Fourteenth International Conference on Domain Decomposition Methods
Editors: Ismael Herrera , David E. Keyes, Olof B. Widlund, Robert Yates

45. Fast Solvers and Schwarz Preconditioners for Spectral Nédélec Elements for a Model Problem in H(curl)

Bernhard Hientzsch[1]

1. Introduction. In this paper, we present fast solvers and overlapping additive Schwarz methods for spectral Nédélec element discretizations for a model problem in $H(\text{CURL})$. Results of numerical experiments in two dimensions and arguments for a theoretical condition number bound in two and three dimensions are presented. We first show the derivation of the model problem in the implicit time discretization of one of the forms of Maxwell's equation and explain our discretization. We next present the main ideas of two fast direct solvers for such discretizations. We define the overlapping Schwarz methods considered in this paper, present numerical results in two dimensions and then, finally, explain the derivation of a condition number estimate. The estimate obtained for element-wise overlap is quadratic in the relative overlap and the number of colors. In our estimate for general overlap, we obtain an additional factor of N, the degree of the polynomials inside each spectral element.

The model problem is: Find $\mathbf{u} \in H_0(\text{CURL}, \Omega)$ such that for all $\mathbf{v} \in H_0(\text{CURL}, \Omega)$

$$a(\mathbf{u}, \mathbf{v}) := (\alpha \mathbf{u}, \mathbf{v}) + (\beta \text{ CURL } \mathbf{u}, \text{ CURL } \mathbf{v}) = (\mathbf{f}, \mathbf{v}) \tag{1.1}$$

Here, Ω is a bounded, open, connected polyhedron in $\mathbb{R}^3$ or polygon in $\mathbb{R}^2$, $H(\text{CURL}, \Omega)$ is the space of vectors in $(L^2(\Omega))^2$ or $(L^2(\Omega))^3$ with CURL in $L^2(\Omega)$ or $(L^2(\Omega))^3$, respectively; $H_0(\text{CURL}, \Omega)$ is its subspace of vectors with vanishing tangential components on $\partial\Omega$; $\mathbf{f} \in (L^2(\Omega))^d$ for $d = 2, 3$, and $(\cdot, \cdot)$ denotes the inner product in $L^2(\Omega)$ of functions or vector fields. For simplicity, we will assume that α and β are piecewise constant.

2. The model problem, its discretization. The model problem is obtained in several problems in mathematical physics, among them in the time-discretization of several formulations of Maxwell's equation [3, Chapter 3]. For instance, the second order evolution equation for the electric field reads:

$$\epsilon \partial_t^2 \mathbf{E} + \sigma \partial_t \mathbf{E} + \text{ CURL } \left(\frac{1}{\mu} \text{ CURL } \mathbf{E} \right) = \partial_t \mathbf{j}_i$$

If we use an implicit discretization in time, substituting finite differences for the time derivatives and evaluating the terms without time derivatives at the time levels t^k ($k = n+1, n, \cdots$), we obtain

$$\alpha \mathbf{E}^{n+1} + \text{ CURL } \left(\beta \text{ CURL } \mathbf{E}^{n+1} \right) = \mathbf{f}^n$$

where $\mathbf{f}^n$ is constructed from $\mathbf{j}_i$, $\mathbf{E}^k$ and CURL $\mathbf{E}^k$. For instance using leapfrog for the first derivative, a central difference for the second derivative, and Backward Euler for the right hand side, we obtain

$$\alpha = \epsilon + \frac{\sigma}{2}\Delta t \qquad \beta = \frac{1}{\mu}\Delta t^2 \qquad \mathbf{f}^n = \Delta t^2 (\partial_t \mathbf{j}_i)|_{t=t^{n+1}} + 2\epsilon \mathbf{E}^n + \left(\frac{\sigma}{2}\Delta t - \epsilon \right) \mathbf{E}^{n-1}.$$

To use spectral elements, we use the variational form

$$(\alpha \mathbf{E}^{n+1}, \mathbf{F}) + (\beta \text{ CURL } \mathbf{E}^{n+1}, \text{CURL } \mathbf{F}) = (\mathbf{f}^n, \mathbf{F})$$

[1]Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012, E-mail: hientzsc@cims.nyu.edu *or* Bernhard.Hientzsch@na-net.ornl.gov, Homepage: http://www.math.nyu.edu/∼hientzsc. This work was supported in part by the U.S. Department of Energy under Contract DE-FC02-01ER25482.

and we also approximate the L^2-inner products by Gauss-Lobatto-Legendre quadrature.

