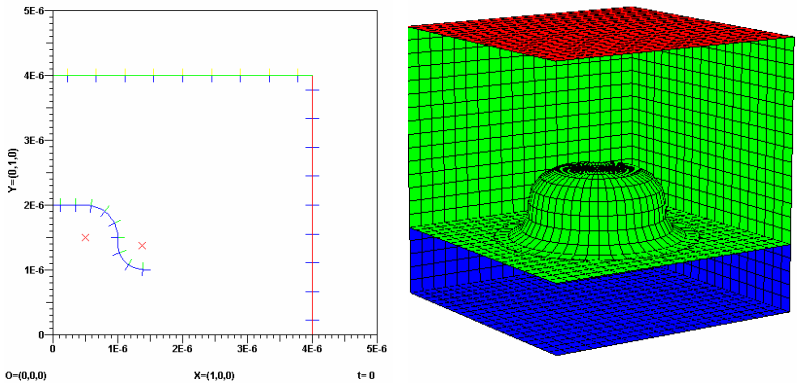

Fig. 17: Grating with "hat like" bumps, periodic in x and y directions (right hand side) and 2D constructions required (left hand side) for generating the 3D model.

bumps that are periodic both in the x and y directions. A single cell of this grating shall be a cube that contains a single "hat" as shown in Figure 17.
The "hat" is an axisymmetric part that is obtained from a the blue C-polygon in Fig. 17 by rotation about the y axis. For this C-polygon, MaX-1 could generate a set of appropriate multipoles. Since only a rough model shall be computed first, only 2 multipoles (red crosses in Fig. 17) are introduced manually. Beside the "hat", the original cell of the grating to be modeled consists of several rectangles that may be obtained by cylindrical translations from the red and green line in Fig. 17. When the

"hat" and the 7 rectangular boundary parts shown in Fig. 17 have been defined, one should delete the circular area of the green rectangle below the hat. The MaX-1 user can do this by the simple directive "ADD INHIBIT D8R6" which means: Add the inhibit directive D8R6 to the list of inhibit directives that are performed before the 3D matching points are generated. D8R6 means "Delete all points of object 8 (the "hat") on the Right hand side of object 6 (the green rectangle mentioned above)".

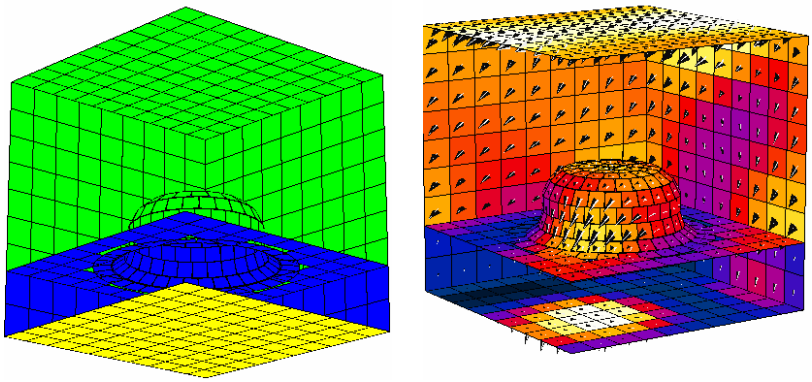

Fig. 18: Grating with "hat like" bumps. Left hand side: 3D matching points. Right hand side: Time-average of the Poynting vector field in the matching points.

Figure 18 shows the matching points that will be generated.
As soon as the 3D boundaries and expansions are defined together with additional data (material properties, frequency, periodic symmetry, etc.), MaX-1 will generate the 3D matching points, impose the boundary conditions, set up and compute the resulting MMP matrix. After that, MaX-1 can evaluate the electromagnetic field and derived fields such as the Poynting vector field, everywhere in space, for example on the matching points themselves as shown in Figure 18.

2.2. Validation, visualization and animation

Visualizations are especially helpful for the model validation, for the analysis of results, and for understanding how electromagnetic fields propagate. Since MMP is a semi-analytical technique, it allows one to compute the electromagnetic fields and arbitrary quantities derived from this field with a high precision everywhere

in space, in the far field as well as in the near field or even on the boundaries. For doing that neither interpolation nor extrapolation routines are required.
Since we also can compute the field in the matching points on all

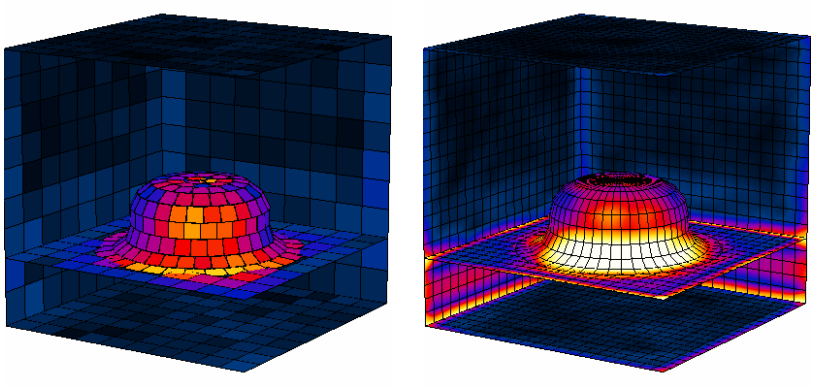

Fig. 19: Distribution of the mismatching error along the boundaries of the problem shown in Figs. 17 and 18. Left hand side: error distribution in the matching points, right hand side: error distribution on a grid. Note that the maximum relative error of this model is 3.3% and the average error is 0.8%.

boundaries, we can easily obtain the mismatching error as in Figure 19.
Although the error distributions in the matching points are useful for a quick overview and for an adaptive model improvement, the errors along the boundaries between the matching points may be higher than between the matching points, namely when simple models with a weak overdetermination are used. Therefore, one can also compute the error distribution on an arbitrary grid as shown in Figure 19.
For more detailed model validations, one can easily modify the MaX-1 models by changing the MMP expansions by increasing the multipole orders. From the comparison of different models one can also estimate the errors in points that are not on the boundaries.
For the visualization of the surfaces of 3D objects, OpenGL is known to be very convenient. However, since the electromagnetic field is a quantity defined everywhere in space and not only on the

surfaces of 3D objects, its visualization is much more demanding. Therefore, one usually visualizes the field only on one or a few rectangular patches. Since MaX-1 allows one to visualize arbitrary, curved boundaries in 3D space, it allows one to also display the field and derived fields on such curved surfaces as shown, for example, in Figure 18. Such visualizations often give a much better insight than visualizations on rectangular patches.
When the field is visualized on several surfaces, hiding becomes an important issue. The standard OpenGL features allow one to draw the fields once and to rotate, pan, zoom while one is observing the field. Although this is very helpful, it still allows one to observe some portions of the field only. In order to obtain a deeper insight and more information, animations are required. In addition to standard animations that show the time-dependence of the field, MaX-1 allows one to easily generate advanced animations that show the field dependency on some other variables or model parameters. For defining advanced animations, the user must instruct MaX-1 what shall be shown on the frames of a movie. This is done by so-called movie directives that define actions such as "increase time", "increase frequency", "rotate object", "move expansion", and many more.

2.3. Solving complicated projects

The movie directives are essentially a set of batch commands for the 3D MaX-1 modeling that are similar to the well-known batch directives of MS DOS. These directives are not only useful for generating advanced animations, but also for solving complicated projects. For example, when we want to compute and draw the band diagram of the first N modes of a 2D photonic crystal, we have 3 sections of the band diagram that correspond to the three sides of the 1st Irreducible Brilloin Zone (IBZ). Since degenerate modes often occur in the tree corners of the IBZ, we usually start the rough eigenvalue search in the centers of the three sections. Once we have detected the first N modes for a center point, we perform a fine eigenvalue search with the desired accuracy for

each mode. Then we track each eigenvalue from the center point to the left and to the right end of the corresponding section. For each point of the section we save the data of all *N* eigenvalues on a file. As soon as all eigenvalues are known along a section, we read the data file and let MaX-1 plot the diagram. The following MaX-1 directives are used for drawing a section of the band diagram:

```
write function band010.fun header 0 2
write function !
end pre processing directives
!
set period constant 3141592.6 0 0 0
write function + header 0 2
write function / text "n"
write function / text "fa/c"
end inner loop pre directives
!
mmp solve
write function / frequency real
increase period constant -78539.815 0 0 0
end inner loop directives
!
write function !
read function /
process function div(v1,3e14)
draw function 1 3 1
increase eigenvalue 1
end inner loop post directives
!
read directive band002.dir
end post processing directives
```

This set of directives consists of 5 blocks 1) pre processing directives, 2) inner loop pre directives, 3) inner loop directives, 4) inner loop post directives, 5) post processing directives. In fact, these blocks define two nested loops.
First, the block 1) is computed. Here, a "function" file prototype is defined. This file will be used for storing the band data. Its name is band010.fun. The second directive of block 1) closes the file.

