Accelerating simulation of ensembles of locally differing optical structures via a Schur complement domain decomposition

Sacha Verweij,[1,2,*] Victor Liu,[3] and Shanhui Fan[4]

[1]*Department of Applied Physics, Stanford University, Stanford, California 94305, USA*

[2]*Institute for Computational & Mathematical Engineering, Stanford University, Stanford, California 94305, USA*

[3]*Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, California 94304, USA*

[4]*Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA*

**Corresponding author: sacha@stanford.edu*

Received August 22, 2014; accepted September 17, 2014;
posted October 14, 2014 (Doc. ID 221501); published November 10, 2014

We present a Schur complement domain decomposition method that can significantly accelerate simulation of ensembles of locally differing optical structures. We apply the method to design a multi-spatial-mode photonic crystal waveguide splitter that exhibits high transmission and preservation of modal content, showing design acceleration by more than a factor of 20. © 2014 Optical Society of America

OCIS codes: (000.4430) Numerical approximation and analysis; (050.1755) Computational electromagnetic methods; (220.0220) Optical design and fabrication; (350.4238) Nanophotonics and photonic crystals.
http://dx.doi.org/10.1364/OL.39.006458

Efficient simulation of ensembles of closely related optical structures is broadly important in optical structure design [1–8]. In such ensembles, often all member structures are identical on some subdomain; the ensembles of [1–8] evidence this, and Fig. 1 provides an exemplar in photonic crystal structure design. In spite of this feature's prevalence, its exploitation appears uncommon. Typically, each member structure is simulated in its entirety, including the subdomain on which all member structures are identical. Repeated simulation of this *constant* subdomain constitutes a significant inefficiency. Ideally this inefficiency would be eliminated, leaving only computation associated with the *variable* subdomain to be performed for each member structure.

In this Letter, we illustrate mitigation of this inefficiency via a domain decomposition strategy. Specifically, we present a classical Schur complement domain decomposition method which, where suitable, alleviates this inefficiency, and apply it to simulation of an ensemble of photonic crystal structures arising in a design problem. As our primary aim is to illuminate the concept, we begin by presenting the method somewhat abstractly and thereafter describe its application more concretely. The ideas and mechanics discussed have a long history outside optics; we discuss this in closing.

Conceptually, the method proceeds as follows: For each member structure, we decompose simulation of the complete structure into a pair of subproblems, one associated with the constant subdomain and the other with the variable subdomain. Being the same for all member structures, the subproblem associated with the constant subdomain needs to be solved only once; evaluating a given member structure thereby reduces to solution of only the subproblem associated with the variable subdomain.

To begin, we assume that each member structure is simulated by solving a square linear system of block form

$$\begin{bmatrix} A_{vv} & A_{vp} & 0 \\ A_{pv} & A_{pp} & A_{pc} \\ 0 & A_{cp} & A_{cc} \end{bmatrix} \begin{pmatrix} v \\ p \\ c \end{pmatrix} = \begin{pmatrix} b_v \\ b_p \\ b_c \end{pmatrix}, \tag{1}$$

where the partitioning and last block row and column do not vary between member structures, and the unknown vectors v, p, and c are the degrees of freedom of the discretization scheme applied to the member structures. In this formalism, v and c are generally those degrees of freedom associated with the variable and constant subdomains, respectively, and p those associated with boundaries that partition the domain into these two subdomains. Typically, such description is possible when the discretization scheme is local in three senses: (1) the discretization scheme's degrees of freedom are associated with spatially local entities; (2) the discretization scheme couples only the degrees of freedom associated with some common spatial locale; and (3) the couplings' dependencies on the structure's physical properties are also confined to some common spatial locale. This is the case, for example, with common frequency-domain finite difference and finite element discretizations.