The problem is posed in $H(\text{CURL})$ and the use of H^1-conforming elements introduces spurious eigenvalues and unphysical continuity conditions. There are approaches to regularize the problem, so that H^1-conforming elements can be used. But in general, for non-convex domains, the solution is not in H^1; so that complicated, weighted formulations have to be used. Nédélec [5, 6] introduced $H(\text{CURL})$-conforming elements, with edge, face and interior moments as degree of freedom; and Ben Belgacem and Bernardi [1] studied methods with spectral element degrees of freedom for Maxwell's evolution equations in a theoretical paper. We will use spectral element type nodal degrees of freedom and enforce only the continuity of the tangential components across element interfaces. For the quadrature, we use variable order Gauss-Lobatto-Legendre formulae.

The discretization on one element results in a block tensor product matrix in both the two-dimensional and three-dimensional case. For simplicity, we consider the two-dimensional case, similar statements hold about the three-dimensional case. The system on one element corresponding to (1.1) reads:

$$K_E\mathbf{u} = \tilde{\mathbf{f}} \qquad \text{or} \qquad \begin{pmatrix} M_1^x \otimes A^y & B^x \otimes C^y \\ B^{x,T} \otimes C^{y,T} & A^x \otimes M_2^y \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = \begin{pmatrix} \tilde{f}_1 \\ \tilde{f}_2 \end{pmatrix}$$

The M are mass matrices, the A are spectral discretizations of scaled Helmholtz operators, B and C are coupling terms between the components involving derivatives and mass matrices.

If we subassemble such elements on a rectangular arrangement with matching polynomial degrees enforcing only tangential continuity, we obtain a global matrix of the same structure, allowing the same solvers as in the element case. For further discussion of the discretization and subassembly procedures we refer to [3, Chapter 8].

The element-by-element (or rectangle-by-rectangle) computation of the matrix-vector product $K\mathbf{u}$ can be implemented by dense matrix-matrix multiplications of the factors of the tensor products with the vector $\mathbf{u}$ laid out in matrix form. It can be implemented by calls to a highly optimized BLAS 3 kernel, and will therefore run at close to peak performance on modern computer architectures.

3. Fast direct solvers. First we treat the case in which we work on one element or a rectangular arrangement of elements with matching polynomial degrees. We transform the block tensor product matrix system on all components into a generalized Sylvester equation for one of the components. Thus, in the two-dimensional case, we eliminate one of the components, say u_2, and obtain:

$$(M_1^x \otimes A^y + C_T \otimes D_T)u_1 = \tilde{f}_1 + (F_T \otimes G_T)\tilde{f}_2$$

with

$$F_T = -B^x(A^x)^{-1} \quad G_T = C^y(M_2^y)^{-1} \quad C_T = F_T C^x \quad D_T = G_T B^y$$

This generalized Sylvester matrix equation can be solved in several different ways. For one, we could solve generalized eigenvalue problems for the matrix pairs (M_1^x, C_T) and (A^y, D_T) and use the eigenvector matrices to diagonalize all factors in the tensor products, and the solution would reduce to a componentwise multiplication with a matrix of the same size as u_1. This is the so-called fast diagonalization method proposed by Lynch et al [4] and used for instance by Tufo and Fischer [9] for the Navier-Stokes equations. There are also Hessenberg and Schur reduction algorithms, see, e.g., Gardiner et al [2], that do not involve the possibly unstable use of transformations to and from the eigenbasis. In our implementation, we used a fast diagonalization method that performs well on most of the examples; we are in the process of testing the method of Gardiner et al.

Once A^x and M_2^y have been factored and the eigenbases and associated transformations have been found, the solution of the Sylvester matrix equation and the back-solve for u_2