Then the outer loop is performed for each of the *N* eigenvalues. This starts with the pre-processing directive block 2) for the inner loop. In the first line of this directive, the start point on the IBZ for the rough eigenvalue search is defined. The remaining directives of the block 2) open a new function file and write header data on it.
Now, the inner loop is started. In this loop, the directive block 3) is used. First, the MMP solver is called. Afterwards, the real part of the frequency (eigenvalue = resonance frequency) is written on the function file. Then, the search point on the segment of the IBZ is moved. This is repeated in the inner loop until the end of the search interval is reached.
When the inner loop is finished, the block 4) will be computed: First, the function file is closed. Then, it is opened again and all data are read. The third directive of block 4) scales the function, i.e., it divides the value by 1E14. Then, the function is drawn. After that, the eigenvalue number is increased by one and the outer loop is repeated until all eigenvalues are computed and plotted.
After the outer loop is terminated, the directive block 5) is performed. Here, a new set of directives, specified in the file band002.dir is read. This set of directives will replace the current one. It contains the directives for plotting the next section of the band diagram.
The movie directives of MaX-1 allow one to automatically perform complicated investigations not only of a single numerical model, but also of entire sets of models. Therefore, it may be used for the extensive analysis of problems.

2.4. Linking with optimizers

Because of its high accuracy and reliability, MMP is very advantageous not only for the analysis, but also for the synthesis. For the latter, MaX-1 must be combined with a numeric optimizer. Typically, an optimizer will "propose" some model parameters (material properties, geometric parameters, etc.) an it then expects the fitness value of this model from the field analysis tool. Within MaX-1, the model analysis can be automated as we have seen in

the previous section. Furthermore, the movie directives allow one to precisely define the fitness function to be returned to the optimizer. MaX-1 and an external optimizer exchange information simply through ASCII files. The optimizer writes the model parameters on a specific input file for MaX-1. As soon as this file is available, MaX-1 reads and deletes it. After that, it performs all directives. When the model is analyzed and the fitness is computed, it is written on a specific output file. While MaX-1 is running, the optimizer waits for the output file. As soon as this file is present, the optimizer reads and deletes it. After that it generates a new model while MaX-1 is waiting for the new input file.

Since the model analysis is by far more time-consuming than the model generation of the optimizer, this procedure can easily be extended for parallel processing or for computer networks.

Note that the choice of the optimizer and its speed depends very much on the quality of the forward solver, i.e., of the field analysis tool. The numerical inaccuracy of the latter can be considered as some noise that can heavily disturb the optimizer. Efficient optimizers therefore require accurate field solvers, whereas slow optimizers that are close to random search strategies may be best suited when an inaccurate field solver is used. Namely simple finite differences solvers often generate fitness functions with some stair-case behavior that is very inconvenient for most of the optimizers. This behavior considerably disturbs even robust stochastic optimizers such as evolutionary strategies and genetic algorithms. Even when the optimizer does not use gradient information of the fitness function, it will work more efficiently when the first (or even higher) derivative of the fitness function is continuous. Together with an optimizer, a slow but accurate forward solver therefore may outperform a fast forward solver with limited accuracy. Therefore, it makes not much sense to compare the speed of different forward solvers without linking them to an appropriate optimizer, when one wants to obtain an efficient synthesis package.

Although current personal computers are not fast enough for the efficient optimization of complicated structures, we expect that the

numerical optimization will soon be very important for the design of new structures and devices, namely for finite photonic crystals that are very promising but often counter-intuitive.

3. Applications

The range of applications of MaX-1 and MMP is very wide. Although these codes are mainly used in computational optics, they may also be applied to computational electromagnetics at much lower frequencies, even down to electrostatics. However, in the following, we focus on applications in the optical regime.

First of all, there are two classes of applications for Maxwell solvers: Scattering and eigenvalue problems. The former are characterized by an inhomogeneous system of equations, where the excitation or incident field defines the right hand side or inhomogenity of the system. The latter is characterized by a homogeneous system. As we have already seen, one can introduce fictitious excitations in eigenvalue problems in order to transform them into a scattering problem.

Beside the excitation, the geometric symmetries play an important role and allow one some classification. For computational optics, three types of symmetries are most important: 1) cylindrical symmetry, 2) rotational symmetry or axi-symmetry, 3) periodic symmetry in 1, 2 or 3 directions of space.

Cylindrical symmetry is very natural for waveguides such as optical fibers. Furthermore, this symmetry is often assumed in simplified models, because it allows one to work on 2D models. This especially holds for photonic crystals. Realistic 2D photonic crystals consist of arrays of cylindrical rods of finite length or of cylindrical holes in a plate of finite thickness. Obviously, all realistic "cylindrical" objects must have some finite length. Therefore, cylindrical (2D) models are always idealized models that ignore the effects caused by the finite length. Codes for scattering at cylindrical structures often assume that the incident wave propagates in a direction perpendicular to the symmetry axis. This is an additional assumption that leads to an additional

simplification of the codes. Since this does not allow one to consider oblique incidence and cannot be used for solving eigenvalue problems, the MMP implementation of MaX-1 considers the more general case of cylindrical symmetry with an electromagnetic field that is not constant in the longitudinal direction.

Although the theoretical handling of rotational symmetry is similar to cylindrical symmetry, its implementation is considerably more demanding for the following reason: In scattering problems, the excitation is most frequently assumed to be a plane wave. The longitudinal (z) propagation of a plane wave in the cylindrical case is simply described by the function $\exp(\mathrm{i}k_z z)$, where k_z is the longitudinal component of the wave vector. This longitudinal dependency of the incident plane wave matches perfectly the symmetry decomposition $\exp(\mathrm{i}\gamma z)$ of cylindrical problems, where γ is the propagation constant. For axisymmetric problems, the symmetry decomposition in angular (φ) direction (about the z axis) is described by $\exp(\mathrm{i}n\varphi)$, where n is an integer number. This essentially is a Fourier decomposition in φ direction. The fact that n is an integer number, whereas γ is real or even complex (in lossy problems) simplifies the procedure, but the angular dependency of a plane wave is not described by $\exp(\mathrm{i}m\varphi)$ in general. Therefore, one must first decompose the incident plane wave with an angular Fourier series and solve the problem for each Fourier component. For reasons of simplicity, one often considers the special case of axisymmetric scattering problems with an incident plane wave propagating along the z axis. Such a wave has only one angular Fourier component of the form $\cos(\varphi)$ or $\sin(\varphi)$, i.e., $m=\pm 1$. For each Fourier component of the field one can find a set of basis functions with the appropriate symmetry. For multipole expansions located on the symmetry axis, this symmetry decomposition is trivial. Since we also want to use off-axis multipoles, we must decompose them, which leads to the concept of ring multipoles [6]. As an alternative to ring multipoles, complex origin multipoles [6] may also be used, but these are restricted to non-toroidal topology and have not been implemented in MaX-1.

A spatial Fourier decomposition of the field is also used for periodic symmetries that are present in grating applications and in models of perfect photonic crystals. Since these models are infinite, they are idealizations that approximate realistic gratings and photonic crystals. There are several methods for handling periodic symmetries. First of all, one could proceed as for axisymmetric problems, i.e., one could implement symmetry adapted, periodic multipole expansions. This leads to convergence problems. Therefore, periodic symmetries are handled in MaX-1 with an alternative method, by the definition of periodic boundary conditions [6, 8]. These are very important for the modeling of gratings and photonic crystals.

3.1. Scattering and antenna

A standard scattering problem is defined by a set of domains with the corresponding material properties and the incident wave. In most cases, the incident wave is considered to be a plane wave. It is well known that realistic waves generated by an antenna are not plane waves, but at a sufficient distance from the antenna and within a sufficiently small area, the antenna field may be approximated very precisely by a plane wave. As an example, we consider the scattering of a plane wave at an array of 6 times 6 cylindrical dielectric rods. This configuration may be considered as a finite 2D photonic crystal. Of course, the response of this structure depends not only on the orientation of the incident plane wave, but also on the material properties of the rods and on the size and spacing of the array. As one can see in figure 20, the block of 6 times 6 rods behaves very much as a square dielectric block, i.e., one can describe it quite well by some kind of "macroscopic" dielectric model that ignores the geometry of the crystal. However, when we sweep the frequency, we find that there are some frequency ranges where the block behaves like perfect conductor and totally reflects the incident wave for all angles of incidence as shown in Figure 21.

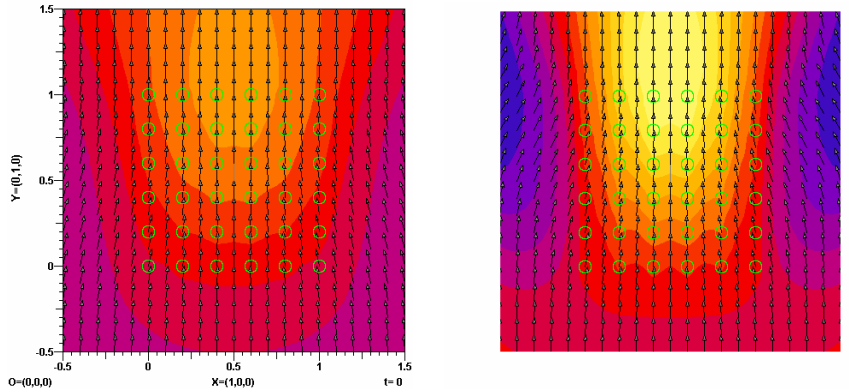

Fig. 20: Finite photonic crystal consisting of 6 times 6 dielectric rods, illuminated by a plane wave. Left hand side: low frequency, right hand side high frequency solution, time average of the Poynting vector field.

