A block pseudospectral method for Maxwell's equations: I. One-dimensional case

Tobin A. Driscoll* Bengt Fornberg†

Abstract

A block pseudospectral (BPS) method is a proposed as a new way to couple pseudospectral discretizations across interfaces in computations for a linear hyperbolic system. The coupling is achieved via discretized derivative-matching conditions obtained from the system. Compared to the standard technique of imposing compatibility conditions based on characteristics of the system, the BPS method offers better stability and accuracy, especially in the case where equation coefficients are discontinuous. Computational examples for Maxwell's equations in nonhomogeneous media demonstrate that BPS retains high accuracy over times orders of magnitude larger than not only low-order methods (such as Yee's), but also high-order methods, such as characteristic-based spectral elements.

Subject classifications: 65M70, 65M55, 78-08, 78A40

1 Introduction

For a linear hyperbolic system, such as Maxwell's equations, waves change little as they propagate over homogeneous regions that in applications are often large. Conceptually, a high-order finite-difference method on a uniform grid would appear to be useful. Two circumstances complicate matters.

- To accomodate irregular geometries and maximize parallelizability, the domain should be split into fairly autonomous blocks.
- If approximations of spatial derivatives are based strictly on data interior to a block, severe loss of accuracy may occur near the block edges.

Note that the latter concern is valid even in a global (single-block) scheme. The traditional answer to this concern is to cluster the grid quadratically as the block edge is approached, leading to Chebyshev (or more general Jacobi-type) grids. The usual pseudospectral-element method employs block Chebyshev grids and interface coupling via characteristic inflow/outflow conditions [3, 4, 11, 12, 9, 10, 14, 15, 18]. However, grid clustering introduces two drawbacks:

*Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, tad@colorado.edu. This work was funded by an NSF Postdoctoral Research Fellowship.

†Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, fornberg@colorado.edu. Partial support provided by NSF and AFOSR.

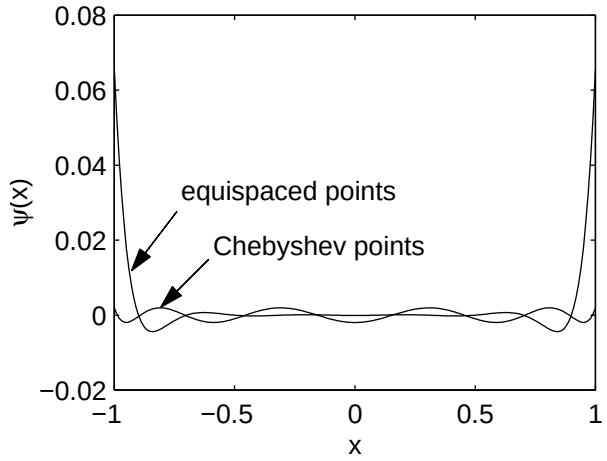

Figure 1: Interpolation remainder polynomial $\psi(x)$ for 10 equispaced and 10 Chebyshev nodes.

- A much stricter CFL condition, and
- Depletion of the grids away from interfaces

Our block psuedospectral (BPS) method is designed to allow the blocking of the domain while reducing these drawbacks. (A very different approach to reducing grid nonuniformity is the mapping technique due to Kosloff and Tal-Ezer [13].) The key idea in the BPS approach is:

> *Accuracy can be gained by clustering a grid less and using the approximations only well within the interior.*

Nonperiodic pseudospectral (PS) methods are based on polynomial interpolation. Given interpolation nodes x_j, $j = 1, \ldots, N$, the interpolation error for a smooth function f is

$$f(x) - p_{N-1}(x) = \frac{f^{(N)}(\xi)}{N!} \prod_{j=1}^{N} (x - x_j) = \frac{f^{(N)}(\xi)}{N!} \psi(x)$$

for some $\xi \in [x_1, x_N]$. It is a classical result of approximation theory that $\psi(x)$ is minimized in the max-norm over $[-1, 1]$ when the nodes are at Chebyshev locations. (Other Jacobi-type distributions are essentially equivalent in this respect.) In Figure 1 we compare $\psi(x)$ for the Chebyshev and equispaced nodes for $N = 10$. The equispaced nodes lead to a $\psi(x)$ that is large near the ends of the interval, as in the well-known Runge phenomenon.

However, we note that the equispaced approximation is actually much better over *most* of the interval. We will exploit this region of high accuracy by adjoining *fictitious points* to the computational domain, making it appear to be the interior part of an extended domain. We will never compute derivative values at the fictitious points. (In practice, our grids will be a compromise between equispaced and Chebyshev grids.)

The key issue in BPS becomes how to assign function values at the fictitious points for a given block. We consider two alternatives, depending on the nature of the interface between blocks.

1. *Blocks meet where the medium is continuous.* The fictitious values can be obtained by interpolation from the interior of the neighboring block.

2. *Blocks meet at a material interface (medium discontinuity).* The derivatives of the solution at the interface will generally be discontinuous, so straightforward interpolation is inappropriate. However, we can derive analytic conditions that relate the solution derivatives as taken from the two sides of the interface. By extending the one-sided finite-difference stencils to include the fictitious points, we can impose discrete forms of the derivative conditions to define the fictitious values. Similar techniques apply at domain boundaries.

To be specific about the second case, suppose in a 1-D problem that each block (subinterval) has two fictitious points extending over a common material interface. To determine the four unknown function values, we compute finite-difference weights (including the fictitious points) for one-sided derivatives up to order three at the interface. Imposing the analytic relationships for those derivatives leads to a total of 4 discrete linear conditions, which can be solved for the unknown function values. All derivatives within a block are then computed using both the block-interior and fictitious values.

The derivative-matching strategy also works well in the continuous-medium case, and we have used it for all the 1-D examples in this paper. For 2-D and 3-D simulations we expect that interpolation will be necessary, as outlined in the final section.

2 Finite-difference viewpoint of PS methods

We consider nonperiodic discretizations of the computational domain $[-1, 1]$ using grid points $-1 \le x_1 < \cdots < x_N \le 1$. Given the values of a function $f(x)$ at these points, we want to approximate $f'(x_j)$, $j = 1, \ldots, N$. One way to execute the FD approach is to choose a stencil at each x_j, interpolate the given values of f at its points by a minimal degree polynomial, and then differentiate the interpolant at x_j. More points in the stencil translate into a higher order of accuracy in the method. When the grid points coincide with those of a Jacobi polynomial pseudospectral method, and the stencil to be used at each grid point extends over all the grid points, then the FD approach is equivalent to the classical PS method [8]. A similar result holds for periodic discretizations. However, the FD method is defined for any grid, whereas the traditional PS methods are meaningful only on the Jacobi-type grids (although one can make special modifications [5, 13]).

In practice, the FD algorithm is not implemented by repeated polynomial interpolation. Instead, one computes each weight w_{jk}^m that multiplies the value $f(x_k)$ in the approximation to $f^{(m)}(x_j)$. Often the weights are assembled into a differentiation matrix that operates on vectors of function values to produce vectors of (approximate) derivative values. There are fast, stable algorithms both for computing the weights for a given stencil and point, and for computing a complete differentiation matrix [8, 17]. A pseudospectral FD method has a dense differentiation matrix, and computing a

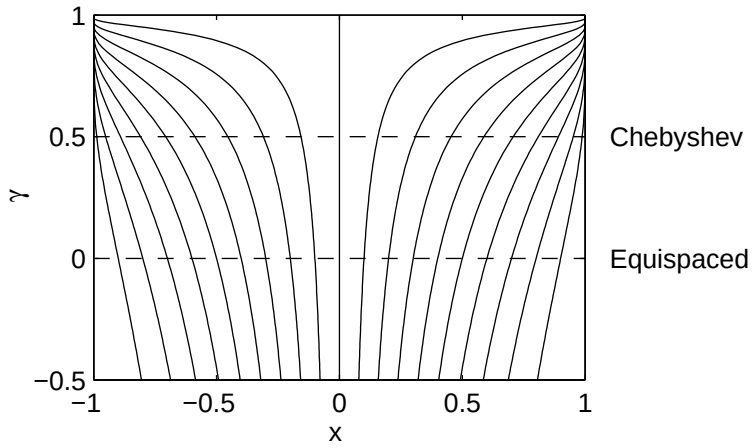

Figure 2: Positions of points in the closed grid defined by (1) with $N = 20$. The points are equispaced when $\gamma = 0$ and become increasingly clustered near the ends as γ increases towards 1.

