12

Effective Index Method

Because the computational complexity of a beam-propagation problem increases sharply from one to two transverse dimensions it would be good to have at least an approximate method applicable in quasi-one dimensional geometries. Propagation in planar waveguides and also in some integrated optical devices are examples of situations in which such a simplification could be possible. In a planar waveguide, the optical field is localized in the dimension normal to the plane of the structure. It propagates confined to one or more guided modes in this direction, while the beam is free to diffract in the dimension parallel to the waveguide plane. The effective index method has been designed for these situations. It makes it possible to eliminate one transverse dimension based on the assumptions that independently of the dynamics experienced by the propagating beam, the "vertical" profile of the field is always given by the fundamental mode of the waveguide.

12.1 Propagation in a planar waveguide

We will first consider beam propagation within a planar waveguide. To fix our frame of reference, assume that the refractive index $n(x)$ of the whole structure only depends on coordinate x, and that the beam propagation direction is z. Let us further assume that the geometry and material properties are such that guided modes confined to the vicinity of $x = 0$ exist. We will denote by $\boldsymbol{\mathcal{E}}_m(x)$ the electric field modal function that corresponds to the fundamental guided mode. Also let $\beta_m(\omega) = k^2 n_{\text{eff}}^2$ be the (frequency dependent) propagation of this mode. Obviously, $\mathcal{E}_m$ can be obtained as a solution of a one-dimensional Helmholtz-type problem with appropriate boundary conditions along material interfaces (see the Exercises below). As a consequence of the symmetry, guided modes come in two polarizations, namely TE (in which it is the electric field amplitude that is parallel to the waveguide plane) and TM (in which case the magnetic field amplitude is parallel to the waveguide). It is not important for our proposes in this section which is the polarization of the beam under consideration, but we suppose that the polarization does not mix TE and TM modes.

Since "all" beam propagation equations can be derived from the Helmholtz equation, it will be our point of departure. Our goal is to eliminate the"vertical" dimension x from the consideration. For simplicity, let us consider the TE or TM polarization and use scalar notation with appropriate boundary conditions at material interfaces:

$$\left[\partial_{xx} + \partial_{yy} + \partial_{zz} + k^2 n(x)^2\right] E(x, y, z) = 0 \ . \tag{12.1}$$

Next, assume that the field can be written as

$$E(x,y,z) = \mathcal{E}_m(x)\mathcal{E}(y,z) \ , \tag{12.2}$$

and recall that the modal field alone obeys the following eigenvalue equation

$$\left[\partial_{xx} + k^2 n(x)^2 - \beta_m^2\right] \mathcal{E}_m(x) = 0 \ . \tag{12.3}$$

Inserting this, together with the ansatz in the Helmholtz, we obtain

$$\left[\partial_{yy} + \partial_{zz} + k^2 n_{\text{eff}}^2\right] \mathcal{E}(y,z) = 0 \ . \tag{12.4}$$

This is nothing but a two-dimensional equation in a "medium" characterized by the effective refractive index. This simple procedure can be repeated even if the properties of the planar waveguide depend slowly on the remaining spatial coordinates. The problem is thus reduced by one spatial dimension, and the resulting reduced Helmholtz equation can serve as a starting point to derive desired beam propagation equations along the same lines as for three spatial dimensions. The waveguide properties only enter the resulting BPM equations through the effective index.

Naturally, this method can be straightforwardly generalized for superpositions of several guided modes. It then results in a system of BPM equations that may be mutually coupled if, for example, nonlinear interaction mediate the interaction. Such coupling must be weak, otherwise the underlying assumption of the effective index method would not hold, since the interaction would "destroy" the very modes on which it is based. This approach is is in fact often utilized for slab-geometry waveguides filled with gasses.

In brief, the effective index method provides a simple approximation to simulate near-planar geometries, when the whole machinery of BPM can be applied as if in a two-dimensional world with an index of refraction replaced by its effective value. An example of an application is simulation of edge-emitting semiconductor lasers, especially the broad-area types. There, effective-index BPM versions can be utilized for static modal calculations and/or as part of time-resolving simulators.