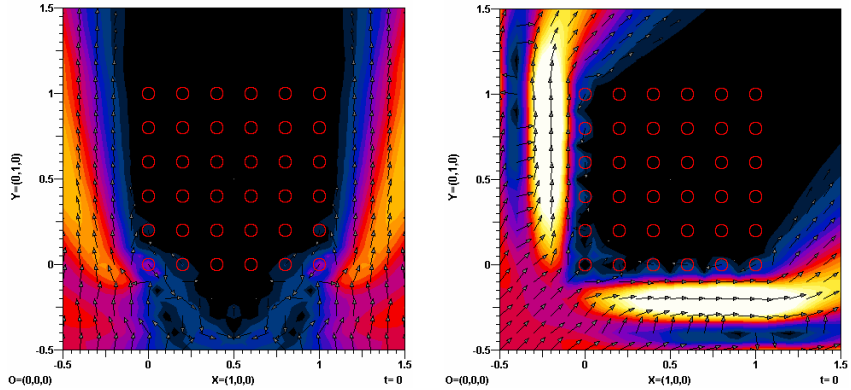

Fig. 21: Finite photonic crystal consisting of 6 times 6 dielectric rods, illuminated by a plane wave. Frequency within the band gap. Left hand side: perpendicular incidence, right hand side: diagonal incidence.

Since we know that the incident plane wave is a part of an antenna field, we would like to know how to model an antenna, although this term usually is not used in optics. A standard antenna is a more or less complicated structure with some waveguide that feeds the antenna. Antenna design programs often do not explicitly model the antenna feed. Instead of this, some impressed current

usually acts as the source of the field. This concept is not appropriate in optics. Here, we prefer the correct modeling of the feeding waveguide. As a simple example of an optical antenna, we can consider a finite photonic waveguide that contains a defect waveguide (see Figure 22). The defect waveguide acts as the antenna feed and it is considered to be infinitely long. MaX-1 allows us to compute the modes of such waveguides with different methods that will be outlined below. Once the modes of the waveguide are known, one can pack the expansions that describe the field of a mode in a connection. This allows us to truncate the infinite feeding waveguide and to model only a short section of it. Note that not only the mode that shall excite the antenna must be modeled, but also the modes that are reflected back into the waveguide feed. In addition to the guided modes, one may also have evanescent, higher order modes. When the waveguide section that is explicitly modeled is long enough, all evanescent modes will be damped sufficiently well in this section and need not be modeled by connections. On the other hand, in order to reduce the model size, one can also compute the relevant evanescent modes and pack them in connections.

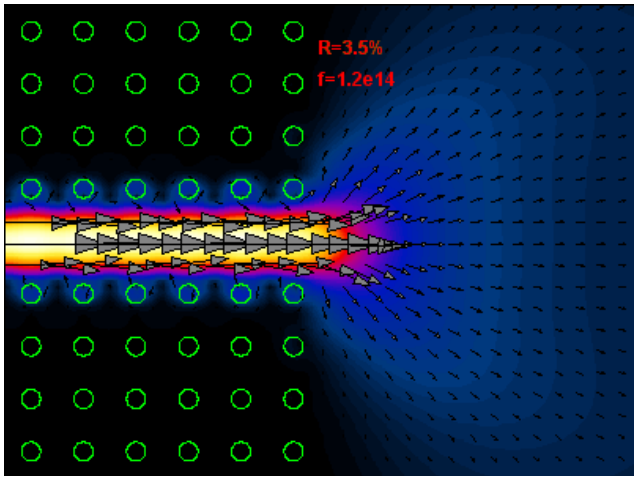

Fig. 22: Time average of the Poynting vector field for a simple photonic crystal antenna fed by a single mode defect waveguide.

3.2. Axisymmetric structures

The handling of axisymmetric structures is relatively simple, when these structures have the same symmetry as one of the Fourier components of the angular symmetry decomposition. Such a symmetry is obtained when a plane wave is incident along the axis, but also when the structure is fed by a circular waveguide structure (for example, an optical fiber) along the axis. Typical examples for this are found in Scanning Nearfield Optical Microscopy (SNOM). Standard SNOMs use optical fibers with a metallic cladding and with a small tip or aperture. The fundamental fiber mode is the HE11 mode that has $\cos(\varphi)$ symmetry. This mode is much damped by the cladding [17]. In classical electromagnetics, a strong field confinement is obtained with simple multiconductor transmission lines. Such structures cannot be fabricated in the optical regime because metals lose their conductivity. However, at optical frequencies, some of the metals may be described by a complex permittivity with a negative real part, which can cause surface plasmons. When the imaginary part of the permittivity is small enough, one can use this for obtaining interesting waveguides and resonators. Figure 23 shows a concept study of a metallic resonator with a sharp tip illuminated by an axisymmetric light source. At some frequency range, an extremely bright spot is observed near the tip. The macroscopic model that was used shows no limitations concerning the spot size – provided that the radius of curvature near the tip is small enough. Therefore, the spot size will probably only be limited by atomic effects that cannot be handled with a pure Maxwell solver.

The MMP solution of such axisymmetric models is very efficient, because only a 2D section of the 3D object is modeled. More precisely, only 1D boundary lines are discretized. Therefore, one can work with small matrices of a few hundred unknowns even when the object is not small compared with the wavelength. For example, the model shown in Figure 23 has been computed with 252 unknowns. The average error is 0.05% and even at the most critical sections, the relative error is in the order of 1%.

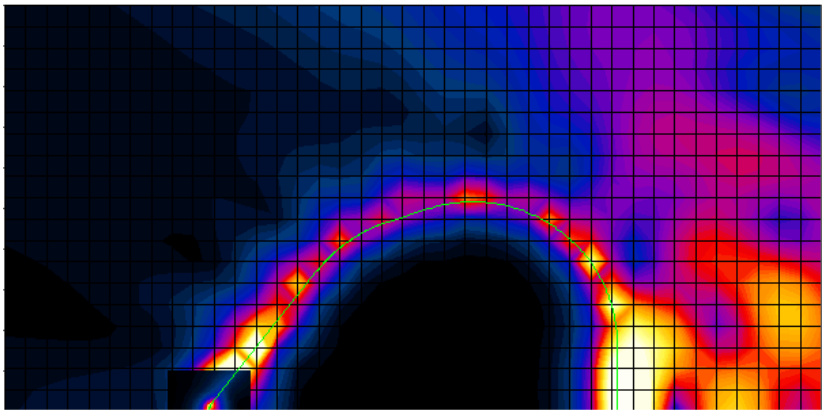

Fig. 23: Concept study of a resonant, metallic SNOM tip with an extremely localized electric field near the tip. The rotational symmetry axis is on the bottom of the figure. A special window with a finer grid and another scaling has been introduced around the tip in order to show the strength of the confinement together with the overall field distribution that is much weaker.

Since the MMP matrix is small, one can easily perform extensive studies that show the dependence on the frequency, on the illumination, on the material properties, and on the geometric shape.

Another interesting structure is the axisymmetric, circular sinusoidal grating shown in Figure 24. As one can see, it generates a nice beam when it is illuminated by a plane wave along the axis. For this example, 521 unknowns were used and a relative error of 0.2% in the average along the boundary was obtained. The maximum relative error is in the order of 1%.

Although axisymmetric problems can usually be solved with small MMP matrices it should be mentioned that the matrix setup time is much longer than the matrix setup time for cylindrical MMP models because the computation time of ring multipoles is much longer than the computation of cylindrical (2D) multipole expansions.

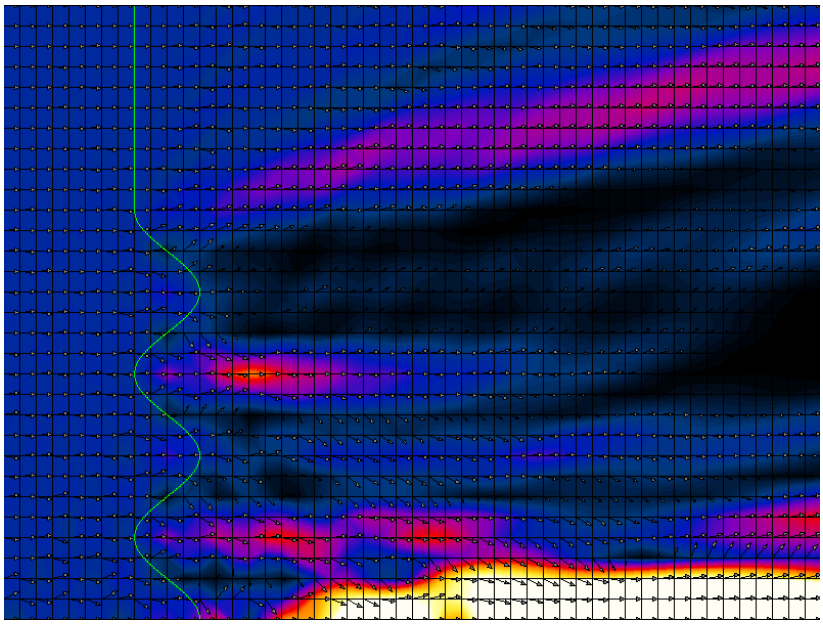

Fig. 24: Axisymmetric, finite sinusoidal grating illuminated by a plane wave incident from the left hand side. Time average of the Poynting vector field. The symmetry axis is on the bottom of the figure.