derivative with it takes $O(N^2)$ operations.[1]

For the discussion that follows, we want to control the amount of grid clustering near the ends of the domain. Following [8], we introduce a parameter $\gamma < 1$ and define the grid points x_j for $j = 0, \ldots, N$ implicitly by

$$\frac{j}{N} = \int_{-1}^{x_j} c_\gamma (1 - x^2)^{-\gamma}\, dx, \qquad (1)$$
$$c_\gamma = \frac{\Gamma(\frac{3}{2} - \gamma)}{\sqrt{\pi}\,\Gamma(\frac{1}{2} - \gamma)}.$$

The grid defined by (1) is closed; i.e., ± 1 are included. For the BPS method it will be useful to have open grids, in which the endpoints are excluded. These are defined for $j = 1, \ldots, N$ by

$$\frac{j - \frac{1}{2}}{N} = \int_{-1}^{x_j} c_\gamma (1 - x^2)^{-\gamma}\, dx. \qquad (2)$$

When $\gamma = 0$ the points become equispaced throughout the interval. When $\gamma = 1/2$, the open and closed grids coincide with the Gauss–Chebyshev and Gauss–Lobatto–Chebyshev points resp., which are quadratically clustered near ± 1. Figure 2 shows the positions of the closed-grid points for $N = 20$ and continuously varying γ. Traditionally one chooses $\gamma = 1/2$, which leads to optimal error over $[-1, 1]$ as $N \to \infty$. We will effectively choose $0 < \gamma < 1/2$ and restrict the use of the interpolant to a proper subinterval of $[-1, 1]$.

[1] An advantage of Chebyshev grids is that the FFT can be used to evaluate derivatives, requiring only $O(N \log N)$ work. However, in block methods, the values of N are typically far too small for FFTs to have any benefit.

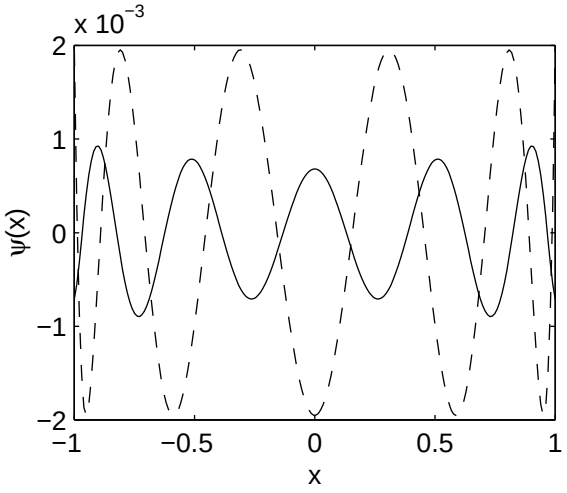

Figure 3: Interpolation remainder polynomial $\psi(x)$ for $N = 10$. The dashed curve is for the Chebyshev grid (cf.Figure 1); the solid curve was produced with $\gamma = 0.35$ and one fictitious node per side.

3 Fictitious point methods

The idea of fictitious points is to augment a grid with additional nodes at which the function values are not directly known nor updated. The effect is to only use the interpolant well within the interior of an extended interval. Figure 3 illustrates this by comparing the interpolation remainder $\psi(x)$ for a fictitious-point grid to those shown in Figure 1.

To construct this particular grid, we start with an 10-node open grid defined by (2) with $\gamma = 0.35$. Then we add one fictitious point per side of the interval by reflecting the outer points about ± 1. Compared to a Chebyshev grid ($\gamma = 0.5$), the accuracy of the fictitious point method is twice as good, although the minimum spacing between true (non-fictitious) nodes increases by a factor of about 1.5.

The fictitious point method must include a means of determining function values at the fictitious points. At a domain boundary, the values are determined by supplementary boundary conditions, typically involving derivatives of the solution. A study of this technique has been done [7, 8] for the model eigenvalue problem $u'' = \lambda u$ with $u(\pm 1) = 0$, where the additional conditions take the form $u''(\pm 1) = 0$, $u''''(\pm 1) = 0$, ... The fictitious values are those which make an appropriate number of the finite-difference versions of the conditions hold. It is found that as the number of fictitious points increases, the acceptable γ decreases and the outlying eigenvalues of the differentiation matrix become less severe. Another example of using fictitious points with high-order methods is described in [].

At an interface between two subdomains, two techniques are considered for determining the fictitious values. The first technique is to obtain the required values by interpolation from the

neighboring subdomain. This reduces to a standard overlapping method, except for the choice of grids. If the subdomain interface coincides with a material interface (discontinuous coefficients), however, the values obtained from direct interpolation would be incorrect due to a loss of smoothness in the solution at the interface. In this case we return to the idea of derivative conditions derived from the equations of the system. These conditions will determine fictitious values, so the interpolants used will obey high-order relations at the interface. Derivative matching works even when the coefficients are smooth, and we have found it to have slightly better accuracy and stability properties than the interpolation technique. Hence, derivative matching is used for all of the BPS simulations of this paper.

4 Derivative matching at an element interface

We consider the standard interval $x \in [-1,1]$, divided into two subintervals, $I_1 = [-1,\xi]$ and $I_2 = [\xi,1]$. We derive the BPS method for linear hyperbolic systems of the form

$$u_t(t,x) + A(x)u_x(t,x) = B(x)u, \qquad -1 \le x \le 1, \qquad u \in \mathbf{R}^d, \tag{3}$$

where

$$A(x) = \begin{cases} A_1 & \text{if } x \in I_1, \\ A_2 & \text{if } x \in I_2, \end{cases} \tag{4}$$

and the matrix $B(x)$ has an analogous form. (We defer discussion of boundary conditions until Section 5.) The prototypical situation we have in mind is electromagnetic radiation passing through two simple dielectric materials. A more general spatial or time dependence of the coefficient matrices requires a straightforward generalization of the derivation that follows. The special case of an artificial interface occurs when $A_1 = A_2$ and $B_1 = B_2$.

Given the values of u at the spatial grid points of each subinterval at some time, our goal is to find approximate values of the time derivative u_t at those grid points. We want to use primarily information local to a subinterval, with some coupling to its neighbors. While u is required to be continuous, spatial derivatives of u will not necessarily be continuous at $x = \xi$. We denote the one-sided derivatives at the interface by

$$u^{(q)}(t,\xi^-) = \left.\frac{\partial^q u}{\partial x^q}\right|_{x\to\xi^-}, \qquad u^{(q)}(t,\xi^+) = \left.\frac{\partial^q u}{\partial x^q}\right|_{x\to\xi^+}.$$

By assumption, we have that

$$u^{(0)}(t,\xi^-) = u^{(0)}(t,\xi^+).$$

We also require that u_t be continuous. Differentiating with respect to time, we find that

$$-A_1u^{(1)}(t,\xi^-) + B_1u^{(0)}(t,\xi^-) = -A_2u^{(1)}(t,\xi^+) + B_2u^{(0)}(t,\xi^+).$$

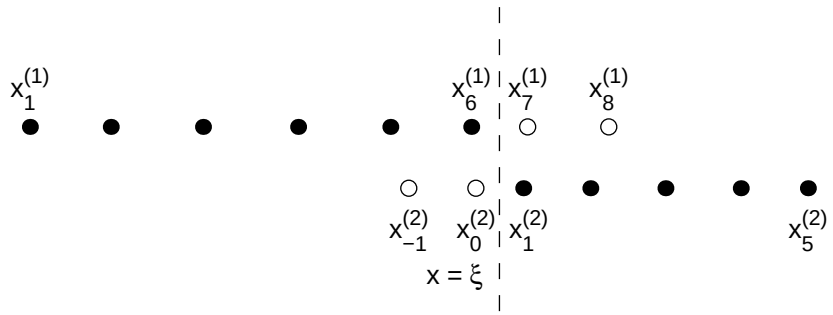

Figure 4: Grid points for $N_1 = 6$, $N_2 = 5$, and $m = 2$. The fictitious points are shown as open circles.