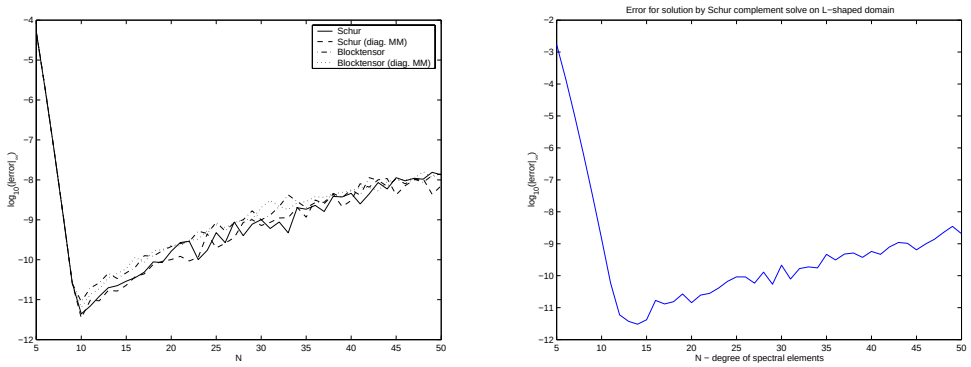

Figure 3.1: Direct solution of (1.1) with $\alpha = 1$, $\beta = 1$. Left: Fast diagonalization solver and Schur interface solver on 5×5 elements of degree $N \times N$, with an exact and with a diagonal mass matrix. Right: Schur interface solver on a L-shaped domain made from three spectral elements.

can be implemented by dense matrix-matrix multiplications and elementwise matrix-matrix multiplication, both kernels from BLAS that are usually available in highly optimized form for the common platforms and run at close to peak performance.

For the case that the above solver is not applicable because the domain is not a rectangle with appropriate polynomial degrees in the elements, we have implemented a direct substructuring method as a direct solver. The local Schur complements can be formed by solving local problems for each (tangential) degree of freedom on the element boundary. We can subassemble the Schur complement and the right hand side for the Schur complement system using the local Schur complements and local solves. The Schur complement system for the tangential components on the interface is then solved, and interior values are found by local (fast) interior solves.

The convergence for two examples is shown in figure 3.1. For further discussion of the implementation, timings, and experiments, we refer to [3, Chapter 9].

4. Overlapping Schwarz methods. To define Schwarz preconditioners in the standard abstract framework (see, e.g., [7]), we have to specify subspaces and solvers on them. We start with a collection of subdomains Ω_i that are either spectral elements themselves or rectangular arrangements thereof. Each subdomain is of size H, each spectral element has uniform degree N in all components. (The analysis goes through for more complicated settings, we chose this case here for simplicity and ease of presentation.) We also define overlapping subregions $\Omega'_{j,\delta} \subset \Omega$ with an overlap of δ. These subregions can be constructed in several ways, e.g., by extending subdomains by a fixed overlap δ in all directions, or by finding vertex centered subdomains that overlap by δ. Most of our early computations (and the numerical results that we show in this paper) were performed on 2×2 vertex centered assemblies of subdomains (taken as single spectral elements), but we will present numerical results for other layouts of $\Omega'_{j,\delta}$ and small overlap in a forthcoming paper.

The local spaces V_j are the linear span of the basis functions associated to Gauss-Lobatto-Legendre points in $\Omega'_{j,\delta}$. In general, the support of functions in V_j will be larger than $\Omega'_{j,\delta}$, but if one only considers the Gauss-Lobatto-Legendre grid, they vanish on grid points outside $\Omega'_{j,\delta}$. On the local spaces we use exact solvers which corresponds to inversion of a submatrix of K. In the 2×2 case the local solve corresponds to the solution of a standard tangential value problem on 2×2 element patches. In any case, the local solve can be implemented

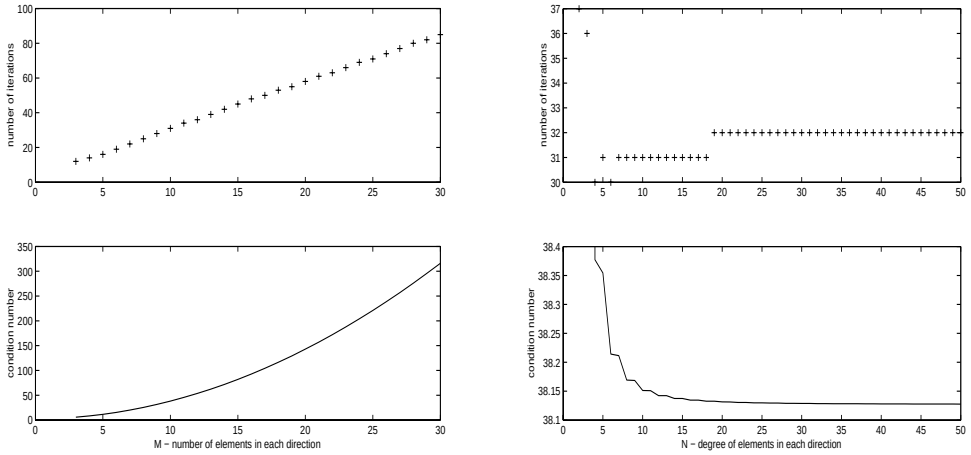

Figure 5.1: One-level method, $\alpha = 1$, $\beta = 1$. Left: Scaling with respect to M, the number of subdomains of degree 10×10. Right: Scaling with respect to N, the degree inside the 10×10 spectral elements.

using the direct fast solvers introduced in the previous section.