3.3. Periodic structures

It has already been mentioned that periodic structures are handled with periodic boundary condition s in the MMP implementation of MaX-1. This allows one to separate a single cell of the structure. Once this has been done, the field inside the cell may be modeled exactly as the field in any other domain. For this purpose, one can use standard multipole expansions, Bessel expansions, and all other expansions available.

A special modeling is required for gratings that may be periodic in one or two directions. Typically, one has a sandwich-like geometry. On top and bottom, one has a half space and the grating structure is somewhere in between as shown in Figure 25. When the grating is illuminated by a plane wave incident in the top half space, the field in this half space will also contain the reflected

waves that may be expanded by special harmonic expansions, so-called Rayleigh expansions or Floquet modes. Similarly, the transmitted waves in the bottom half space may be modeled with Rayleigh expansions. Such expansions are available in the MMP code for 2D and 3D models. Figure 25 shows a 2D example. An example of a 3D model with a grating that is periodic in two perpendicular directions has already been given in Figure 18.

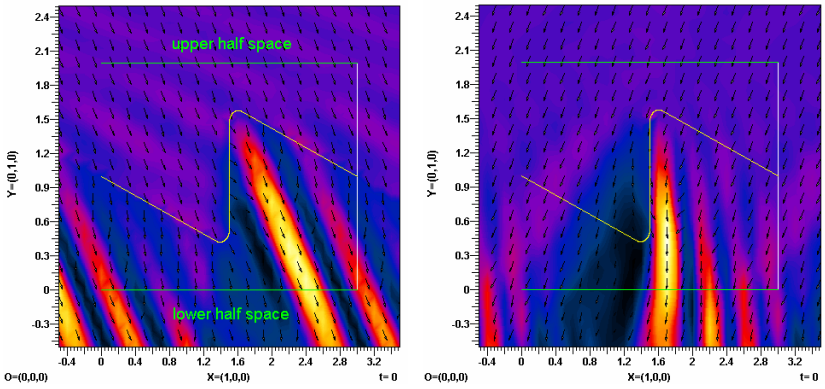

Fig. 25: 2D blazed grating illuminated by a plane wave. Left hand side angle of incidence –25°, right hand side: +25°. Time average of the Poynting vector field. Green lines: fictitious boundaries that separate the upper and lower half space, white line: fictitious boundary with periodic boundary conditions.

Note that the half spaces of a grating must usually be separated from the grating structure itself by fictitious boundaries [6, 8] because the Rayleigh expansions alone are not complete when the half space is not limited by a planar boundary.

Perfect photonic crystals are a class of periodic problems that is different from gratings because they cover the entire space. In fact, there is no excitation when perfect photonic crystals are studied, i.e., these structures are considered as eigenvalue problems whereas gratings are handled as scattering problems. Therefore, we first consider eigenvalue problems before we return to photonic crystals.

3.4. Guided waves and resonators

Guided waves and resonators are typical eigenvalue problems. The resonator problem is simpler because the resonance frequencies are the eigenvalues that are characteristic for the modes, whereas the guided waves are characterized by the frequency dependency of the propagation constants of the modes. As we have seen, the PET is well suited for speeding up the calculation of the frequency dependency. Furthermore, we have seen that different methods may be used for defining eigenvalues. Within MaX-1, several methods are available, which provide a high degree of freedom. This is important because eigenvalue problems turn out to be very tricky when the geometry is not very simple, when many modes are present, and when losses are present. In order to illustrate, we consider a so-called holey optical fiber consisting of 6 circular holes on a hexagonal lattice. Figure 26 shows the first quadrant of this structure.

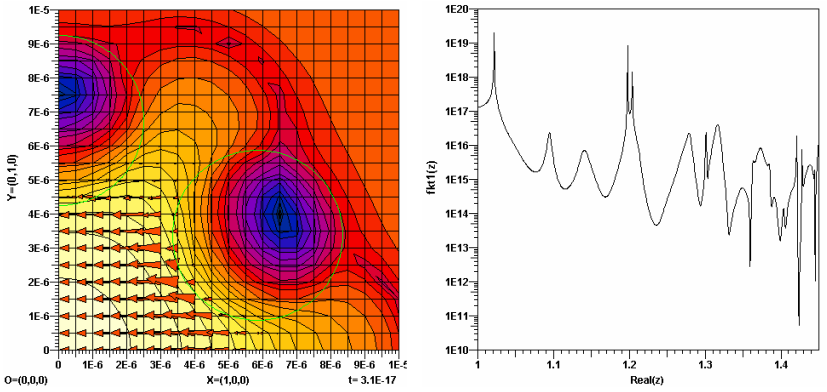

Fig. 26: Holey optical fiber. Left hand side: Electric field of the fundamental mode. Right hand side: Eigenvalue search function for the modes that exhibit the symmetry of the fundamental (pseudo HE11) mode.

The fundamental mode of the holey fiber resembles the HE11 mode of the circular step index fiber, but the classification of the modes is extremely difficult because the holey fiber does not exhibit rotational symmetry. A partial classification might be obtained from the symmetry that is still there: The holey fiber has

6 symmetrical axes, i.e., diedric D6 symmetry. The symmetry decomposition for D6 is not very difficult, but it is not trivial as well, because the corresponding symmetry group is two-dimensional. Furthermore – as most codes for computational electromagnetics – the MMP implementation of MaX-1 does not allow one to take D6 symmetry into account. It only considers up to 2 perpendicular symmetrical axes, i.e., D2 symmetry. When we only consider D2, we obtain an incomplete symmetry decomposition, but this will still allow us to distinguish four classes of modes with different symmetrical properties. This considerably simplifies the detection for the following reasons. Assume that 100 modes are present. Without symmetry consideration, one will have to perform a rough search with at least 10000 search points, because some of the modes can have very similar eigenvalues. When D2 symmetry is considered, each of the four classes of modes will have about 25 modes and the distances between them will be considerably bigger. Therefore, 4 times 1000 search points should be sufficient. Furthermore, the size of the matrix that describes the modes is considerably smaller (N/4 times M/4 instead of N times M) when D2 symmetry is considered. Therefore, the computation time for the rough search will be reduced by a factor of approximately 160.

As one can see from Figure 26, the eigenvalue search function exhibits extremely sharp minims that characterize the modes mainly near the right hand side of the diagram. In order to accurately compute the field of the corresponding mode, we must explore the corresponding minims with a high precision. For example, the rightmost minimum that corresponds to the fundamental mode goes down to 1E6.

The minims that characterize higher order modes are more flat and high. Although this simplifies the search for the corresponding modes, it also indicates that an insufficient accuracy will be reached because the orders of the multipoles and Bessel expansions are not high enough for a good approximation of these modes with their complicated field patterns.

3.5. Perfect photonic crystals

Perfect photonic crystals are structures that are periodic in up to three directions in space. For reasons of simplicity, we focus on cylindrical 2D crystals that are periodic in two directions of the *xy* plane that is perpendicular to the cylinder (*z*) axis. Such crystals usually consist of dielectric rods or holes in a dielectric, but metallic rods may also be considered [15, 16]. For the analysis of photonic crystals, methods known from crystallography are applied.

Its special symmetries are first explored. For the MMP analysis, fictitious boundaries are introduced for separating a single cell of the original lattice. Along these boundaries, periodic boundary conditions are imposed. Here, it is important to note that the fictitious boundaries may be straight lines – which seems to be natural – but curved lines are admitted as well. If one already has some idea about the field pattern of the mode to be analyzed, it is reasonable to define the fictitious boundaries in an area where the field is relatively smooth because this will minimize the model size. As shown in Figure 27, only a single original cell of the original lattice must be modeled explicitly. Since this cell often exhibits additional symmetrical axes, it would be possible to model only a part of it and to apply appropriate symmetry decomposition, but this has not been implemented in the current MMP version. As a consequence, one may encounter degenerate modes that cause some numerical problems.

In the second step, the reciprocal lattice space is introduced which is essentially obtained from a spatial Fourier analysis. This space is illustrated in Figure 28. It is spanned by the components of the wave vector. Because of the symmetry, it is sufficient to analyze the modes only over a finite part of the reciprocal space, the 1st Irreducible Brillouin Zone (IBZ) as shown in Figure 28. For the most frequently used square and triangular (hexagonal) lattices, the IBZ is a triangle.

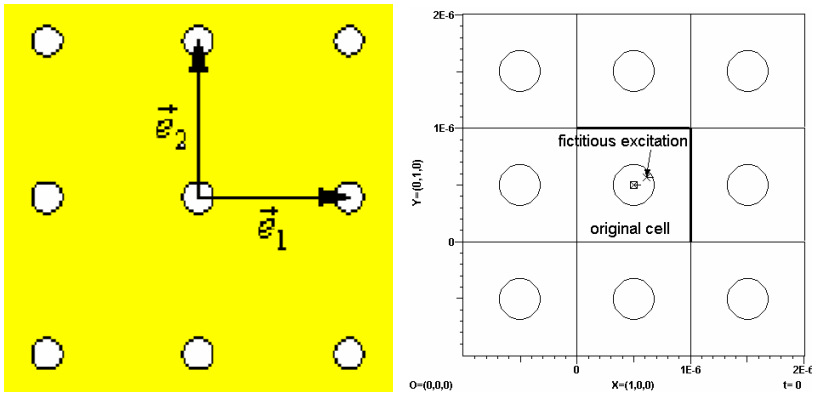

Fig. 27: Simple 2D photonic crystal with a square lattice. Right hand side: MMP-MAS model with the square original cell, periodic boundary conditions imposed along the fictitious boundaries shown by thick lines, and fictitious excitation of the MMP-MAS eigenvalue search.