Staying strictly on one side of the interface at $x = \xi$, time differentiation is equivalent to the operator $(-A\partial_x + B)$, and the spatial differentiation operator commutes with the matrices. Given a positive integer m, we make the definitions

$$\mathcal{A}_1 = \begin{bmatrix} I & & & & \\ B_1 & -A_1 & & & \\ B_1^2 & -(A_1B_1 + B_1A_1) & A_1^2 & & \\ & & & \ddots & \\ B_1^{2m-1} & \cdots & \cdots & (-A_1)^{2m-1} \end{bmatrix}, \qquad \mathcal{U}_1 = \begin{bmatrix} u^{(0)}(t,\xi^-) \\ u^{(1)}(t,\xi^-) \\ \vdots \\ u^{(2m-1)}(t,\xi^-) \end{bmatrix}, \tag{5a}$$

where the block entries of row q are the matrix coefficients of powers of ∂_x in $(-A\partial_x + B)^q$. With similar definitions for $\mathcal{A}_2$ and $\mathcal{U}_2$, we summarize the continuity of time derivatives up to order $2m-1$ of u at ξ by the system

$$\mathcal{A}_1\mathcal{U}_1 = \mathcal{A}_2\mathcal{U}_2. \tag{5b}$$

We impose these derivative conditions on the interpolants for I_1 and I_2. For $k = 1, 2$, let $x_1^{(k)}, \ldots, x_{N_k}^{(k)}$ be the N_k grid points in the interior of I_k. Our interpolant for I_k will extend over the N_k grid points and m fictitious points determined by reflection of grid points about the line $x = \xi$, as illustrated in Figure 4. To find the function values at the fictitious points, we enforce discrete versions of the conditions (5). For I_1, we denote the FD weights $w_{qj}^{(1)}$ in the approximations to derivatives at ξ as follows:

$$\frac{d^q f}{dx^q}(\xi) \approx \sum_{j=1}^{N_1+m} w_{qj}^{(1)} f(x_j^{(1)}),$$

and define

$$W_1 = \begin{bmatrix} w_{0,1}^{(1)} & \cdots & w_{0,N_1+m}^{(1)} \\ \vdots & & \vdots \\ w_{2m-1,1}^{(1)} & \cdots & w_{2m-1,N_1+m}^{(1)} \end{bmatrix}, \qquad v_1 = \begin{bmatrix} u(t, x_1^{(1)}) \\ \vdots \\ u(t, x_{N_1+m}^{(1)}) \end{bmatrix}.$$

Recall that u is a d-vector. To map v_1 to a discrete approximation of $\mathcal{U}_1$, we must find d zeroth derivatives, d first derivatives, etc. Thus, we want to multiply by the matrix

$$\begin{bmatrix} w^{(1)}_{0,1} I_d & w^{(1)}_{0,2} I_d & \cdots & w^{(1)}_{0,N_1+m} I_d \\ w^{(1)}_{1,1} I_d & w^{(1)}_{1,2} I_d & \cdots & w^{(1)}_{1,N_1+m} I_d \\ \vdots & \vdots & & \vdots \\ w^{(1)}_{2m-1,1} I_d & w^{(1)}_{2m-1,2} I_d & \cdots & w^{(1)}_{2m-1,N_1+m} I_d \end{bmatrix},$$

where I_d is the d-dimensional identity matrix. Making similar definitions for I_2 we write compactly

$$\mathcal{A}_1 (W_1 \otimes I_d) v_1 = \mathcal{A}_2 (W_2 \otimes I_d) v_2 \tag{6}$$

as the discrete form of (5). Here $\otimes$ is Kronecker matrix multiplication.

Some remarks:

- The unknowns in equation (6) are $u(t, x^{(1)}_{N_1+1})$, ... , $u(t, x^{(1)}_{N_1+m})$, $u(t, x^{(2)}_{-m+1})$, ... , $u(t, x^{(2)}_0)$. These are determined by the $2md$ linear equations.
- The values of any component of u on the fictitious points may be coupled to all the values of all the components of u.
- The procedure can be seen as an extrapolation based on derivative values at $x = \xi$, an unstable procedure as the numbers of variables go to infinity. However, in practice one would use few variables and treat the subintervals near the interface as a "boundary layer" having small spatial extent. This also prevents the Kronecker-product matrices from becoming large.
- Because the weights for higher derivatives increase exponentially, the rows of the weight matrices $W_{1,2}$ should be rescaled to prevent severe ill-conditioning. We choose to make the largest element in each row equal to one.

In a large application one might have many subintervals, each with an interior grid. For a fixed subinterval I_k, we add m fictitious points outside each end of I_k. The function values for each group of fictitious points are determined as above (or boundary conditions, as described in Section 5) independently of other interfaces. Then the values are interpolated over I_k and both sets of fictitious points simultaneously and differentiated. Hence, the derivatives over I_k depend on data only over I_k and its immediate neighbors.

4.1 Example: Maxwell's equations

For Maxwell's equations in nonconducting media, the derivative relations for the BPS method take a particular, simple form. The equations are, for transverse electric (TE) waves

$$u = \begin{bmatrix} E \\ H \end{bmatrix}, \qquad A_k = \begin{bmatrix} 0 & \frac{1}{\epsilon_k} \\ \frac{1}{\mu_k} & 0 \end{bmatrix}, \qquad B_k = 0, \qquad k = 1, 2. \tag{7}$$

As a result, the derivative relations for the electric component E and magnetic component H are decoupled. If

$$\rho = \frac{\mu_2 \epsilon_2}{\mu_1 \epsilon_1} = \left(\frac{c_1}{c_2}\right)^2,$$

then

$$\frac{E_2^{(q)}(t,\xi)}{E_1^{(q)}(t,\xi)} = \rho^{\lfloor q/2 \rfloor}, \qquad \frac{H_2^{(q)}(t,\xi)}{H_1^{(q)}(t,\xi)} = \rho^{\lfloor (q+1)/2 \rfloor}, \tag{8}$$

where c_1, c_2 are the propagation speeds and $\lfloor \cdot \rfloor$ represents rounding down to the nearest integer. The differentiation matrices for the field components can be computed separately.

If either material has a finite nonzero conductivity, the relationships are more algebraically complex. In this case the right-hand side matrix is of the form

$$B_k = \begin{bmatrix} -\sigma_k/\epsilon_k & 0 \\ 0 & -\sigma_k^*/\mu_k \end{bmatrix}, \qquad k = 1, 2, \tag{9}$$

where σ is electric conductivity and σ^* is an analogous magnetic loss (often zero for real materials). The derivative relations now involve E and H simultaneously.

5 Boundary conditions

We shall illustrate the treatment of boundary conditions by considering in detail the two most important types for numerical simulations of Maxwell's equations in CEM: the perfect electric (resp. magnetic) conductor, at which $E = 0$ (resp. $H = 0$); and an absorbing layer meant to simulate radiation that leaves the computational domain.

5.1 Perfect conductor

Without loss of generality, we consider $I_1 = [-1, \xi]$ and perfect electric conductor at $x = -1$. Our given condition is $E(-1) = 0$. Clearly, at $x = -1$,

$$\begin{aligned} 0 &= E_t \\ &= -\frac{1}{\epsilon} H_x - \frac{\sigma}{\epsilon} E \\ &= -\frac{1}{\epsilon} H_x, \end{aligned}$$

so $H_x(-1) = 0$. Similarly,

$$0 = E_{tt} = -\frac{1}{\epsilon} H_{xt} = c^2 E_{xx} + \frac{\sigma^*}{\epsilon \mu} H_x = c^2 E_{xx}.$$

Continuing, we find that E is an odd function and H is an even function about $x = -1$.