Given a set of such linear systems, our aim is to efficiently construct a corresponding set of condensed linear systems in which the blocks that do not vary between member structures have been eliminated. Considering a given system in the set, we accomplish this aim via block elimination as follows:

Where A_{cc} is nonsingular, solving the system's last block row for c provides

$$c = A_{cc}^{-1} b_c - A_{cc}^{-1} A_{cp} p. \tag{2}$$

After substituting this expression for c in the system's first two block rows, minor rearrangement yields the condensed system

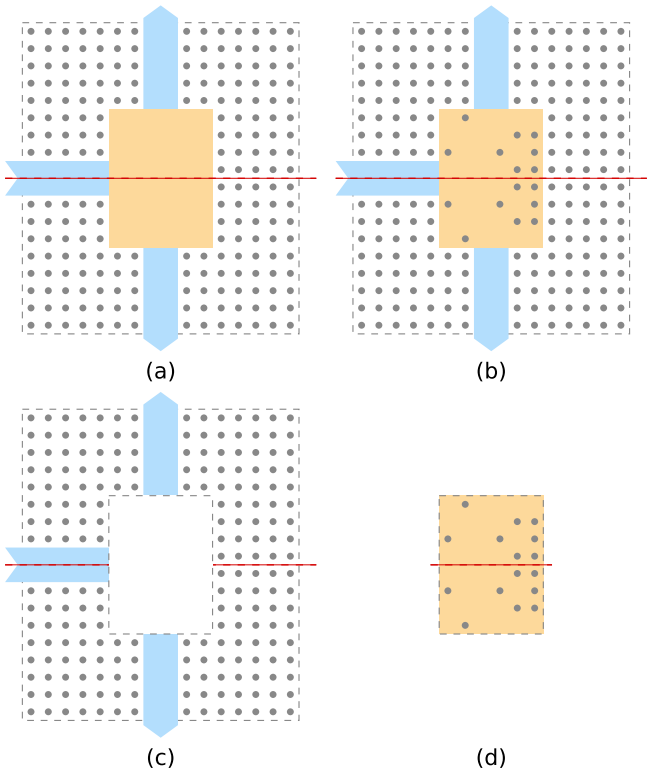

Fig. 1. (a) Design prototype for a multi-spatial-mode photonic crystal waveguide splitter that preserves modal content. A square lattice of silicon cylinders (gray circles, $r/a = 0.2$, $n = 3.4$) in air (white backdrop), which exhibits a transverse-magnetic (TM) mode bandgap over frequencies $0.29{\cdot}2\pi c/a$ to $0.42{\cdot}2\pi c/a$, comprises the background crystal. Highlighted in blue, three line-defect waveguides each support two TM spatial modes over frequencies $0.36{\cdot}2\pi c/a$ to $0.42{\cdot}2\pi c/a$. Highlighted in orange, a junction region couples the horizontally oriented input waveguide on the left to the vertically oriented output waveguides at top and bottom. Indicated by the red dashed line, a horizontal mirror plane bisects the prototype. We generate an instance of this prototype, such as that shown in (b), by placing silicon cylinders, identical to those in the prototype, at a set of lattice sites in the junction region. Between instances of the prototype, only the junction region (d) varies; the rest of the structure (c) remains the same.

$$\left(\begin{bmatrix} A_{vv} & A_{vp} \\ A_{pv} & A_{pp} \end{bmatrix} - \begin{bmatrix} 0 \\ A_{pc} \end{bmatrix} A_{cc}^{-1} \begin{bmatrix} 0 & A_{cp} \end{bmatrix}\right)\begin{pmatrix} v \\ p \end{pmatrix} = \begin{pmatrix} b_v \\ b_p \end{pmatrix} - \begin{bmatrix} 0 \\ A_{pc} \end{bmatrix} A_{cc}^{-1} b_c. \tag{3}$$

The matrix appearing above on the left is known as the *Schur complement* of A_{cc} in A. The Schur complement appears routinely in block elimination [9] and forms the heart of classical non-overlapping domain decomposition methods. Exploiting the zero blocks, the condensed system simplifies to

$$\begin{bmatrix} A_{vv} & A_{vp} \\ A_{pv} & A_{pp} - A_{pc}A_{cc}^{-1}A_{cp} \end{bmatrix}\begin{pmatrix} v \\ p \end{pmatrix} = \begin{pmatrix} b_v \\ b_p - A_{pc}A_{cc}^{-1}b_c \end{pmatrix}. \tag{4}$$