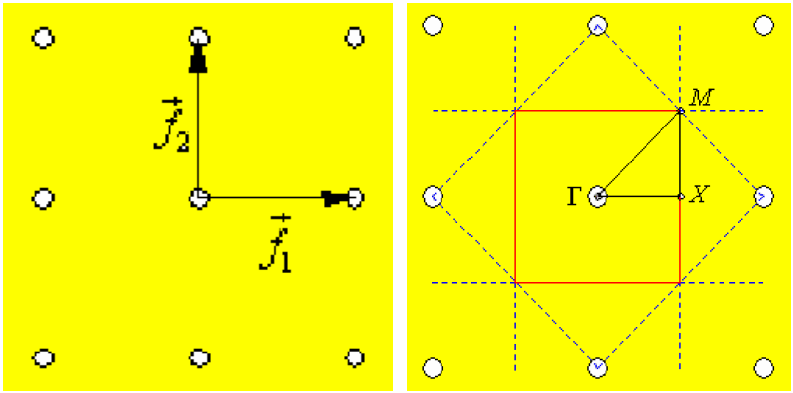

Fig. 28: Reciprocal lattice space (left hand side) and 1st Irreducible Brillouin Zone (IBZ) for the case shown in Figure 27.

In order to compute the band gap that is of highest interest for photonic crystals, it is sufficient to explore the borders of the IBZ only, i.e., to find the resonance frequencies along the lines in Figure 28 from Γ to X, from X to M, and from M back to Γ. Figure 29 shows a typical band diagram.

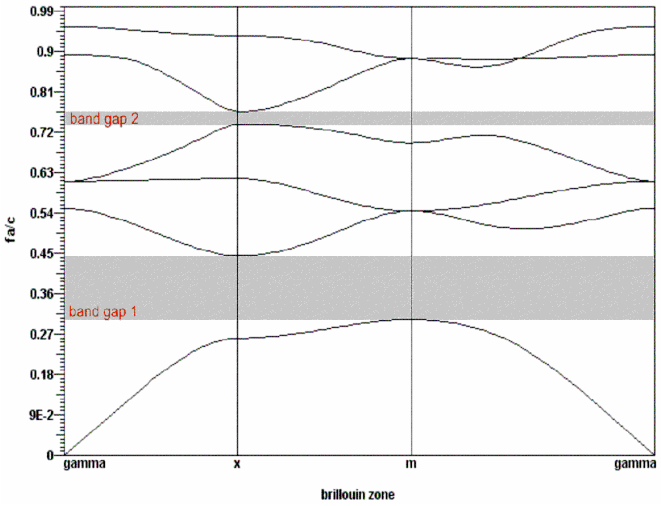

Fig. 29: Band diagram for the model shown in Figure 27.

As one can see, the band gaps are zones or frequency ranges (vertical direction in Figure 27) where no mode exists. When a finite photonic crystal is illuminated by a wave incident from an arbitrary direction, it will totally reflect the wave. This interesting and useful effect is only obtained when the dielectric contrast between the rods and the background si sufficiently high and when the rods are sufficiently big. Furthermore, the geometric shape of the rods and the symmetry of the lattice play important roles.
Since the band gaps define the working area for photonic crystals, the main reason for the computation of the band diagrams is to find crystals with a sufficiently wide bandgap around the desired frequency range. This is a rather difficult optimization task, especially because the computation of the band diagram exhibits several numerical difficulties. Fortunately, one usually works in the lower band gap. In order to find it, only the lowest order modes need to be considered. These modes have a relatively simple field and can be computed with relatively low computational effort. MaX-1 provides all tools required for the efficient computation of band diagrams and for dealing with numerical difficulties.

3.6. Waveguides in photonic crystals

Although it is nice to fabricate a photonic crystal that totally reflects all waves within a certain frequency range, defects make photonic crystals much more attractive and promising. In fact, photonic crystals were first proposed as an optical counterpart of semiconductor crystals [19]. As we know, doping makes semiconductors really interesting. Since the cells of photonic crystals are much bigger than the atoms of a semiconductor, one can "dope" photonic crystals very precisely by changing the geometry and material properties of one or several cells. For example, when a photonic crystal consists of a set of dielectric rods on a square lattice, one can remove all rods along a line [20, 21]. As a result, a simple defect waveguide is obtained as shown in Figure 22. Note that it is very easy to obtain many different types of waveguides simply by modifications of the cells of a photonic crystal along one or several parallel lines [22]. This gives one much freedom in the design of waveguides, but it also poses a problem to design and optimize waveguides with desired properties, which is a hard synthesis problem.

The standard approach to compute waveguides in photonic crystals is the so-called super cell approach [23, 24] that considers an array of parallel waveguides instead of a single one. When the waveguides are separated by sufficiently many cells of the photonic crystal and when one operates sufficiently well within the band gap, there is almost no interaction of the waves in neighboring waveguides. Therefore, one essentially obtains the same field distribution as in the single waveguide. The set of infinitely many parallel waveguides may be considered as a bigger periodic structure with an original cell that contains several cells of the former, undisturbed photonic crystal. Figure 30 shows an example of a super cell with a defect waveguide in the center and four rods on each sides of the channel, i.e., neighboring channels are separated by 8 rows of rods.

As one can see in Figure 30, each of the modes of the undisturbed crystal on both sides of the band gap splits in several modes. The

number of modes increases with the size of the super cell, i.e., with the number of rods on each side of the channel. Therefore, the computational effort is drastically increased, which makes the super cell method inefficient.

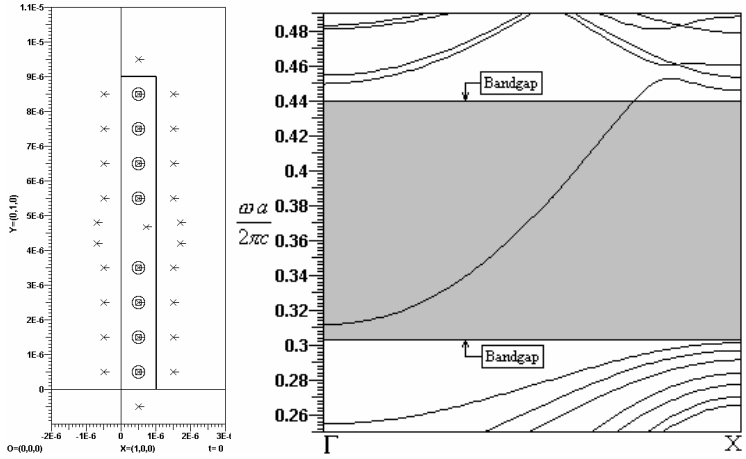

Fig. 30: MMP model of the super cell (left hand side) of a defect waveguide and band diagram of the modes near the lowest band gap of the undisturbed crystal.

Beside the split modes, one also observes a completely new mode within the original band gap in Figure 30. This is the defect waveguide mode. Note that this mode is not completely within the band gap, but it covers a wide area of it. Depending on the model parameters, one can obtain single mode waveguides as in Figure 30, but multimode waveguides as well.

Although the super cell method allows one to use the same procedures as for the analysis of perfect photonic crystals, this method is obviously neither very realistic nor numerically efficient. A more realistic waveguide model would consist of a channel with a finite number of rows of rods on both sides, i.e., one might simply remove the fictitious boundaries on top and bottom of the super cell shown in Figure 30 and keep the left and right fictitious boundaries. When one does this, the structure is no longer periodic in the vertical direction, but it is still periodic in the horizontal direction. Thus, one obtains a one-dimensional lattice

space and a one-dimensional reciprocal space as well. The reciprocal space now essentially is the space of the propagation constant of the waveguide, i.e., the wave number in horizontal direction. This makes the eigenvalue search much easier.

When one removes the fictitious boundaries that separate the super cell from the neighboring cells with the parallel neighboring waveguides, one should recognize that the resulting waveguide with finite walls may radiate in the upper and lower half space. When the number of rows of rods on both sides of the channel is big enough, one may expect that the radiation is weak and may be neglected. In order to obtain information about the radiated power, we must take it into account. Within our MMP model, we can easily do this by a model similar to the grating model that separates the upper and lower half space by fictitious boundaries. Therefore, we can re-introduce the fictitious boundaries on top and bottom of Figure 30. Instead of imposing periodic boundary conditions, we match all field components and model the radiated field in the top and bottom half space with simple multipole expansions.

The radiation losses have an important consequence: the propagation constant now must be complex even if all dielectrics are loss-free. This requires a complex eigenvalue search instead of a real search. When the radiation loss is not very high, one can first ignore it and perform a real eigenvalue search. As soon as the field of the mode is known, one can also compute the radiated power and the total power transmitted along the channel. From these data, one can easily approximate the attenuation constant of the mode.

In order to illustrate the procedure, we consider a defect waveguide with only one row of rods on each side of the channel as shown in Figure 31. Obviously such a waveguide must exhibit the strongest possible radiation. Therefore, we perform a complex eigenvalue search. In order to get an overview over the eigenvalue search function in the complex plane (see Figure 32), we start the rough search on a rectangular grid defined on the complex plane. When we want to have a fine resolution, this is time-consuming. Note that this is not necessary in general because there are not many

different eigenvalues of interest, i.e., eigenvalues with a reasonably small propagation constant. Therefore, we could also start with a rough search along the real axis and perform the fine search from the resulting start points into the complex plane, which would be numerically much more efficient.