Now let us add the m fictitious points $x^{(1)}_{-m+1} < \cdots < x^{(1)}_0 < -1$. We extend the interpolants of E and H to these points. To determine the field values at the extra points, we impose the finite-difference forms of

$$\left.\frac{\partial^q E}{\partial x^q}\right|_{x=-1} = 0, \qquad q = 0, 2, \ldots, 2m-2,$$
$$\left.\frac{\partial^q H}{\partial x^q}\right|_{x=-1} = 0, \qquad q = 1, 3, \ldots, 2m-1.$$

The mechanics of this process are completely analogous to the algebra described in §4 for derivative matching. These $2m$ conditions determine the $2m$ unknown extra field values, which are in turn used in the standard finite-difference procedure to find E_x and H_x. This procedure is easily extended to the general hyperbolic system (3) with linear homogeneous boundary conditions.

5.2 Absorbing layer

A common technique for implementing an absorbing boundary is to impose a one-way wave equation (OWWE) at the boundary point. It seems clear that no extension of grid points past the boundary is appropriate for such a technique, since by premise there is no additional information available at the boundary. To maintain accuracy and stability, we would need to maintain a highly clustered grid near the boundary. One way to do this is to replace the weight function that defines grid points in (2) by $(1+x)^{-1/2}(1-x)^{-\gamma}$ for the boundary at $x = -1$, for example.

An attractive alternative to the OWWE, developed by Berenger [2], is the perfectly matched layer (PML). The PML offers substantially less reflection at oblique angles in multiple dimensions than do OWWE techniques. In one dimension the PML boils down to padding the computational domain with a fictional lossy material satisfying

$$\frac{\sigma}{\epsilon_0} = \frac{\sigma^*}{\mu_0},$$

where ϵ_0 and μ_0 are properties of free space. The wave impedance of the PML medium then matches that of free space perfectly, mathematically preventing reflections at normal incidence. Wave amplitudes in the PML material decay exponentially in space.

Implementing the PML method in BPS is therefore simply a matter of constructing a material interface at the physical boundary as described in §4 and adding, say, a perfect conductor to the far end of the layer.

6 Numerical accuracy and stability

There are several parameters available in the BPS method. Given a hyperbolic problem and boundary conditions, one may select subinterval sizes, the number of grid points in each subinterval (N_k), grid clustering (γ), and the number of fictitious points (m). We shall only attempt to indicate accuracy and stability in a few cases. Our measure will be the eigenvalues of the discretized

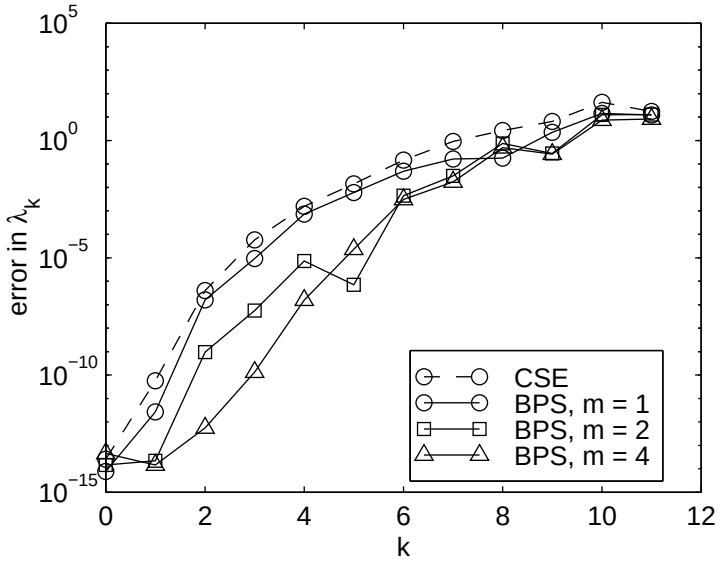

Figure 5: Accuracy of BPS for $u_t + u_x = 0$. The parameters are $N_1 = N_2 = 12$, $\gamma = 0.4$, and $m = 1, 2, 4$. Eigenvalue accuracy is shown for 12 nonnegative values of k. The BPS method is much more accurate than standard CSE.

approximation to the operator $-A\partial_x + B$, which would be applied to u to find u_t in a method-of-lines solution.

We first consider the simplest scalar problem with $A \equiv 1$, $B \equiv 0$, and periodic boundary conditions. The exact eigenvalues for this problem are $\lambda_k = i\pi/k$ for $k \in \mathbf{Z}$. Figure 5 shows the accuracy of the spectrum using the BPS method with $\xi = 0$, $N_1 = N_2 = 12$, $\gamma = 0.4$, and $m = 1, 2, 4$. For comparison, we also plot the accuracy using the differentiation matrix with $\gamma = 0.5$ (standard Chebyshev grid) and coupling via characteristic variables. Henceforth we refer to this method as characteristic spectral elements (CSE). The eigenvalues of the BPS method are more accurate than those for CSE for each choice of m, and the accuracy improves substantially with m. Surprisingly, the CSE method is slightly unstable, as the spectral abscissa (maximum real part of the spectrum) is about 7×10^{-3} here. This instability could manifest itself in a long-time integration or a domain with very many elements.

In Figure 6 we plot the spectra of the discretized operators in the complex plane for $m = 4$ and for the CSE method. The two methods feature roughly the same number of accurate eigenvalues, but the outliers are located quite differently. In the BPS method, the outliers also lie on the imaginary axis. This reflects the fact that the same differentiation matrix would be used for the problems $u_t - u_x = 0$ and $u_t + u_x = 0$; the derivative matching conditions would be unchanged. Hence the discrete spectrum is symmetric about the imaginary axis. The CSE outliers, on the other hand, lie far in the left half-plane, because this differentiation matrix (for $u_t + u_x = 0$) would be "downwind" for $u_t - u_x = 0$, and therefore unstable in time (i.e., the differentiation matrix would

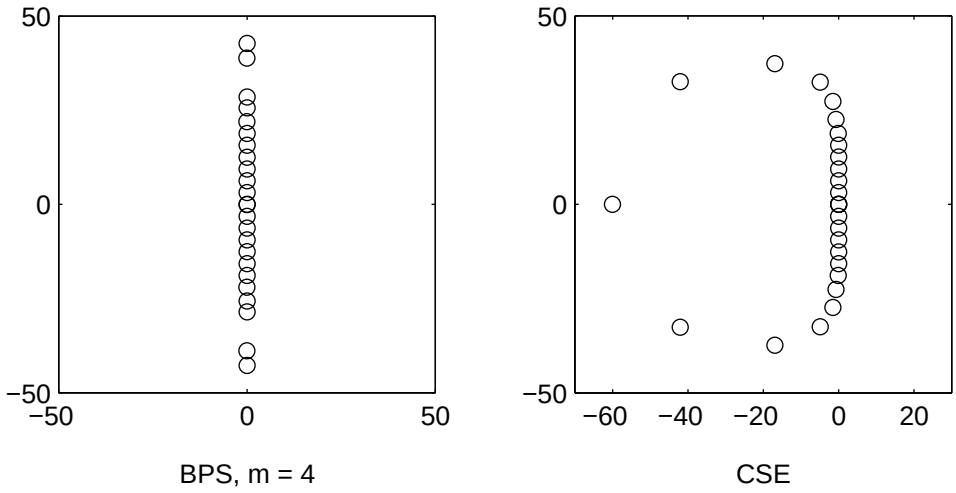

Figure 6: Spectra of BPS ($m = 4$) and CSE differentiation matrices for the problem of Figure 5. The BPS method is nondissipative because its treatment of interfaces is independent of the direction of flow, while the CSE method is strongly dependent on upwinding.

be changed to reflect the new flow direction).