Exercise: Derive equations for the fundamental mode of a symmetric slab waveguide.

Solution: Note that this solution presents a simplified treatment aiming solely for the fundamental mode, and as such "leaves out" many other solutions.

Let the waveguide core has the width of 2ρ, and the refractive indices are n_{co} and n_{cl} in the core and cladding, respectively. As usual, let k stands for the vacuum wavenumber corresponding the given wavelength λ.

It is convenient to choose a frame of reference such that the waveguide is parallel to plane y, z, and its center is positioned at $x = 0$. It will be further assumed that propagation is in direction of z. Then, we are looking for a solution to the Helmholtz equation in the form

$$\boldsymbol{\mathcal{E}} = e^{i\beta z}\boldsymbol{E}(x).$$

Due to the symmetry of the problem, and properties of the fundamental mode, we only need even solutions, such that $\boldsymbol{E}(x) = \boldsymbol{E}(-x)$. This means that considerations can be restricted to positive values of x.

Inserting the ansatz in into Helmholtz, gives the following equations to be satisfied by the modal amplitude $\boldsymbol{E}(x)$:

$$\rho^2 \frac{d^2\boldsymbol{E}(x)}{dx^2} + U^2\boldsymbol{E}(x) = 0 \quad , \quad 0 \le |x| < \rho$$

$$\rho^2 \frac{d^2 \boldsymbol{E}(x)}{dx^2} - W^2 \boldsymbol{E}(x) = 0 \quad , \quad \rho \leq |x| < \infty$$

where the core modal parameter is

$$U = \rho \sqrt{k^2 n_{co}^2 - \beta^2}$$

and the cladding parameter reads

$$W = \rho \sqrt{\beta^2 - k^2 n_{cl}^2}$$

Obviously,

$$U^2 + W^2 = V^2 \quad , \quad V = k\rho \sqrt{n_{co}^2 - n_{cl}^2}$$

In these expressions, it is β that is our unknown to be determined.

Utilizing the symmetry of the problem further, we can split tho solution into independent TE and TM polarizations. In the TE polarization, the only nonzero electric field components is E_y. In the TM polarization, both E_x and E_z are nonzero. This follows from the Maxwell equations when written explicitly in the component form, but is is also easy to see when one considers the divergence constraint our modal field must obey. Clearly, if $E_y(x)$ is the only component, and because it does not depend on y, divergence of such field is zero as it must be. In case of TM mode, zero divergence implies

$$\partial_x E_x(x) + i\beta E_z(x) = 0 \ ,$$

which can be used to calculate the longitudinal component once $E_x(x)$ is known.

In both of the sub-problems, we use boundary conditions for the electric field at a material interface to fully determine the thought after solutions. Let us start with the TM solution. The in-core portion can be written as

$$E_x = \frac{\cos(Ux/\rho)}{\cos(U)} ,$$

which is normalized such that its value at the interface is one. This clearly satisfied the in-core equation. The solution in cladding must then look like this

$$E_x = \frac{n_{co}^2}{n_{cl}^2} \frac{\exp(Wx/\rho)}{\exp(W)} .$$

This satisfies the in-cladding equation, and also fulfills the boundary condition that requires that the normal component of electric induction is continuous, $D_x(\rho^-) = D_x(\rho^+)$. Now all that is needed is an additional equation to relate waveguide parameters U and W, and thus "select" the propagation constant β. This condition is obtained from the divergence constraint. Since E_z is continuous because it is parallel to the interface, so must be the normal derivative of E_x,

$$\partial_x E(\rho^-) = \partial_x E(\rho^+) \ .$$

Inserting the above ansatz, this gives the eigenvalue equation

$$U \tan U = \frac{n_{co}^2}{n_{cl}^2} W$$

which restricts allowed values of β.