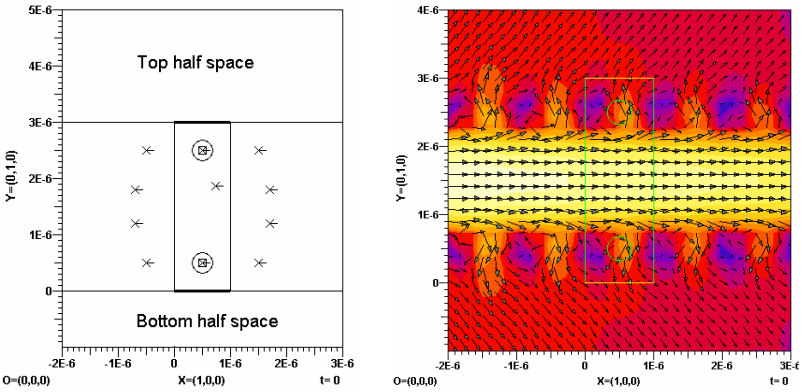

Fig. 31: MMP model of a defect waveguide with only one row of rods on each side of the channel. Right hand side: Time average of the Poynting vector field of the fundamental mode. Note that a logarithmic scale is used here, because the field is much weaker outside the waveguide. With linear scaling the area outside would be almost black.

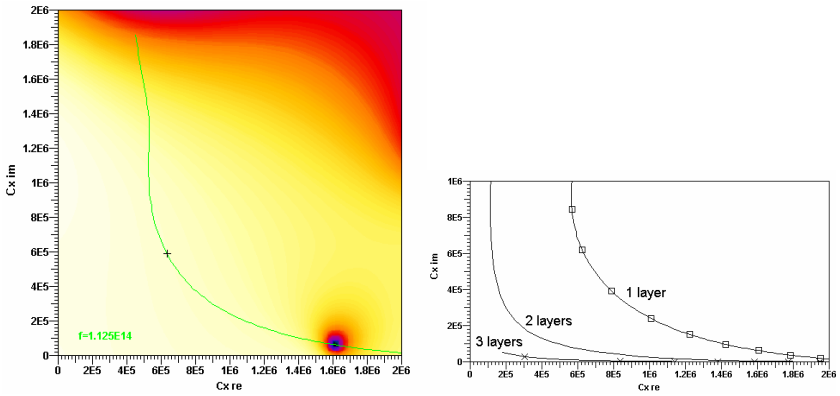

Fig. 32: Left: MMP eigenvalue search function in the complex plane. The minimum that characterizes the mode is near the lower right corner. The green line is the trace of the eigenvalue for varying frequency. The black + marker shows the lower end of the band gap. The upper end of the band gap is out side the clip area. Right: Complex eigenvalue trace for waveguides with 1, 2, 3 layers of rods on each side of the channel.

As soon as we have found the desired eigenvalue that characterizes the fundamental (radiating) mode, we may compute and plot its field (see Figure 31) and we may track its frequency dependency in the complex plane using the PET (see Figure 32). As one can see, the attenuation constant (vertical direction in Figure 32) drastically grows when the frequency is reduced and it is quite big when the lower end of the band gap is reached. Furthermore, we can trace the mode also outside the band gap although this makes not much sense for practical applications. However, we observe that the behavior of these photonic crystal waveguide modes is similar to the behavior of lossy fiber modes [17] although the former exhibit radiation loss whereas the latter exhibit Ohmic loss in the coating. From the practical point of view, it is important to know that one should not operate near the lower limit of the frequency band in order to avoid strong attenuation due to strong radiation loss. Of course, we may drastically reduce the radiation loss by adding one or several rows of photonic crystal cells on both sides of the channel (see Figure 32).

When the radiation loss is not very high, the complex eigenvalue search can be avoided. When we first assume that the attenuation constant is zero, we can perform a real eigenvalue search (along the real axis of the complex plane) and obtain a real valued approximation of the complex propagation constant, i.e., an approximation of the phase constant. As one can see from Figure 33, a good approximation is obtained even in the worst case with a single layer of rods on both sides of the channel. As soon as the phase constant is known, one can also compute the electromagnetic field and from this, one can estimate the power loss along the guide due to radiation. From this, one can easily obtain an estimate of the attenuation constant that is also shown in Figure 33. As one can see, the estimation is quite good - even for the worst case - over the entire frequency range of the band gap. Obviously, the error is higher for lower frequencies. It should be mentioned that the defect mode considered here has a cutoff near the lower end of the band gap. One can see this from models with a sufficiently large number of layers on both sides of the channel. Note that the

cutoff frequency of any waveguide mode is strictly defined only in the loss-free case. When Ohmic losses or radiation losses are present, some sort of smooth transition is observed as shown in Figure 32.

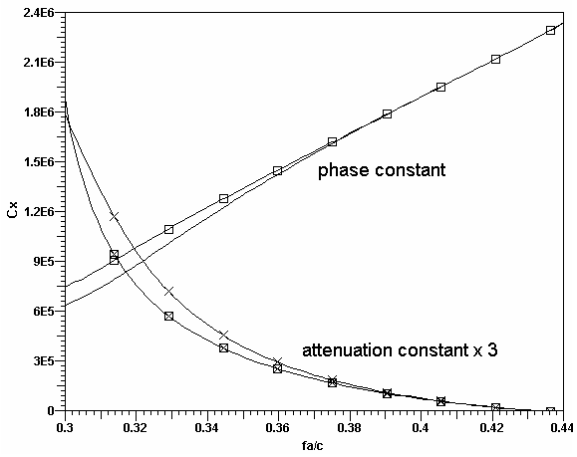

Fig. 33: Frequency dependency of the real and imaginary (x markers) part of a radiating defect waveguide with only one layer of rods on each side of the channel. The square markers indicate the results obtained from a real eigenvalue search with an estimation of the attenuation constant obtained from the radiated power computation.

3.7. Discontinuities in photonic crystals

Although it is important to have waveguides in photonic crystals, the existence of such waveguides would not make the photonic crystal concept sufficiently attractive for fabrication although such waveguides may be smaller than traditional waveguides used in integrated optics. The photonic crystal concept shows its huge potential when one studies discontinuities of waveguides. Theoretically, periodic discontinuities might be studied with the super cell method. As we have already seen, this method is not efficient at all. Therefore, we search for an alternative approach.

Although finite photonic crystal waveguide discontinuities are somehow different from the discontinuities of traditional waveguides, we may use the same concept for a study with MMP [25, 26]: 1) Essentially all waveguide discontinuity problems consist of at least one input waveguide and usually of one or several output waveguides. These waveguides are assumed to extend to infinity. The modes in these waveguides can be analyzed using the MMP eigenvalue solver. 2) The waveguide sections may be separated with fictitious boundaries from the discontinuity region. The fictitious boundaries play the role of reference planes known from measurement. Let us call them ports. In each port one has a finite number of propagating modes and an infinite number of evanescent modes. The field in the discontinuity region can be modeled by a standard MMP expansions. It is linked to the waveguide modes by the fictitious boundaries at the ports. The field inside the discontinuity region is excited by the incoming modes and it excites the outgoing modes. In MMP, fictitious boundaries and connections allow one to easily link the field with the modes (see Figure 34).

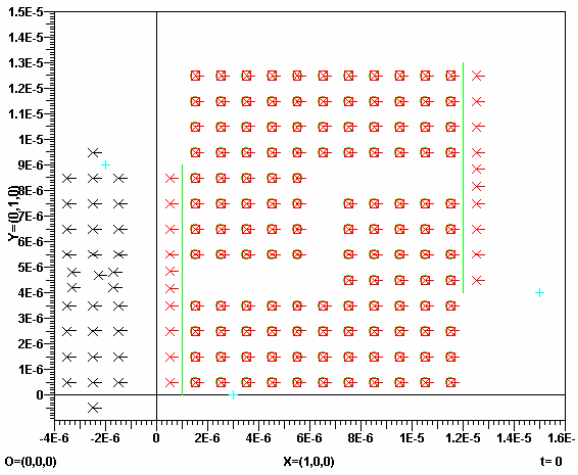

Fig. 34: Defect waveguide with two sharp 90^{o} bends. The green lines indicate fictitious boundaries that describe the input (left hand side) and output ports. Multipoles are indicated by crosses. The black crosses indicate the multipoles used for the connections that describe the waveguide modes.

Since the evanescent modes are excited in the discontinuity region and decay exponentially, one can omit them when the waveguide port is sufficiently far away from the discontinuity. Moving ports away from the discontinuity increases the size of the discontinuity region that is explicitly modeled and this considerably increases the computation time. Therefore, it is sometimes more reasonable to move the ports close to the discontinuity region, where the lowest evanescent modes are still strong. If this is done, the field in the waveguides must be approximated by a superposition of the N_g guided waves and of the first N_e evanescent modes. As we have seen, the MMP eigenvalue solver can find all guided waves of an infinite waveguide, even when it is lossy. Since MMP can search eigenvalues in the complex plane, it also can find evanescent modes. Once a waveguide mode is known, one can pack the expansions that model its field in a connection. One then can model the field in each port with $N_g + N_e$ connections. This allows one to match the field at the ports with a high precision in such a way that almost no reflections at the ports are observed. Figure 35 illustrates that for a simple case.