We next examine Maxwell's equations in dielectric media having $\mu_1 = \mu_2 = 1$, $\epsilon_1 = 1$, $\epsilon_2 = 4$. We first set the interface $\xi = 0$ and make the problem periodic, thus including an interface at ± 1 as well. We will construct discrete approximations to the operator

$$-\begin{bmatrix} 0 & \frac{1}{\epsilon(x)}\partial_x \\ \partial_x & 0 \end{bmatrix}. \tag{10}$$

The exact eigenvalues include $2k\pi i$ for any integer k, but there are also other eigenvalues that can be easily determined numerically to high precision.

The characteristic-based CSE method can be extended to this problem. Within I_1, the characteristic variables $E + c_1 H$ and $E - c_1 H$ propagate with speed c_1 to the right and left, resp. The characteristic variables and speeds are discontinuous at the interface $x = \xi$, with c_1 replaced by c_2. At the interface, there are two values each of E_x and H_x computed. One takes the quantity $E_x + c_1 H_x$ from I_1 and $E_x - c_2 H_x$ from I_2 to define the values of E_x and H_x used in both subintervals. This amounts to modifying two rows in the analog of (10). A similar procedure is applied at $x = \pm 1$.

Figure 7 displays the accuracy of the spectrum obtained when $N_1 = 10$, $N_2 = 14$ and $\gamma = 0.4$, for $m = 1, 2, 4$. (More grid points are required where the wave speed is smaller.) The BPS method is now much more accurate for all choices of m than CSE (which always uses $\gamma = 0.5$). The accuracy again improves with m. As before, the CSE matrix is somewhat time-unstable, having spectral abscissa around 8×10^{-3}. The BPS eigenvalues are very nearly on the imaginary axis—up to 10^{-8}

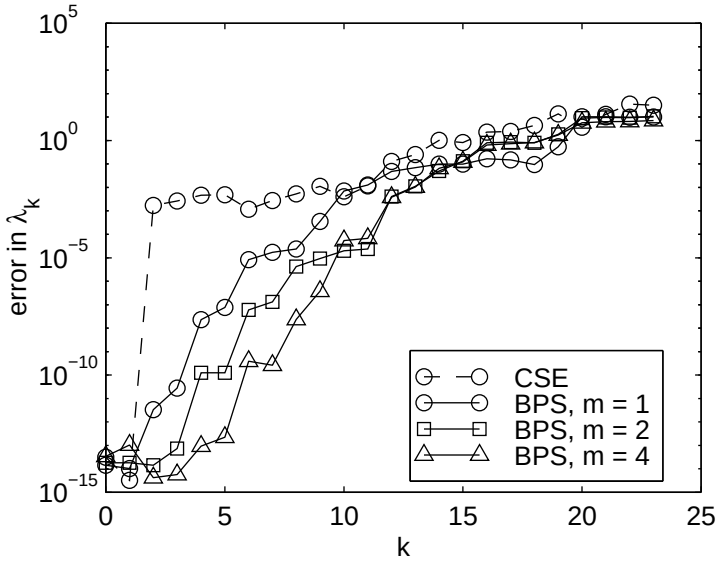

Figure 7: Accuracy of derivative-matched BPS method for Maxwell's equations through a medium discontinuity. The parameters are $N_1 = 10$, $N_2 = 14$, $\gamma = 0.4$, and $m = 1, 2, 4$. Eigenvalue accuracy is shown for eigenvalues in the complex upper half-plane. The BPS method is much more accurate than standard CSE (which is slightly time unstable).

or better.[2]

In Figure 8 we show the corresponding results for the unequal-interval case where $\xi = 1/3$, and with $N_1 = N_2 = 12$. The results are qualitatively similar. CSE is still sightly unstable, with roughly the same spectral abscissa as before, while BPS remains stable.

We do not wish to give the impression that the BPS method is stable for all situations. There is a danger that the outlying eigenvalues, which are nominally on the imaginary axis, will move off into the two half-planes. Factors causing this behavior certainly include the parameters N_k and m. Indeed, we have seen this behavior even with a smoothly varying $A(x)$. We are considering a number of ways to deal with this instability, but the simplest is to add a small artificial damping term such as ϵu_{xxxx} to the differential equation. High-order damping affects the high-frequency modes most strongly, and those are the troublesome ones. Accuracy is hardly affected, because these modes are dispersively incorrect in any case, and we can limit the damping in space to the vicinity of the interfaces. In contrast, the unstable CSE modes are at low wavenumbers; therefore damping would have to be much more destructive to be effective.

[2]The number would be more like 10^{-14} except that the eigenvalue at zero is not simple and exhibits great sensitivity to rounding errors.

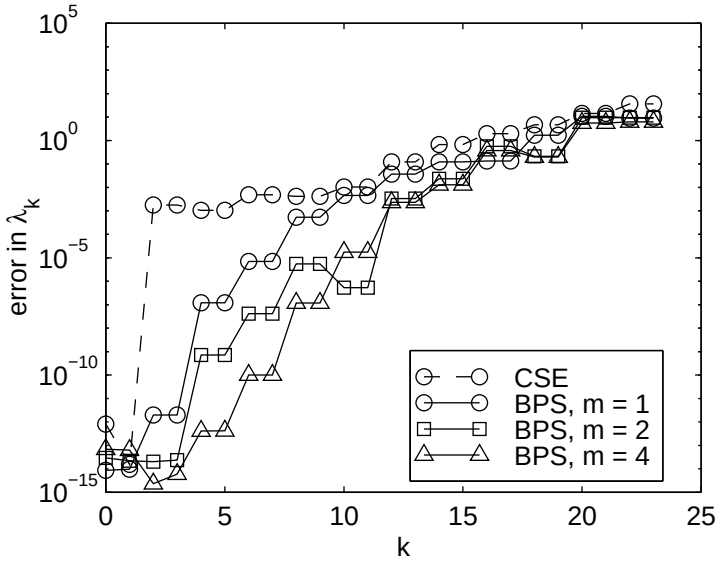

Figure 8: Same as Figure 7, but with interval lengths in the ratio 2:1 and $N_1 = N_2 = 12$. Stability is about the same as before.

7 Computational examples

It is well known that low-order finite-difference methods suffer from serious numerical dispersion [8]. Prohibitively many grid points are necessary for long-time simulations. Pseudospectral methods propagate more wavenumbers accurately. The improved eigenvalue accuracy of BPS over CSE implies that BPS should be more accurate for very long times.

This is illustrated by our first example, the propagation of electromagnetic waves through two idealized dielectric media:

$$\begin{aligned} E_t &= -\frac{1}{\epsilon(x)} H_x \\ H_t &= -\frac{1}{\mu} E_x, \end{aligned} \qquad -1 \le x \le 1, \tag{11}$$

where $\mu \equiv 1$ and the permittivity $\epsilon(x)$ is given by

$$\epsilon(x) = \begin{cases} 1, & \text{if } -1 \le x \le 0, \\ 4, & \text{if } 0 < x \le 1. \end{cases}$$

The wave speed is unity for $x \in [-1, 0]$ and $\frac{1}{2}$ for $x \in (0, 1]$. We apply periodic boundary conditions. The initial electric field is a cosine bell centered at $x = -0.5$ and with half-width 0.3; the initial magnetic field is chosen so that all energy initially moves rightward.

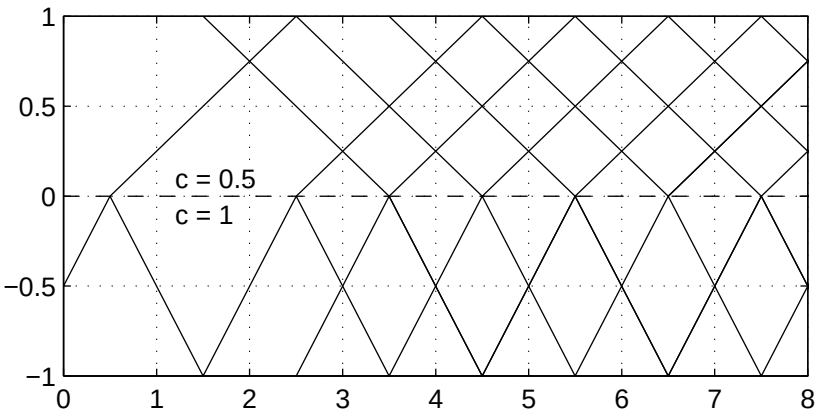

Figure 9: Schematic of the exact time history for the problem given by (11). The lines trace the paths of local pulse peaks. At integer times, there are only three distinct pulses.