Being the same for all systems in the set, the terms $\Sigma = A_{pc}A_{cc}^{-1}A_{cp}$ and $\sigma = A_{pc}A_{cc}^{-1}b_c$ in Eq. (4) and $\Gamma = A_{cc}^{-1}A_{cp}$ and $\gamma = A_{cc}^{-1}b_c$ in Eq. (2) need be computed only once. Having precomputed these objects, evaluating a given member structure reduces to solving the condensed system

$$\begin{bmatrix} A_{vv} & A_{vp} \\ A_{pv} & A_{pp} - \Sigma \end{bmatrix}\begin{pmatrix} v \\ p \end{pmatrix} = \begin{pmatrix} b_v \\ b_p - \sigma \end{pmatrix} \tag{5}$$

and recovering the remaining unknowns from $c = \gamma - \Gamma p$. Often, however, only a small subset of the unknowns c associated with the constant subdomain are of interest. This occurs, for example, in design optimization problems where the figure of merit is dependent solely on the structure's behavior on a fragment of the constant subdomain. In such a case, one need only evaluate the restriction of this expression for c to the desired subset r,

$$r = \gamma|_r - \Gamma|_r p, \tag{6}$$

possibly significantly reducing cost. Furthermore, if only the result of applying a linear map to r is desired, for example to compute a figure of merit, then precomputing application of that linear map to $\gamma|_r$ and $\Gamma|_r$ in Eq. (6) may yield additional cost reduction. Naturally, similar statements hold for maps of r that are not linear.

The presented method can significantly decrease the cost of simulating each member structure: Without the presented method, simulating each member structure requires solving Eq. (1), a size $(N_v + N_p + N_c)$ square linear system (N_x denotes the number of elements in x). The presented method reduces this computation to solving Eq. (5), a size $(N_v + N_p)$ square linear system, and evaluating Eq. (6), an $N_r \times N_p$ matrix-vector product.

To demonstrate the presented approach's utility, we show an accelerated design of a multi-spatial-mode photonic crystal waveguide splitter that preserves modal content. Multi-spatial-mode photonic crystal devices that preserve modal content, for instance the waveguide bends designed in [8], hold promise for mode-division-multiplexing in future optical interconnects.

Specifically, here we consider the design prototype illustrated and described in Fig. 1(a). We generate an instance of this prototype, such as that shown in Fig. 1(b), by placing silicon cylinders, identical to those in the prototype, at a set of lattice sites in the prototype's junction region. 2^{24} instances are consistent with the prototype's horizontal mirror symmetry; this ensemble forms our design space. Whereas the junction region [Fig. 1(d)] varies between instances, the rest of the structure [Fig. 1(c)] remains the same; accordingly, the former and latter regions, respectively, constitute the presented method's variable and constant subdomains.

We used the Dirichlet-to-Neumann map method detailed in [8,10,11] for discretization. This scheme's degrees of freedom are field values on crystal cell boundaries. Additionally, this scheme couples only degrees of freedom which share a crystal cell interior, and the couplings depend only on the shared crystal cell's physical properties. Respectively, these three features make the discretization scheme spatially local in the three senses discussed above. As a result, by applying this scheme to the ensemble of instances and identifying the unknowns on the junction region (the variable subdomain) with v,

those on the rest of the structure (the constant subdomain) with c, and those on the boundary separating these two subdomains with p, we obtained a set of linear systems of the form of Eq. (1).

By casting these systems as low-rank modifications of an initial system, we solved individual systems in $\mathcal{O}(N^2)$ time via a low-rank updating method (see, for example, [7,12]). This was the case whether we solved the original or condensed systems, having sizes $(N_v + N_p + N_c)$ and $(N_v + N_p)$, respectively. As solving these systems dominated the computation in either case, we expected the presented method to provide a performance improvement approaching the ratio $(N_v + N_p + N_c)^2/(N_v + N_p)^2$.

Here, $(N_v + N_p + N_c)$ and $(N_v + N_p)$ are proportional, with the same constant of proportionality, to the number of discretized cell edges, respectively, associated with the complete structure and variable subdomain, in this case 578 and 110. As such, the above performance improvement estimate evaluates to ~27.6.

Characterizing each instance required only the unknowns associated with the waveguide ports; these unknowns constituted the restricted set r of recovered unknowns on the constant subdomain. Because r was small relative to v, the cost of recovering r was negligible relative to the cost of solving the condensed systems.

When solving the original systems, our design code characterized ~7.7 instances per second per test frequency on a single Xeon E5-2680 core; though our implementation of the presented method was suboptimal, this rate increased to ~181.8 when solving the condensed systems, a factor of ~23.6 improvement. Characterizing all 2^{24} admissible instances required only ~34 core-hours per test frequency on 10 compute nodes with two eight-core Xeon E5-2680 processors each; without benefit of the method, this would have required ~800 core-hours. Differences between powers calculated via the condensed and original systems generally remained several orders below the accuracy of the discretization scheme.