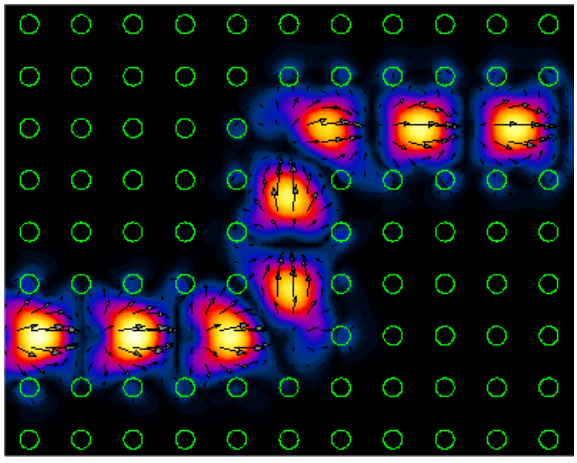

Fig. 35: Poynting vector field of a defect waveguide with two sharp 90^o bends.

As soon as one knows how to handle waveguide discontinuities, one can start designing interesting structures in photonic crystals, for example, multiplexers [27], resonators, filters [28], etc. It is well known that the device synthesis is much more demanding than the analysis of an existing device or structure. The main problem in the photonic crystal design is that almost no design rules are currently available. For example, we know that we can obtain filters, but we don't have rules how to design a filter with specified characteristics. In order to illustrate a possible design procedure, we now consider a filtering T junction with one input port and two output ports [27]. We would like to design this structure in such a way that it transmits almost all energy from the input port to the first output port at one frequency f_1 and to the second output port at another frequency f_2.

First, we select a perfect photonic crystal with a sufficiently wide bandgap that contains both frequencies f_1 and f_2. Since we would like to have waveguides in two perpendicular directions, we use a square lattice. Furthermore, we consider a crystal consisting of dielectric rods rather than of holes in a dielectric. Note that the latter would be better for some practical reasons, but the former is numerically simpler and we already used it in the examples shown above.

In the second step, we design three different single mode waveguides that shall be used for the three ports. When the input waveguide propagates at all frequencies of the entire band gap, whereas the first output port mode only propagates within a limited frequency band around f_1 but not at f_2, we can expect that most of the power is transmitted to the first output port at f_1. Similarly, we can design the second output waveguide in such a way that most of the power is transmitted to the second output port at f_2. Figure 36 shows results obtained with this concept.

Although the design with three different waveguides is promising, we do not want to have structures with different waveguide types. When we want a filtering T junction with three identical waveguide ports, we can use the model above and connect the two

output waveguides with waveguides of the same geometry as the input waveguide after a certain length.

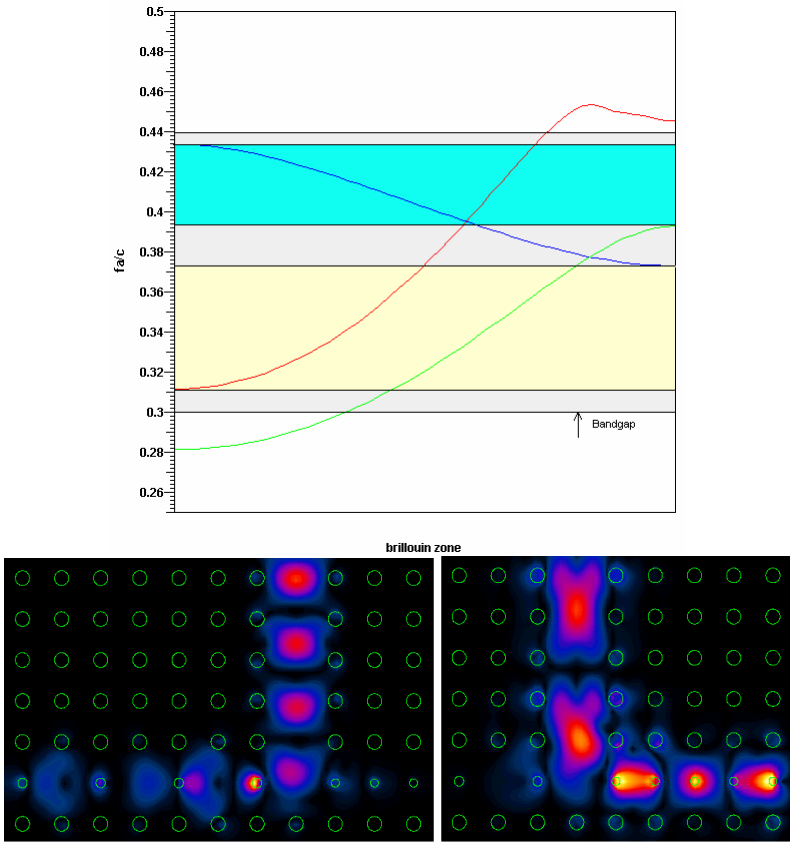

Fig. 36: Filtering T junction with three different waveguides. Top: Waveguide mode 1) red: input port (from top), 2) blue: left port, 3) green: right port. Bottom: Poynting vector plots for the two transmission frequencies.

It is obvious that this will cause some additional reflections at the transitions between the different waveguides and that the quality of the T junction will considerably depend on the lengths of the two output waveguides. Since the main goal of photonic crystals is miniaturization, we try to make the T junction as small as possible. Figure 36 shows such a T junction that is obtained after some optimization of the two output sections. The main ideas of this

optimization are that the propagation constants of the modes of the input and first output section should be the same at f_1. Furthermore, the group velocities of these modes should be similar at f_1. The same shall hold for the second output section at f_2. Finally, we know that the evanescent field of the undesired mode will reach far into the wrong output port when one operates close to the cutoff frequency. Therefore, the cutoff frequency of the second output mode should not be close to f_1 and vice versa. Figure 37 shows the resulting structure.

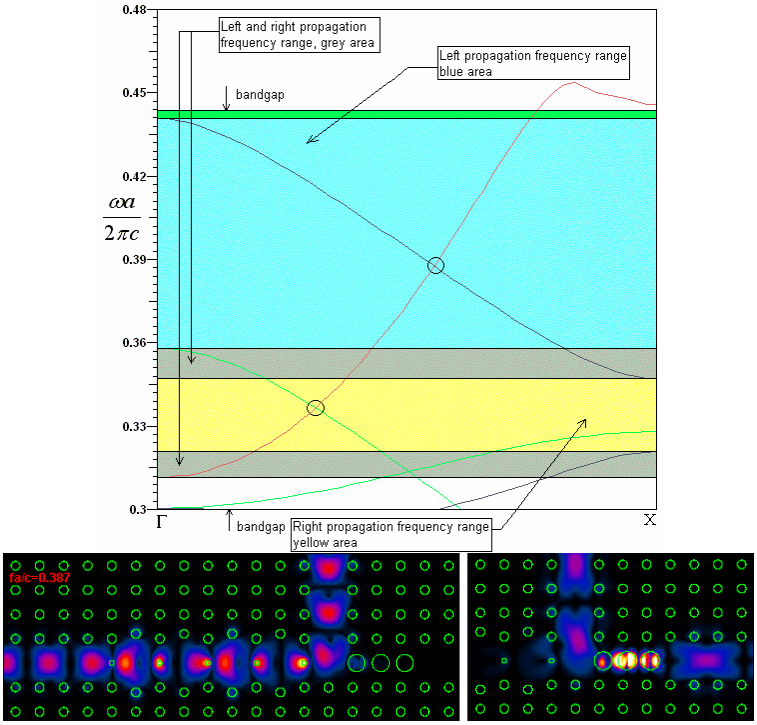

Fig. 37: Filtering T junction with identical waveguide ports. Top: Waveguide mode 1) red: input port (from top), 2) blue: left section, 3) green: right section. Bottom: Poynting vector plots for the two transmission frequencies.

With this simple design, a rather good performance (approximately 89% transmission to port 1 for f_1 and 93% to port 2 for f_2) is obtained with a simple tuning of the waveguide sections. In order

to obtain even better performance, one must either increase the size of the T junction or one must perform more comprehensive optimization of the entire T junction area, which is numerically very time consuming.
Obviously, the left channel that is similar to a coupled cavity waveguide structure is geometrically more complicated and also considerably longer than the right channel that simply consists of three thick rods. In order to obtain a more compact design, we replace the left waveguide section by a waveguide section similar to the right waveguide section as shown in Figure 38.