Each time the pulse encounters a change of medium at $x = 0$ or $x = \pm 1$, it is split into reflected and transmitted parts. As time evolves, there arise many reflections and transmissions, but in this example such pulses coalesce with others centered at the same location. Thus only a few distinct pulses can be seen at any one time. Figure 9 traces the local pulse peaks and shows how they move. At each integer time, there are only three distinct pulses present, centered at $-1/2$, $1/4$, and $3/4$. The analytical solution can easily be recovered at these times.

In Figure 10 we compare fragments of the time histories of E for four numerical solutions. Each uses 34 spatial grid points and a time step of 0.01. The fragments begin at times 0, 100, and 1000, respectively. The leftmost history is for the Yee or FDTD algorithm [16, 19], using centered second-order time and space differences on staggered equispaced grids. The next history uses fourth-order differences in space, again with staggered second-order time stepping. The third plot shows the results of CSE, and the rightmost, for BPS, each using a fourth-order Runge–Kutta method in time. In all cases, the time step is sufficiently small so that time errors are insignificant.

The low-order methods are acceptable for very short times, but they bear little resemblance to the exact solution by $t = 100$. The fourth-order method is somewhat more accurate than second-order Yee. The CSE and BPS methods are nearly identical for small times. For times around 100, CSE still bears some relation to the solution, but the noise level is high. As mentioned above, CSE is weakly unstable for this problem, and instability dominates by time 1000 (the vertical scale of the figure is 1000 times greater in this fragment). In contrast, BPS still shows no significant degradation at time 1000. The infinity-norm errors are shown as functions of time in Figure 11. The instability of CSE is evident. Over the early times $0 \le t \le 10$, CSE and BPS are similar, while the Yee and FD4 methods quickly lose accuracy. Also shown are the results of a 7-point spatial finite difference scheme which has been "optimized" for propagation of intermediate wavenumbers [20]. Due to the low resolution of the initial pulse, FD4 outperforms this method. Figure 12 compares, at certain times, the analytical and numerical solutions for the four methods. The times are chosen

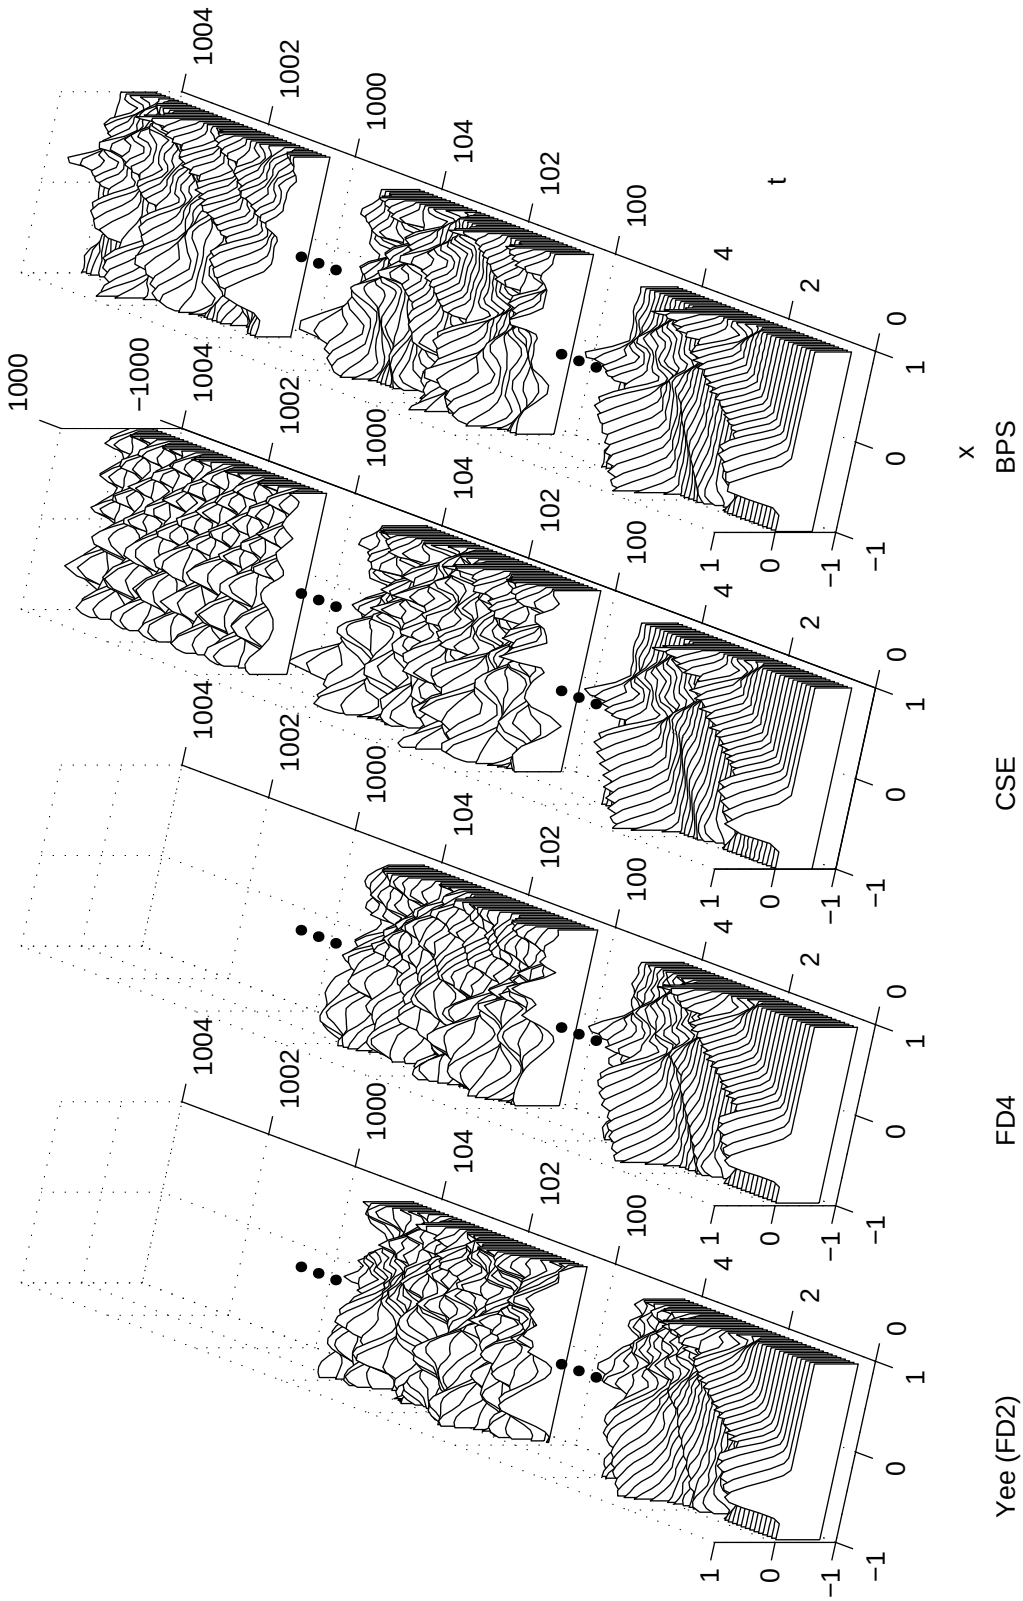

Figure 10: Time history fragments of E for four numerical methods for (11). The first two use 2nd- and 4th-order spatial differences on staggered equispaced grids, and second-order time stepping. The other two use CSE and BPS, respectively, with fourth-order Runge–Kutta in time. The last fragment of CSE is on the scale -1000 to 1000. There are 34 spatial grid points in all cases. Only for BPS are the pulses clearly visible throughout.