The coarse space V_0 is a low-order Nédélec spectral element space of uniform degree N_0 defined on the coarse (subdomain level) mesh. We use the direct solvers of the last section as exact solvers. In the standard way, the local and the coarse solve define local projections T_i and T_0 that can be used to implement different overlapping Schwarz methods. In this paper, we only consider two additive operators: a one-level operator T_{as1} and a two-level operator T_{as2}, defined by

$$T_{as1} = \sum_{i \geq 1} T_i \qquad T_{as2} = T_0 + T_{as1}$$

5. Numerical results in two dimensions. We solve the model problem (1.1) with $\alpha = 1$, $\beta = 1$ on the unit square, decomposed into $M \times M$ subdomains. These subdomains are single spectral Nédélec elements of degree $N \times N$ in all components. The overlapping subregions are patches of 2×2 spectral elements centered around each interior vertex of the spectral element mesh. (Therefore, $\delta = H$.) We have implemented a conjugate gradient method with the fast matrix-vector multiplication $K\mathbf{u}$ by tensor products with the additive Schwarz preconditioners defined by the fast local and coarse solves of section 3. We report the number of iterations needed to decrease the norm of the residual by 10^{-6}, and we also show estimates for the condition number of the preconditioned operator obtained, using the Lanczos connection, from the parameters computed in that conjugate gradient run.

In figure 5.1, we present results for the one-level operator T_{as1}. On the left, we work with degree $N = 10$, and vary the number of subdomains M^2. We see that the iteration numbers grow approximately linearly in M, and that the condition number is growing superlinearly in M. This behavior is to be expected from the absence of a coarse space; the one-level method is not scalable with respect to M. Having fixed the number of subdomains $M^2 = 10 \times 10$, increasing the polynomial degree N improves the condition number slightly, as seen on the right, and the iteration numbers also stay bounded. The condition number stabilizes slightly above 38.

In figure 5.2, we show the results for the two level method T_{as2}. On the left, we show the scaling with M^2 in the case of fixed $N = 10$. Both the iteration numbers and condition

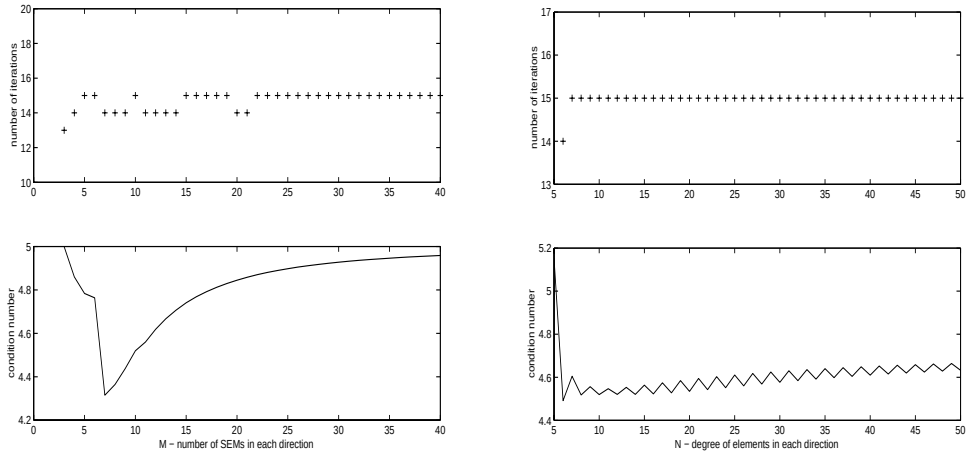

Figure 5.2: Two-level method, $\alpha = 1$, $\beta = 1$, $N_0 = 3$. Left: Scaling with respect to M, the number of subdomains of degree 10×10. Right: Scaling with respect to N, the degree inside the 10×10 spectral elements.

numbers are clearly bounded, the latter being below 5. On the right, we fix the number of subdomains to be 10×10, and vary the polynomial degree N. The condition number is clearly bounded, and less than 4.7 for $N = 50$. This example shows that the two-level additive Schwarz preconditioner T_{as2} is scalable with respect to M, and behaves uniformly in N. The condition numbers (and also the iteration numbers) of the two-level method are considerably smaller than those of the one-level method.