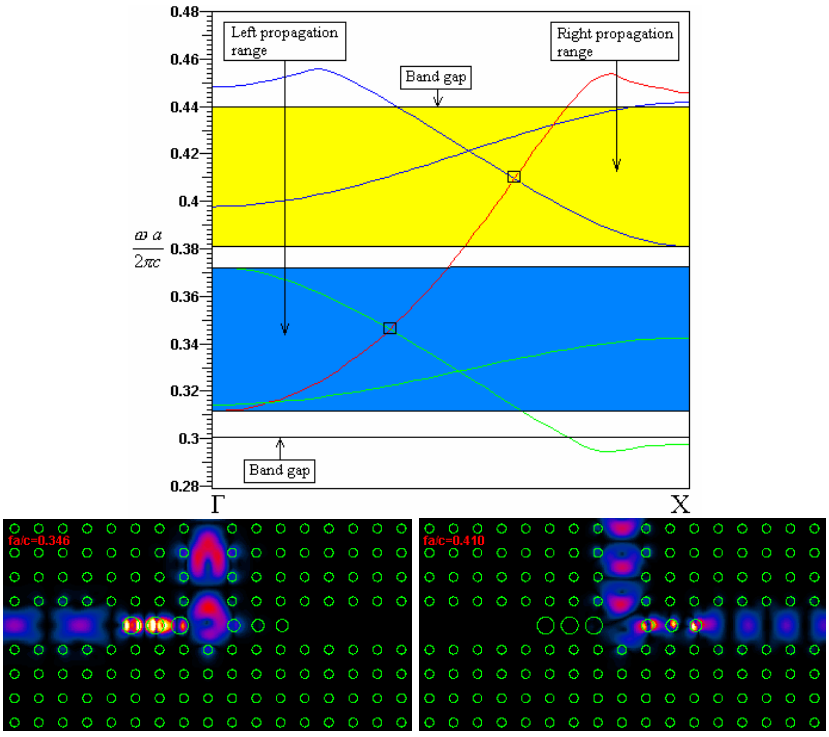

Fig. 38: Compact filtering T junction. Top: Waveguide mode 1) red: input port (from top), 2) blue: left section, 3) green: right section. Bottom: Poynting vector plots for the two transmission frequencies.

The new structure is considerably more compact, but its performance is considerably worse. We now try to improve it with a simple optimization of the transition region. In fact, there would be many model parameters (size and positions of all rods in the transition area) that might be optimized, but such an optimization in a high-dimensional parameter space would be too time-consuming for our personal computers. Therefore, we select only two important parameters and optimize them. These parameters are the distances of the rods in the two output sections. Since the minimum reflection coefficients and the maximum transmission coefficients for the two output ports at f_1 and f_2 are obtained for different values of the two distances, we must find some compromise. The resulting design is shown in Figure 39. This optimization allows us to increase the transmission from 63% (Fig. 38) up to 88% (Fig. 39) for both channels, i.e., the new T junction is almost as good as the one shown in Figure 37, but it is much more compact and easier to fabricate.

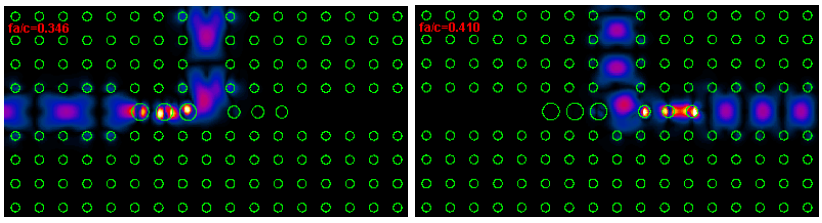

Fig. 39: Compact filtering T junction optimized. Poynting vector plots for the two transmission frequencies.

When we optimize a photonic crystal structure, we see that some of the model parameters may be rather critical. In our filtering T junction, the size and locations of the rods in the transition area are quite critical. This means that a precise fabrication is required. Otherwise the performance of the junction might be considerably worse. Although this may cause problems for the fabrication, it is also interesting for tuning. It is well known, that photonic crystals may be tuned [29, 30], for example, with an external electrical field, when appropriate materials are introduced. When the

material properties in the transition area are modified or when the sizes of the rods in this area are modified, essentially the same effects are obtained. Therefore, it should be possible to tune the properties of the filtering T junction with external electrical field, i.e., we may obtain also optical switches [31] with a similar, compact T-junction design.

Conclusions

We have presented a powerful method for the accurate and efficient full-wave analysis of electromagnetic waves that is very useful for the analysis and design of novel structures for ultra dense integrated optics, namely for photonic crystal structures. Furthermore, we have presented several technique for speeding up the analysis of sequences of similar models and for obtaining efficient and precise eigenvalue solvers. Due to its high reliability and accuracy, our code is very well suited for future optimizations of novel photonic crystal structures.

Acknowledgments

This work was supported by the Swiss National Science Foundation.

References

1. Ch. Hafner, (in German) (Diss. ETH, No. 6683, Zurich, 1980).
2. G. Mie, *Annalen der Physik.* , 2 (1900).
3. G. Mie, *Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen*. (Ann. Phys. 25, 377–445, 1908, An English translation is available as G. Mie, *Contributions to the optics of turbid media, particularly of colloidal metal solutions*, Royal Aircraft Establishment library translation 1873. Her Majesty's Stationery Office, London, 1976).
4. I. N. Vekua, *New methods of solution of elliptic equations. (Russian)* (Gostekhizdat, Moscow-Leningrad, 1948).
5. I. N. Vekua, *New methods of solution of elliptic equations.* (North-Holland Publ. Co., Amsterdam, 1967).
6. Ch. Hafner, *Post-modern Electromagnetics Using Intelligent MaXwell Solvers.* (John Wiley & Sons, 1999).

7. Ch. Hafner, *The GGP Code Web Page.* (IFH, ETH Zürich), http://alphard.ethz.ch/ggp .
8. Ch. Hafner, *MaX-1: A Visual Electromagnetics Platform.* (John Wiley & Sons, 1998).
9. F. G. Bogdanov, D. D. Karkashadze, and R. S. Zaridze, *Generalized Multipole Techniques for Electromagnetic and Light Scattering.* (edited by T. Wriedt, pp. 143-172, Elsevier, Amsterdam, 1999).
10. G. Tayeb, *The Method of Fictitious Sources Applied to Diffraction Gratings.* (Applied Computational Electromagnetics Society Journal, Vol. 9, No. 3, 1994).
11. V. Rokhlin, *Rapid solution of integral equations of scattering theory in two dimensions*. (J. Comp. Phys. 86, 414-439. 1990).
12. G. H. Golub and C. F. Van Loan, *Matrix Computations.* (Johns Hopkins University Press, Baltimore, 1996).
13. M. R. Hestenes, E. Stiefel, *Methods of Conjugate Gradients for Solving Linear Systems.* (Journal of Research of the National Bureau of Standards, Vol. 49, No. 6, December 1952).
14. J. Smajic, Ch. Hafner, D. Erni, *Automatic Calculation of Band Diagrams of Photonic Crystals Using the Multiple Multipole Method.* (ACES Journal, to be published).
15. K. Sakoda, N. Kawai, T. Ito, A. Chutinan, S. Noda, T. Mitsuyu, K. Hirao, *Photonic bands of Metallic Systems. I. Principle of Calculation and Accuracy.* (Phys. Rev. B, 64, pp. 045116, 2001).
16. E. Moreno, D. Erni, Ch Hafner, *Band Structure Computations of Metallic Photonic Crystals with the Multiple Multipole Method.* (Phys. Rev. B, 65, pp. 155120, 2002).
17. L. Novotny, *Light Propagation and Light Confinement in Near-Field Optics* (Diss. ETH, No. 11420, Zurich, 1996).
18. Ch. Hafner, J. Smajic, *The Computational Optics Group Web Page.* (IFH, ETH Zürich), http://alphard.ethz.ch .
19. E. Yablanovich, *Inhibited SpontaneousEemission in Solid-StatePhysiscs and Electronics.* (Phys. Rev. Lett. 58, pp. 2059-2062, 1987).
20. E. Centeno, D. Felbacq, *Guiding Waves with Photonic Crystals*. (Opt. Commun. 160, 57, 1999).
21. H. Benisty, *Modal Analysis of Optical Guides With Two-Dimensional Photonic Band-Gap Boundaries.* (J. Appl. Phys. 79, pp. 7483-7492, 1996).
22. M. Loncar, J. Vuckovic, A. Scherer, *Methods for controlling positions of guided modes of photonic crystal waveguides*. (J. Opt. Soc. Am. B 18, pp. 1362-1368, 2001).
23. K. Sakoda, *Optical Properties of Photonic Crystals.* (Springer, Berlin, 2001).
24. J. D. Joannopoulos, R. D. Meade, J. N. Winn, *Molding the Flow of Light.* (Princeton University Press, 1995).

25. E. Moreno, D. Erni, Ch. Hafner, *Modeling of Discontinuities in Photonic Crystal Waveguides With the Multiple Multipole Method.* (Phys. Rev. E, 66, pp. 036618, 2002).
26. E. Moreno, *MMP Modeling and Optimization of Photonic Crystals, Optical Devices, and Nanostructures.* (Diss. ETH, No. 14553, Zurich, 2002).
27. J. Smajic, Ch. Hafner, D. Erni, *On the Design of Photonic Crystal Multiplexers.* (Optics Express, to be published).
28. S. Fan, P. R. Villeneuve, J. D. Joannopoulos, H. A. Haus, *Channel Drop Filters in Photonic Crystals.* (Opt. Express 4, pp. 4-10, 1998).
29. K. Busch, S. John, *Liquid-Crystal Photonic-Band-Gap Materials: the Tunable Electromagnetic Vacuum.* (Phys. Rev. Letters, 83, pp. 967-970, 1990).
30. A. Figotin, Y. A. Godin, *Two-Dimensional Tunable Photonic Crystals.* (Phys. Rev. B, 57, pp. 2841-2848, 1998).
31. R. L. Sutherland, V. P. Tondiglia, L. V. Natarajan, S. Chandra, *Switchable Orthorombic F Photonic Crystals Formed by Holographic polymerisation-Induced Phase Separation of Liquid Crystal.* (Optics Express, 10, pp. 1074-1082, 2002).