The procedure is similar for the TE mode. Now

$$E_y = \frac{\cos(Ux/\rho)}{\cos(U)} \quad , x < \rho$$

and

$$E_x = \frac{\exp(Wx/\rho)}{\exp(W)} \quad x > \rho \ .$$

Continuity is satisfied as it should for a field parallel to the material boundary. The remaining piece is the eigenvalue equation for β. As in the TM case, boundary condition for the normal derivative must be taken into account. Even without looking at explicit Maxwell equations, one can tell that also in this case the normal derivative of E_y must be continuous. Indeed, $\partial_x E_y(x)$ is coupled to magnetic fields through the Maxwell equations. Since the magnetic components are continuous across the waveguide interface, so must be $\partial_x E_y(x)$. Writing this continuity condition explicitly results in the eigenvalue equation for the propagation constant:

$$U \tan U = W$$

The eigenvalue equations must be solved numerically. Since they in general have multiple solutions, care must be exercised to identify the one that corresponds to the fundamental mode. The propagation constant of the fundamental mode is the largest allowed, and results in modal field that changes slowest w.r.t. transverse coordinate x. In other words it has a small cladding parameter U. So to look for the numerical solution, one can start from small U and search upward — the fundamental mode solution will be the first found.

Exercise: Generalize the results of the previous exercise for the case of a planar waveguide with different refractive indices of substrate and cladding media.

Solution: Only solution for the TE mode is sketched next. This mode has only one non-zero electric field component. With the waveguide and frame of reference oriented as in the previous Exercise, it is the E_y polarization component. Since the symmetry with respect to x is lost if the substrate and cladding refractive indices are different, let us shift the frame of reference such that the material interface between substrate and core is located at $x = 0$. The core total thickness will be denoted by $2a$.

The corresponding Helmholtz equation is

$$\frac{d^2 E_y}{dx^2} + k^2(\epsilon(x) - n_{\text{eff}}^2)E_y = 0 \quad , \quad \beta = n_{\text{eff}} k_0 \ .$$

Let us denote the respective refractive indices of the substrate, core, and cladding by n_1, n_2 and n_3. Also let

$$\gamma_1 = k_0\sqrt{n_{\text{eff}}^2 - n_1^2} \quad , \quad \gamma_2 = k_0\sqrt{n_2^2 - n_{\text{eff}}^2} \quad , \quad \gamma_3 = k_0\sqrt{n_{\text{eff}}^2 - n_3^2}$$

be the transverse wavenumbers of the solution in the substrate, core, and cladding, respectively. The effective index n_{eff} will attain a value for which that all three are real-valued. With these wavenumbers, the ansatz for the solution can be taken in the form

$$E_y(x) = C_1 \exp[\gamma_1 x] \quad , \quad x < 0 \ ,$$

$$E_y(x) = C_2 \cos[\gamma_2 x + \alpha] \quad , \quad 0 < x < 2a \ ,$$

$$E_y(x) = C_3 \exp[-\gamma_3 (x - 2a)] \quad , \quad x > 2a \ .$$

This ensures that the Helmholtz equation is satisfied, and that for large values of x, positive and negative, the field decays exponentially as it should for a guided mode.

There are four so far undetermined constants, which will be fixed by requirement of continuous field and its derivative w.r.t. x. The corresponding continuity equations imply

$$\begin{aligned} C_1 &= C_2 \cos\alpha \\ -\gamma_1 C_1 &= \gamma_2 C_2 \sin\alpha \\ C_2 \cos[2\gamma_2 a + \alpha] &= C_3 \\ -\gamma_2 C_2 \sin[2\gamma_2 a + \alpha] &= -\gamma_3 C_3 \end{aligned}$$

Solubility condition determines the value of the effective index n_{eff}. To obtain the corresponding eigenvalue equation, divide the first two equations to get α:

$$\alpha = -\arctan\left[\frac{\gamma_1}{\gamma_2}\right] + m_1\pi \quad , m_1 = 0, 1, 2, \ldots$$