6. Condition number bound. We use the abstract Schwarz framework in which we obtain an estimate of the condition number in terms of N_C (number of colors), ω (norms of local solvers), and C_0^2 (splitting constant). For a general introduction to the abstract Schwarz framework in the context of preconditioners for the h-version, see Smith et al [7, chapter 5].

Since we use exact solvers, the parameter ω is equal to one. The largest eigenvalue is bounded by the number of colors N_C by a standard argument [7]. The smallest eigenvalue is bounded from below by C_0^{-2} where

$$\sum a(\mathbf{u}_i, \mathbf{u}_i) \leq C_0^2 a(\mathbf{u}, \mathbf{u}) \qquad \text{for } \mathbf{u} = \sum \mathbf{u}_i \qquad \text{and } \mathbf{u}_i \in V_i$$

We will estimate the smallest eigenvalue by a variant of the arguments in Toselli [8]. To complete the argument, we need to make explicit the N-dependence of the different bounds in Toselli's argument. The bounds were proven in the low-order case with a finite-dimensional space argument which only allows fixed N. For varying N the dimension of the space is related to N, and those arguments can not easily be extended to include the dependence on N.

We will sketch the analysis briefly and refer for complete details to a forthcoming paper.

Our variant of Toselli's analysis is expressed in the following theorem:

Theorem 6.1 (N-dependence of the condition number) *Given a Nédélec interpolation estimate on divergence-free functions from H(curl) with polynomial curl of the form:*

$$||(I - \mathbf{\Pi}_N^{ND,I})\mathbf{w}||_0 \leq Chf_1(N)||\mathrm{CURL}\mathbf{w}||_0 \tag{6.1}$$

the L^2-stability of the local splitting:

$$||\mathbf{\Pi}_N^{ND,I}(\chi_i \mathbf{u})||_0 \leq C f_2(N) ||\chi_i \mathbf{u}||_0 \tag{6.2}$$

the curl-stability of the local splitting:

$$||\mathrm{CURL}\left(\mathbf{\Pi}_N^{ND,I}(\chi_i \mathbf{u})\right)||_0 \leq C f_3(N) ||\mathrm{CURL}(\chi_i \mathbf{u})||_0 \tag{6.3}$$

and the standard conditions on a partition of unity χ_i:

$$||\chi_i||_\infty \leq C \qquad ||\nabla \chi_i||_\infty \leq \frac{C}{\delta} \tag{6.4}$$

then, the inverse of the smallest eigenvalue is bounded by

$$\max\left(CN_c\left(1+\frac{H}{\delta}\right), C\frac{\max(\alpha,\beta)}{\min(\alpha,\beta)}(1+N_c f_2^2(N)),\right.$$
$$\left. C\frac{\max(\alpha,\beta)}{\min(\alpha,\beta)}\left\{1+N_c f_3^2(N)\left(1+\left(\frac{H+hf_1(N)}{\delta}\right)^2\right)\right\}\right)$$

We will now discuss the different assumptions and estimates required and their proof.

The interpolation estimate (6.1) can be proven with $f_1(N) = 1 + C(\epsilon)N^{-1+\epsilon} \leq 1 + C(\epsilon)$. In two dimensions, $f_1(N) = C(\epsilon)N^{-1+f(\epsilon)}$. A similar bound should also hold in three dimensions, but is not yet proven.

We can reduce the L^2-stability of the local splitting (6.2) for polynomial χ_i to the L^2-stability of the Nédélec interpolation operator between polynomial spaces. Then we need to use χ_i that are polynomial interpolants of the standard piecewise linear ones. By the definition of the V_i, χ_i needs to be zero on the Gauss-Lobatto-Legendre grid points outside of $\Omega'_{i,\delta}$. We need to prove that they satisfy the conditions on the partition of unity (6.4). By a general theorem, we would obtain bounds involving the Lebesgue constant for the Gauss-Lobatto-Legendre interpolation, i.e., a logarithmic factor in N; numerical experiments and special arguments using the specific form of the Gauss-Lobatto-Legendre point values of the partition of unity result in the same bounds as for the standard piecewise linear partition of unity.