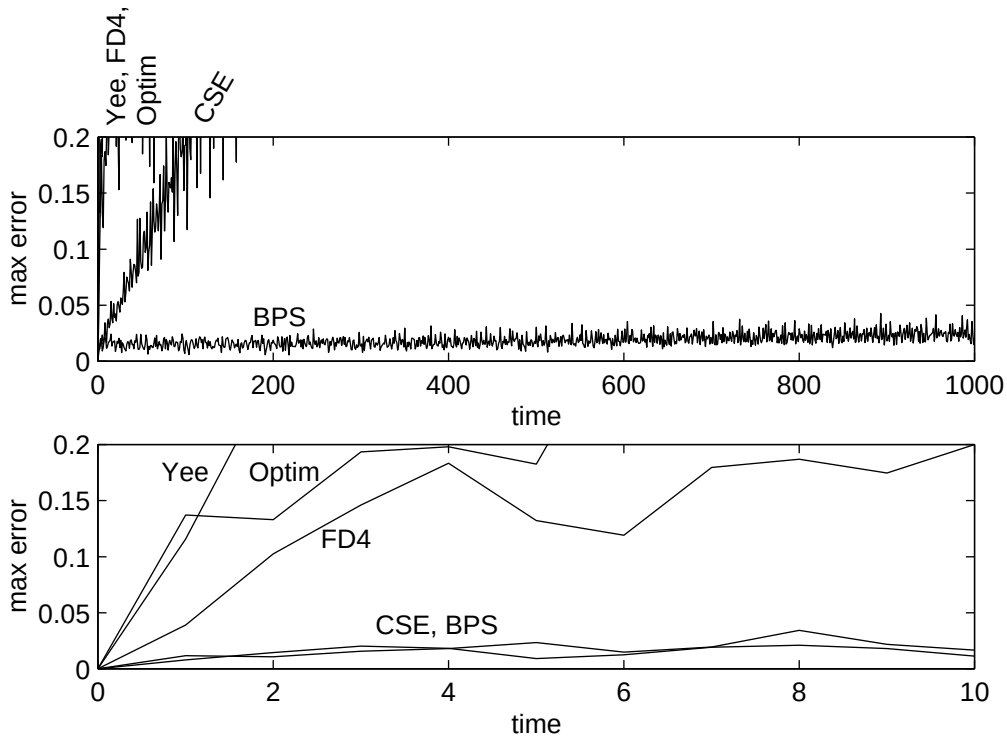

Figure 11: Errors of the numerical solutions in time. On the time scale from 0 to 1000, only the BPS method is acceptable—the instability of CSE is clear. On the shorter time scale from 0 to 10, the Yee, FD4, and "optimized" finite-difference methods quickly lose accuracy compared to the two spectral methods.

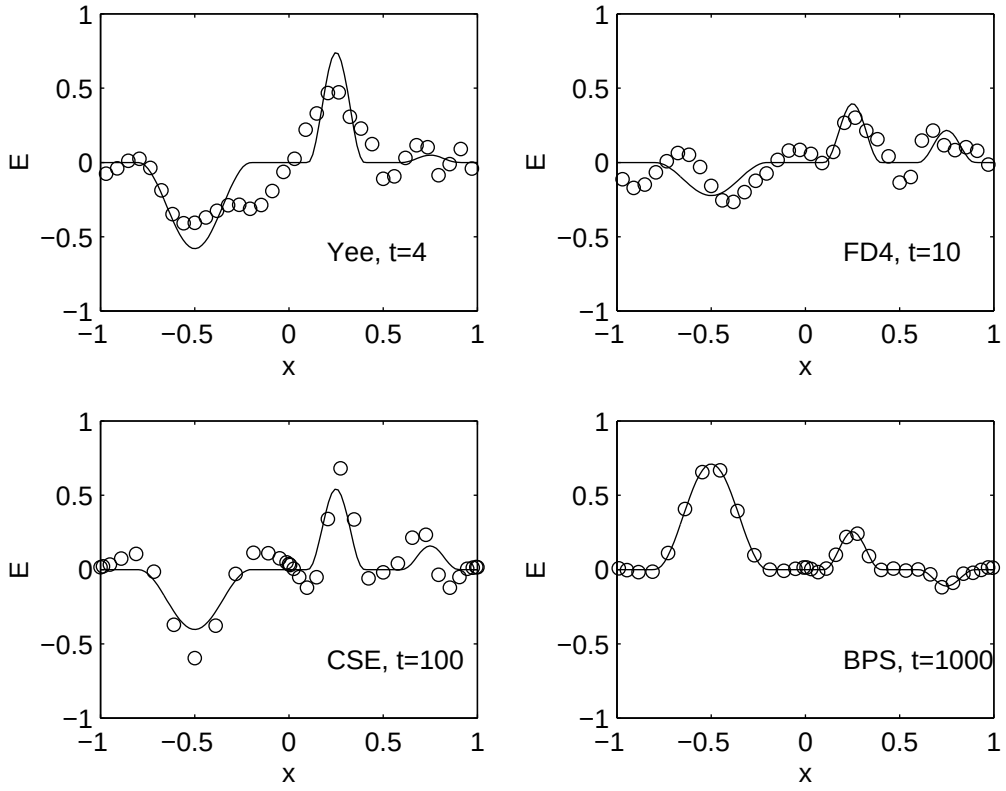

Figure 12: Snapshots of the numerical solutions of Figure 10. Yee, FD4, and CSE have lost substantial accuracy already at times 4, 10, and 100, resp., whereas the error in the BPS method is barely visible, even at time 1000.

such that the numerical solution begins to show significant inaccuracy, except for the BPS case, in which the solution is visually accurate for all $t \leq 1000$.

Our second example illustrates the simultaneous application of many of the devices outlined in previous sections. The underlying equations are again Maxwell's, with four different media, including both an absorbing layer and a perfectly conducting boundary. Figure 13 illustrates the problem and the subintervals used for the BPS solution.

Figure 14 shows the time history for a cosine bell of half-width 0.1 for the first five seconds. The errors again remain small over long times. Since the initial pulse is supported over only 7 of the 64 grid points, this is a fairly severe test case for dispersive errors.

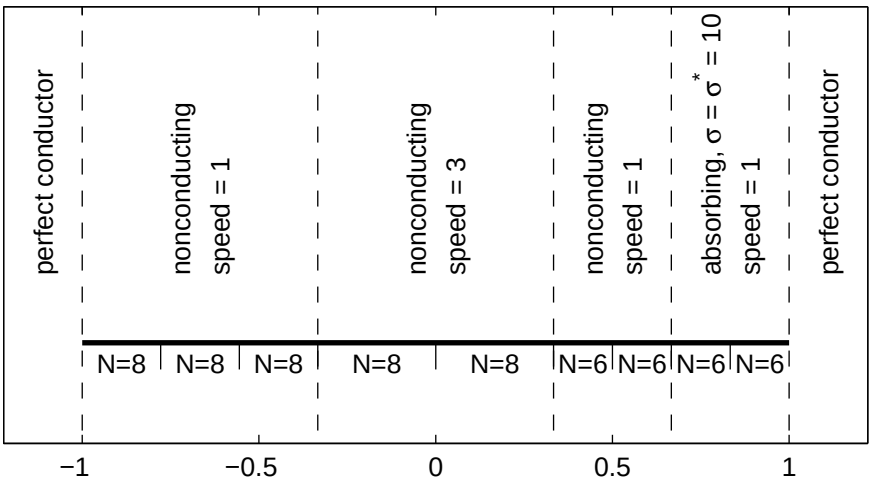

Figure 13: Material types and subintervals for the second computational example. The number of grid points in each subinterval is displayed below it. Each subinterval grid uses spacing parameter $\gamma = 0.35$ and $m = 2$. At every subdomain interface, four matching conditions are imposed.