Dividing the remaining two equation leads to

$$2\gamma_2 a = \arctan\left[\frac{\gamma_3}{\gamma_2}\right] - \alpha + m_2\pi \quad , m_2 = 0, 1, 2, \ldots$$

Elimination of α gives the thought eigenvalue equation

$$2\gamma_2 a = \arctan\left[\frac{\gamma_3}{\gamma_2}\right] + \arctan\left[\frac{\gamma_1}{\gamma_2}\right] + m\pi \quad , m = 0, 1, 2, \ldots$$

As it should, this equation reduces to that derived for the symmetric waveguide when $n_1 = n_3$.

For the TM mode, normal derivatives remain continuous, but there is a discontinuity of the field itself. This only results in a minor modification of the above method and gives

$$2\gamma_2 a = \arctan\left[\frac{n_2^2\gamma_3}{n_3^2\gamma_2}\right] + \arctan\left[\frac{n_2^2\gamma_1}{n_1^2\gamma_2}\right] + m\pi \quad , m = 0, 1, 2, \ldots$$

as the eigenvalue equation for the TM mode.

13

Using BPM techniques to calculate waveguide and resonator modes

While the main usage of beam propagation techniques is in finding spatial profiles of the optical field, say throughout volume of an integrated optics device, the method can also be used to determine modes of waveguides and optical resonators.

One favorable property of the approach described next is that only minimal modifications or additions above the standard implementation are needed. This is very useful in situations that call for initial conditions of a BPM problem that are expressed in term of guided modes. For example, we may require that the initial beam is the fundamental mode of a waveguide that serves as an input port into a device to be simulated. Such an initial condition can be conveniently calculated by the same program, but running along imaginary propagation distance. It is also convenient that this initial solution is obtained on the same grid that is used for subsequent BPM simulation.

The simple approach described in the following sections has the same mathematical justification for both waveguides and resonators. It draws on the simplest possible method to determine first leading eigenvalue(s) of a matrix of moderate sizes, sometimes called "power method." The name reflects that the calculation is essentially equivalent to evaluating the action of a high power A^n of the matrix A in question to an "arbitrary" initial vector.

It should be noted that there exist dedicated numerical algorithms to calculate eigenvalues of large matrices, and they would work even in cases when the methods outlined in what follows fail to converge.

13.1 Finding maximal eigenvalues of a matrix

In many physical problems, symmetric or self-adjoined matrices describe the system under investigation. Hamiltonias in quantum mechanics and operators that arise in the beam propagation framework are good examples. Such matrices have a system of eigenvectors that can serve as a basis to span the whole physical space of possible solutions, and their eigenvalues are real. Suppose all eigenvalues and eigenvectors of a given matrix A are known. For simplicity we will assume that at least a few largest eigenvalues are non-degenerate. The matrix then can be written as a weighted sum of projection operators that project on the subspaces corresponding to this matrix eigenvalues:

$$A = \sum_l \beta_l |v^{(l)}\rangle\langle v^{(l)}| \,, \tag{13.1}$$

where $|v^{(l)}\rangle$ is the eigenvector corresponding to the l-th eigenvalue. In the vector component form this reads

$$A_{ij} = \sum_l \beta_l v_i^{(l)} \bar{v}_j^{(l)} \ . \tag{13.2}$$

Because the eigenvectors corresponding to different eigenvalues are orthogonal,

$$A_{ik}^2 = \sum_j \sum_l \beta_l v_i^{(l)} \bar{v}_j^{(l)} \sum_u \beta_u v_j^{(u)} \bar{v}_k^{(u)} = \sum_{l,u} \beta_l \beta_u v_i^{(l)} \delta_{lu} \bar{v}_k^{(u)} = \sum_l \beta_l^2 v_i^{(l)} \bar{v}_k^{(l)} \ , \tag{13.3}$$