The L^2-stability of the Nédélec interpolant (6.2) can be proven by identifying the Nédélec interpolant as a componentwise tensor product operator having as factors discrete L^2- and modified H^1-projections. The L^2-projections are obviously stable, the modified H^1-projections have been proven, by a recent result of ours, to be stable for fixed differences in the degrees, and linear in the square root of the difference in the degree otherwise. We also confirmed the predicted behavior of the modified H^1-projections in extensive numerical experiments. Combining the results from the factors, we obtain a constant bound for $f_2(N)$ for χ_i of fixed degree (for instance for element-wise overlap), and a bound $f_2(N) = \sqrt{c_1 N + c_2}$ for χ_i of degree N (i.e., for general overlap).

The CURL-stability of the Nédélec interpolant (6.3) of the local splitting can be reduced to the L^2-stability of the Raviart-Thomas interpolant by the commuting diagram property. Therefore, for polynomial χ_i, it is enough to analyze the L^2-stability of the Raviart-Thomas interpolant between polynomial spaces of different degrees. One can identify the Raviart-Thomas interpolant as a componentwise tensor product operator having as factors discrete L^2- and modified H^1-projections. These factors are of the same form as for the Nédélec interpolant. Our recent results for the one-dimensional projections therefore imply, analogously to the results of the previous paragraph, a constant bound for $f_3(N)$ for χ_i of fixed degree and a bound $f_3(N) = \sqrt{c_3 N + c_4}$ for χ_i of degree N (i.e., for general overlap).

Combining the estimates with the above theorem, we obtain two corollaries:

Corollary 6.1 (Element-wise overlap) *In the case of element-wise overlap, the condition number of T_{as2} is bounded by*

$$\kappa(T_{as2}) \leq C(N_c+1)\frac{\max(\alpha,\beta)}{\min(\alpha,\beta)}\left(1+N_c\left(1+\left(\frac{H}{\delta}\right)^2\right)\right)$$

Corollary 6.2 (General overlap) *For general δ, an upper bound of the condition number of T_{as2} is given by*

$$\kappa(T_{as2}) \leq C(N_C+1)N\frac{\max(\alpha,\beta)}{\min(\alpha,\beta)}\left(1+N_C\left(1+\left(\frac{H}{\delta}\right)^2\right)\right)$$

It is not known if the powers of $\frac{H}{\delta}$ and N_C in both corollaries and of N in the second corollary are optimal, and it is not known if the estimates are sharp. We are performing a numerical study of the extreme eigenvalues for minimal and for fixed overlap which should give us some insight about sharpness and exponents.

For the limit cases $\alpha \to 0$ or $\beta \to 0$, one can find improved bounds by alternative splittings and proofs.

REFERENCES

[1] F. B. Belgacem and C. Bernardi. Spectral element discretization of the Maxwell equations. *Math. Comp.*, 68(228):1497–1520, 1999.

[2] J. D. Gardiner, A. J. Laub, J. J. Amato, and C. B. Moler. Solution of the Sylvester matrix equation $AXB^T + CXD^T = E$. *ACM Trans. Math. Software*, 18(2):223–231, 1992.

[3] B. Hientzsch. *Fast Solvers and Domain Decomposition Preconditioners for Spectral Element Discretizations of Problems in H(curl).* PhD thesis, Courant Institute of Mathematical Sciences, September 2001. also Technical Report TR2001-823, Department of Computer Science, Courant Institute.

[4] R. E. Lynch, J. R. Rice, and D. H. Thomas. Direct solution of partial difference equations by tensor product methods. *Numer. Math.*, 6:185–199, 1964.

[5] J.-C. Nédélec. Mixed finite elements in R^3. *Numer. Math.*, 35:315–341, 1980.

[6] J.-C. Nédélec. A new family of mixed finite elements in R^3. *Numer. Math.*, 50:57–81, 1986.

[7] B. F. Smith, P. E. Bjørstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations.* Cambridge University Press, 1996.

[8] A. Toselli. Overlapping Schwarz methods for Maxwell's equations in three dimensions. *Numer. Math.*, 86(4):733–752, 2000.

[9] H. M. Tufo and P. F. Fischer. Terascale spectral element algorithms and implementations. In *Proceedings of the ACM/IEEE SC99 Conference*, 1999.