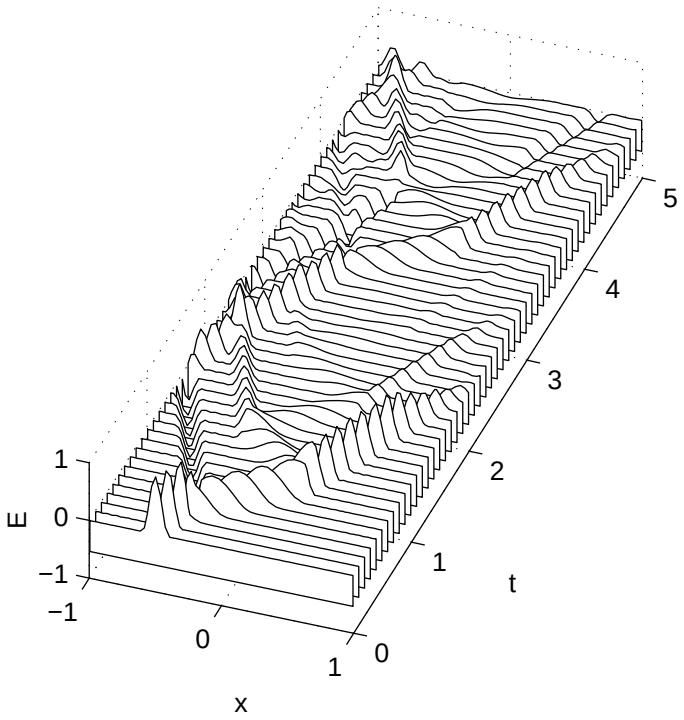

Figure 14: Time history for the BPS numerical solution of the problem illustrated in Figure 13.

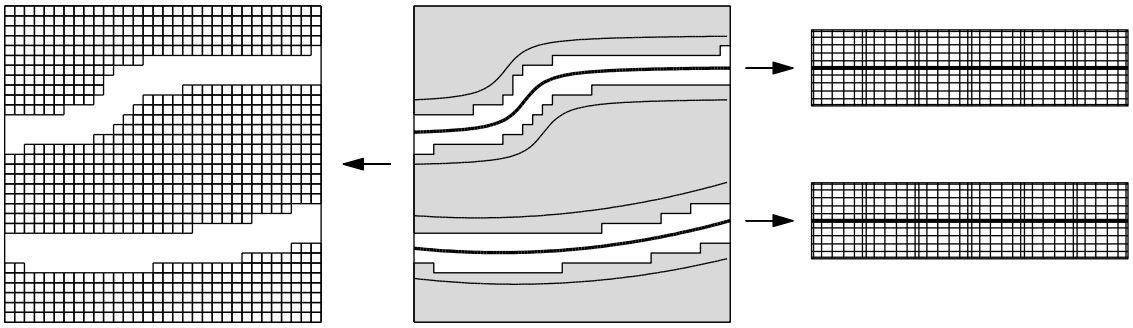

Figure 15: Sketch of a composite grid method using BPS. A high-order finite-difference method is used on the large areas of constant coefficients. In the vicinity of a discontinuity, the region is transformed to a rectangular one in which BPS is applied. The grids are linked by mutual interpolation to determine the extreme grid point values in the overlap.

8 Mulitdimensional applications

It is straightforward to extend the BPS method into tensor-product domains in higher dimensions. The grid lines can be required to line up (i.e., a conforming discretization), and derivative matching is imposed at each line. Because the grids are open rather than closed, no special treatment need be used at the block corners. Experiments with wave propagation over a square confirm that the BPS technique encounters no difficulties. We expect that BPS will be highly competitive with exisitng low-order methods for problems with discontinuous coefficients [1].

For more general domains, the standard spectral element approach is to divide the domain into generalized quadrilaterals that are individually transformed to squares. Derivative matching would require explicit use of high derivatives of these transformations, so it is not easily applicable in this form.

Instead, we propose a combination of overlapping and derivative matching using composite grids [6]. Figure 15 shows a schematic representation of the idea. A high-order finite-difference method on a regular grid is used over the large regions of homogeneous material. Each arbitrary material interface is embedded within a thin region that is separately transformed to a rectangular strip. In this strip, which contains the discontinuity, the BPS method with derivative matching is used. The strips communicate with the background grid via mutual high-order interpolation. Because the interpolation will be used over a relatively small part of the domain, its cost should be negligible. The composite overlapping has at present been successfully tested in 1-D.

We believe that this technique can be used to efficiently solve general linear problems with discontinuous coefficients. We are currently implementing the technique for some 2-D test problems.

References

[1] A. Bamberger, R. Glowinski, and Q. H. Tran, *A domain decomposition method for*

the acoustic wave equation with discontinuous coefficients and grid change, SIAM Journal of Numerical Analysis, 34 (1997), pp. 603–639.

[2] J.-P. Berenger, *A perfectly matched layer for the absorption of electromagnetic waves*, J. Comp. Phys., 114 (1994), pp. 185–200.

[3] M. Bjørhus, *The ODE formulation of hyperbolic PDEs discretized by the spectral collocation method*, SIAM J. Sci. Comput., 16 (1995), pp. 542–557.

[4] C. Canuto and A. Quarteroni, *On the boundary treatment in spectral methods for hyperbolic systems*, J. Comp. Phys., 71 (1987), pp. 100–110.

[5] M. H. Carpenter and D. Gottlieb, *Sepctral methods on arbitrary grids*, J. Comp. Phys., 129 (1996), pp. 74–86.

[6] G. Chesshire and W. D. Henshaw, *Composite overlapping meshes for the solution of partial differential equations*, J. Comp. Phys., 90 (1990), pp. 1–64.

[7] B. Fornberg, *An improved psuedospectral method for initial–boundary value problems*, J. Comp. Phys., 91 (1990), pp. 381–397.

[8] ——, *A Practical Guide to Pseudospectral Methods*, Cambridge Universtiy Press, Cambridge, 1996.

[9] P. Hanley, *A strategy for the efficient simulation of viscous compressible flows using a multi-domain pseudospectral approach*, J. Comp. Phys., 108 (1993), pp. 153–158.

[10] J. S. Hesthaven, *A stable penalty method for the compressible Navier–Stokes equations: II. One-dimensional domain decomposition schemes*, SIAM Journal on Scientific Computing, 18 (1997), pp. 658–685.

[11] D. A. Kopriva, *Computation of hyperbolic equations on complicated domains with patched and overset Chebyshev grids*, SIAM J. Sci. Stat. Comput., 10 (1989), pp. 120–132.

[12] D. A. Kopriva and J. H. Kolias, *A conservative staggered-grid Chebyshev multidomain method for compressible flows*, Tech. Rep. 95–18, ICASE, 1995.

[13] D. Kosloff and H. Tal-Ezer, *A modified Chebyshev pseudospectral method with an $O(N^{-1})$ time step restriction*, J. Comp. Phys., 104 (1993), pp. 457–469.

[14] I. Lie, *Multidomain solution of advection problems by Chebyshev spectral collocation*, J. Sci. Comput., 9 (1994), pp. 39–64.

[15] A. Quarteroni, *Domain decomposition methods for systems of conservation laws: Spectral collocation approximations*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 1029–1052.

[16] A. Taflove, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, Artech House, Boston, 1995.

[17] B. D. Welfert, *Generation of pseudospectral differentiation matrices I*, SIAM Journal of Numerical Analysis, 34 (1997), pp. 1640–1657.

[18] B. Yang, D. Gottlieb, and J. S. Hesthaven, *Spectral simulations of electromagnetic wave scattering*, J. Comp. Phys., 134 (1997), pp. 1–15.

[19] K. S. Yee, *Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media*, IEEE Trans. Antennas and Prop., 14 (1966), pp. 302–307.

[20] D. W. Zingg, H. Lomax, and H. M. Jurgens, *High-accruacy finite-difference schemes for linear wave propagation*, SIAM J. Sci. Comput., 17 (1996), pp. 328–346.