and higher powers of this matrix can be expressed simply as

$$A_{ij}^n = \sum_l \beta_l^n v_i^{(l)} \bar{v}_j^{(l)} \ . \tag{13.4}$$

Now imagine that the matrix acts n times on an "arbitrary" vector x_i. The resulting vector y_i,

$$y_i = A_{ij}^n = \sum_l \beta_l^n v_i^{(l)} \sum_j \bar{v}_j^{(l)} x_j \ , \tag{13.5}$$

is then a superposition of contributions (here they are distinguished by index l) that is dominated by that of the maximal eigenvalue. If one normalizes y, and repeats application of A^n on the normalized vector,

$$y \leftarrow y/|y| \quad , \quad y \leftarrow A^n y \ , \tag{13.6}$$

the relative contribution of the maximal eigenvalue becomes even more dominant. The power method consists in repeating the above multiply/normalize operations until the result converges to what is the eigenvector of A corresponding to its maximal eigenvalue, $y \to v^{max}$.

This procedure not only identifies the eigenvector v^{max}, but also its corresponding eigenvalue. This is because the normalization factor that must be applied to obtain new normalized vector y in the later stages of the iteration procedure converges to this eigenvalue.

Moreover, knowing the leading eigenvalue and its eigenvector, it is in principle possible to obtain the next-to-leading pair as well. This is achieved by the same "power procedure," but at each step one subtracts from the current vector its projection onto the (normalized) maximal eigenvector v^{max}:

$$y \leftarrow y - v^{max}(v^{max}.y) \quad , \quad y \leftarrow y/|y| \quad , \quad y \leftarrow Ay \ , \tag{13.7}$$

Because the contribution of the vector corresponding to the maximal eigenvalue is eliminated at each step, the second leading eigenvalue dominates, and this procedure thus leads to its eigenvector. Having determined the first two eigenvectors, one can determine the third in a similar way, by ensuring at each step that the running vector is perpendicular to both known eigenvectors. In principle this method can be extended to find all eigenvalues. In practice, first few can be obtained for moderately sized matrices.

An important question is how fast the power method converges. From the above formulas it should be evident that it is the ratio of the maximal to next-to-maximal eigenvalue that controls the convergence rate in the "final" approach to the solution. In the BPM context, the method is usually applied to determine the fundamental mode of a waveguide or of a resonator cavity, so it is important how close to the fundamental is the next mode. In case of a waveguide, it is often the eigenvalue difference $\beta^{TE} - \beta^{TM}$ between TE and TM polarizations that gives an inverse length-scale over which an initial guess-vector must be propagated. The convergence can be substantially accelerated by a good choice of the initial condition — by starting from a field polarized along the dominant direction, one can promote convergence to the desired polarization mode.

13.2 Propagation in imaginary distance

13.2.1 Propagation constants as matrix eigenvalues

On the abstract level, beam-propagation method evaluates an action of an operator exponential applied to an initial beam profile. In the numerical representation, the operator in question must be discretized and represented with a finite number of degrees of freedom, so it is in the end no more than a matrix. Thus we can write, without much loss of generality, that what a BPM does boils down to

$$E(z) = \exp[izA]\ E(z=0)\ , \tag{13.8}$$

where A is a matrix that represent the right-hand-size of the BPM propagation equation, and E stands for the vector array that represents the solution (which may include more than one vector components of the electric field). In this notation, the imaginary unit is pulled out of the propagator operators in order to separate the real matrix, and make it evident that when the propagation distance z is made imaginary, what we are left with is a "real" exponential.