Testing at 13 frequencies linearly spaced from $0.36{\cdot}2\pi c/a$ to $0.42{\cdot}2\pi c/a$, we found only seven instances exhibiting, at minimally one test frequency, modal-content-preserving transmission exceeding 47% to each output waveguide. Figure 1(b) shows the defect configuration of a particularly broadband, high-performing instance: Over a 20 nm band at 1.55 μm, this instance exhibits modal-content-preserving transmission exceeding 48% to each output waveguide and maintains less than 2% modal crosstalk. Figure 2 provides this instance's modal power transmission spectrum. Figures 3(a) and 3(b), respectively, illustrate its field pattern on excitation of the first and second spatial modes of the input waveguide. To attain higher performance or bandwidth, one could enlarge the discrete design space, or continuously optimize, for example, cylinder radii [13].

Two primary factors impact the computational cost reduction the method can yield. First, we expect significant computational cost reduction only where both the unknowns c associated with the constant subdomain constitute a significant fraction of each linear system's unknowns, and the size of r recovered on the constant subdomain is small relative to c; this generally corresponds to the constant subdomain comprising a substantial portion of each structure, and interest in the solution on the constant subdomain being confined to small parts of that subdomain. Second, most discretization schemes produce linear systems with properties that enable efficient solution. Where these properties are diminished in the condensed linear systems, solving the condensed linear systems may require methods with higher complexity. In this case, the higher complexity may offset the computational cost reduction resulting from the condensed linear systems' decreased size. That said, the condensed linear systems often preserve the original linear systems' properties. For example, symmetry is preserved in the condensed linear systems. Similarly, only sparsity within the block A_{pp} is impacted, and this block is typically relatively small. As such, where efficient solution methods for the original linear systems exist, it may be possible to adapt them to the condensed linear systems, as the example presented above illustrates.

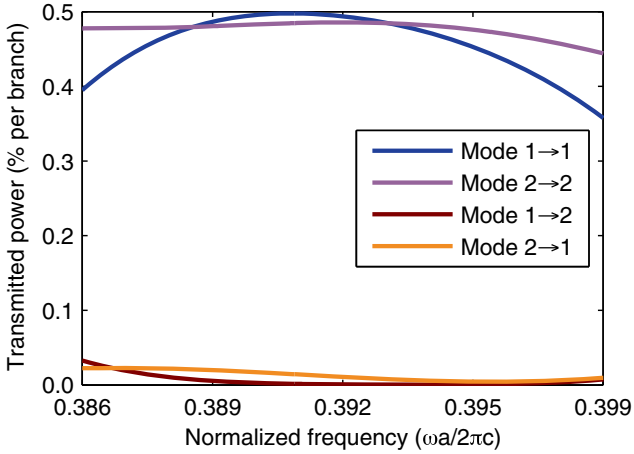

Fig. 2. Modal power transmission spectrum of the selected instance [Fig. 1(b)]. The blue line indicates the proportion of power transmitted from the first TM spatial mode on the input waveguide to the same mode on either of the output waveguides. The purple line indicates the analog for the second TM spatial mode. The red line indicates the proportion of power transmitted from the first TM spatial mode on the input waveguide to the second TM spatial mode on either of the output waveguides (modal crosstalk). The orange line indicates the analog for the second TM spatial mode.

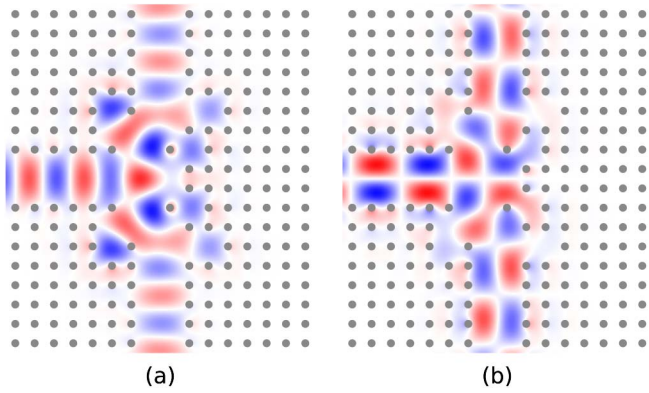

Fig. 3. Out-of-plane electric field patterns resulting from excitation of (a) the first and (b) the second TM spatial modes on the input waveguide of the selected instance [Fig. 1(b)] at frequency $0.392{\cdot}2\pi c/a$. Preservation of modal content is evident.