Suppose that the the initial condition was selected such that it was a superposition of eigenmodes of the waveguide under consideration:

$$E(z=0) = \sum_l a_l v^{(l)}\ . \tag{13.9}$$

Vectors $v^{(l)}$ are also eigenvectors of A, and the propagation therefore results in

$$E(z) = \exp[izA]E(z=0) = \sum_l a_l \exp[izA]v^{(l)} = \sum_l a_l \exp[iz\beta^{(l)}]v^{(l)}\ . \tag{13.10}$$

Now we change the propagation distance and "rotate" it into a purely imaginary direction

$$\tilde{z} \leftarrow -iz\ . \tag{13.11}$$

This will give us real-valued exponentials

$$E(z) = \exp[\tilde{z}A]E(z=0) = \sum_l a_l \exp[\tilde{z}A]v^{(l)} = \sum_l a_l \exp[\tilde{z}\beta^{(l)}]v^{(l)}\ . \tag{13.12}$$

This is the same situation as we discussed in the previous section. With the increasing imaginary distance $\tilde{z}$, the exponential factor separate more and more while the one corresponding to the maximal propagation constant dominates. This leads us to the "power method" based procedure

$$E \leftarrow E/|E| \quad , \quad E \leftarrow \exp[zA]E\ , \tag{13.13}$$

which will converge to the fundamental mode, since that is the one with the maximal propagation constant.

It is worthwhile to note that by restricting the propagation to purely imaginary direction, beam propagation equations often simplify in the sense that real and imaginary parts of the beam amplitudes do not mix anymore (this would not be true if the material exhibited gain or loss, of course). Thus if one starts with a real valued transverse polarization components of the electric field, the imaginary part does not require to be propagated at all. It may be tempting to take advantage of this, but it is hardly worth of the effort needed to modify the code, the danger of introducing errors in the program, and the loss of generality in the implementation. Our aim in this section is to design an "inexpensive" method that only requires next-to-trivial modification of a given BPM program.

13.2.2 Boundary conditions

A couple of notes on boundary conditions is on order. Because the fundamental mode is going to dominate the numerical evolution along the imaginary distance axis, boundary conditions do not play a role as essential as in the free beam propagation simulation. Even if PEC boundary condition apply at the edge of the computational domain that is sufficiently distant from the waveguide core, the method will converge to an approximate fundamental mode. However, convergence study should be performed on the result to estimate the accuracy of the modal shape and its propagation constant. This can be done simply by comparing results from a few runs with different domain sizes, and thus with different distances between the core and artificial PEC boundary. What one should look for, is enough room in the simulated cladding to accommodate the exponentially decaying tail(s) of the mode.

The imaginary-distance beam-propagation method with the perfectly matched layer boundary conditions often yields unstable results for the guided-mode analysis of an optical waveguide. It is therefore useful to include in the simulator a switch that can toggle between PEC and PML. Since PEC condition is normally applied as part of the PML implementation, this does not represent a major change in the code.

The issue of perfectly matched layer and stability in modal calculations can be traced back to the fact that the fundamental mode wave is essentially evanescent in the vicinity of the domain boundary. Thus, there is no phase variation, only decay in the amplitude. This fact can be conveniently used to design yet another version of "stretched coordinate" system (J. Shibayama), this time in real (as opposed to complex-values) space. It boils down to mapping the real axis on a compressed version of itself and thus easier accommodation the extent of the evanescent tail. However, this approach requires a significant change in how we have coded the boundary layer, and we will not discuss this in detail.

Periodic boundary conditions may sometimes be also useful. While their usage goes a little bit against the spirit of this section, namely using a based BPM program with minimal modifications to calculate modes, periodic BCs are not too difficult to implement, so we want to comment on them at least briefly. One example of a situation which calls for application of PBCs is evaluation of modes in index-guiding photonic crystal fibers. Their core can be considered as placed in an infinite cladding that itself has a periodic structure. Rather than breaking this periodicity (or simulating a domain size that can encompass the whole fiber cladding, which usually means several "rows" of holes), it is more advantageous to apply periodic boundary conditions on the computational domain. It is however very important to achieve that the domain size and the period of the cladding "lattice" do agree accurately.

13.2.3 Variations of imaginary distance propagation

Sometimes it is required that a mode and its propagation constant are identified that do not represent the fundamental mode. The extension of the power iteration described above which requires elimination of the leading eigenvectors may not be practical if the sought after mode is "buried" deep in the spectrum of all possible propagation modes.

XXX K. Saitoh XXX

13.3 Fox and Li method for optical resonators

Modes of optical resonators can be often calculated using the same mathematical approach as that of the previous section. If one simply follows what "nature does" in the beam propagation within an optical resonator, the solution will naturally tend to the most sustainable configuration of the

electromagnetic field. This is normally achieved by transforming the beam profile as it exists at mirrors and at other elements that may be present in the cavity. A series of such transformation then represents one round-trip through the cavity. In general light can be amplified or partially lost during such propagation. Continuing in these round-trips will impose different losses or gain on different configurations of the field and gradually "emphasize" the one with the smallest loss or largest gain.

Mathematically, the transformation of the beam profile during a single cavity round-trip can be viewed as an action of a "cavity propagator" operator. For a given discrete representation of the field, this operator is a finite-dimensional matrix, and multiple round-trips correspond to its powers. Hence there is a similarity with the power iteration method discussed in the previous sections, and

The original method of Fox and Li employs the Fresnel-Kirchhoff diffraction integral to "propagate" the profile of the optical beam from one optical element to the next. We are interested in a modification of this approach such that a BPM technique is employed instead of the diffraction integral.

It is fair to say that such an "upgrade" does not make the method more robust. On the contrary, there are practical computational issues that must be reckon with. However, in some situations one may prefer to use a BPM-based cavity propagator, for example when the cavity calculation is a part of a more complex propagation problem. This section therefore briefly discussed a couple of additions that, when added on top of the free-space propagation techniques developed in the introductory sections, will make BPM technique suitable for calculation of modes in simple optical resonators.

We will only consider spectral methods as discussed in Section XXX. The first modification that may be needed has to do with the maximal distance to which a spectral method can propagate a beam without some of the radiation reaching computational domain boundaries. This issue can be dealt with by adding artificial "apertures" to the propagation, arguing that if light reaches an aperture it would surely not return to resonator later. Naturally, in this we face the problem that spectral methods are not too flexible in dealing with the boundary conditions. One has to resort to a version of the poor mans boundary, and employ a soft-edge aperture that absorbs light that gets too close to the domain boundary. The softness of the aperture is meant to eliminate as much as possible unwanted (and unphysical) reflections. Sufficiently large part of the domain must be sacrificed for the boundary layer in order to avoid leakage of light through the periodic boundary conditions inherent to spectral BPM.

Note that artificial apertures may need to be employed also if real apertures are present in the propagation problem. Any hard edge aperture will introduce waves with spatial wavenumbers spanning the whole numerically available bandwidth. Some of them may contribute the the resonator mode, but some may leave the cavity which means that they have to be absorbed in our artificial apertures.

The next ingredient we discuss here is transition of the beam through ideal thin lenses and/or its reflection from mirrors. These are modeled as corresponding "phase masks" that are applied to the propagating beam. Such a phase mask corresponds to the relative phase shift experienced by different parts of the beam wavefront. It should be emphasized that this view is only valid for continuous wave situation. Should the BPM be used as a part of time-domain problem solution one has to be more careful. If a very short duration pulse has to be modeled, and reflected off a mirror with a not too long focal distance, the phase-front curvature alone is not sufficient. In fact, a simple phase-screen model of a mirror results in an artificial pulse prolongation in the focus. For a few-femtosecond pulse, focal distance on the order of one meter and a beam diameter of few millimeters, the additional duration can easily reach ten femtoseconds. If one needs to deal with a situation like that, there is simple way around; Instead of application of a phase-screen

to shape the phase front, one applies a position dependent *temporal* shift. For a parabolic mirror this temporal shift is parabolic in radius, and for a long-duration pulse it exactly corresponds to the simple phase-screen model of the mirror.