Concerning memory consumption, the method replaces the size $(N_v + N_p + N_c)$ square matrix from Eq. (1), which nominally is sparse, with the size $(N_v + N_p)$ square matrix from Eq. (5), which is identically sparse but for fill-in in block A_{pp}, and introduces the $N_r \times N_p$ rectangular matrix from Eq. (6), which is dense. Where the method is likely to provide significant computational cost reduction, such as under the conditions outlined in the preceding paragraph, this transformation typically should lead to an acceptable impact on storage. Where matrices are stored in dense form, memory consumption commonly should decrease.

Finally, we note that the ideas and mechanics discussed in this Letter have a long history outside optics. Non-overlapping domain decomposition primarily grew out of seminal papers by Przemieniecki in structural analysis [14] and Buzbee *et al.* in numerical partial differential equations [15,16]. Originally developed to enable piecemeal treatment of problems that could not be tackled globally, Przemieniecki's *substructuring*, to which the method presented here is mechanically similar, flourished in the structural analysis community, giving rise to and interweaving with conceptually and mathematically similar developments in model reduction such as Guyan's *static condensation* [17] and Hurty's *component mode synthesis* [18]. The work of Buzbee *et al.*, with which the work in structural analysis eventually converged, led to a vast and still-growing body of techniques for solving linear systems arising from discretization of partial differential equations.

As Bjørstad and Widlund note in [19], the structural analysis community recognized domain decomposition's utility in various forms of repeated analysis within 20 years of Przemieniecki's development. This realization has since made its way into the broader computational electromagnetics community; see, for example, similar work in the microwave domain [20]. With few exceptions such as [21], however, this particular utility of domain decomposition does not appear well known in the optics community, though domain decomposition's more traditional uses in parallelization and piecemeal handling of large problems certainly are. In summary, here we have shown the value of domain decomposition techniques for accelerating repeated analysis in optical structure design.

We thank Dr. W. Shin, J. A. Briggs, J. E. Herriman, and the anonymous reviewers for suggestions that considerably improved the manuscript. This research was supported in part by the Department of Energy Office of Science Graduate Fellowship Program (DOE SCGF), made possible in part by the American Recovery and Reinvestment Act of 2009, administered by ORISE-ORAU under contract no. DE-AC05-06OR23100. This research was also supported in part by the National Science Foundation through XSEDE resources provided by the XSEDE Science Gateways program, and the United States Air Force Office of Scientific Research (USAFOSR) grant FA9550-09-1-0704.

References

1. D. Taillaert, P. Bienstman, and R. Baets, Opt. Lett. **29**, 2749 (2004).
2. R. Pala, J. Liu, E. Barnard, D. Askarov, E. Garnett, S. Fan, and M. Brongersma, Nat. Commun. **4**, 2095 (2013).
3. P. Borel, A. Harpøth, L. Frandsen, M. Kristensen, P. Shi, J. Jensen, and O. Sigmund, Opt. Express **12**, 1996 (2004).
4. X. Jiao, J. Goeckeritz, S. Blair, and M. Oldham, Plasmonics **4**, 37 (2009).
5. X. Sheng, S. Johnson, J. Michel, and L. Kimerling, Opt. Express **19**, A841 (2011).
6. H. Kim, J. Park, and B. Lee, Opt. Lett. **34**, 2569 (2009).
7. Y. Jiao, S. Fan, and D. Miller, Opt. Lett. **30**, 141 (2005).
8. V. Liu and S. Fan, Opt. Express **21**, 8069 (2013).
9. R. Cottle, Linear Algebra Appl. **8**, 189 (1974).
10. Z. Hu and Y. Lu, Opt. Express **16**, 17383 (2008).
11. V. Liu, D. Miller, and S. Fan, Proc. IEEE **101**, 484 (2013).
12. W. Hager, SIAM Rev. **31**, 221 (1989).
13. Y. Jiao, "Topology optimization and fabrication of photonic crystal structures," Ph.D. thesis (Stanford University, 2005).
14. J. Przemieniecki, AIAA J. **1**, 138 (1963).
15. B. Buzbee, G. Golub, and C. Nielson, SIAM J. Numer. Anal. **7**, 627 (1970).
16. B. Buzbee, F. Dorr, J. George, and G. Golub, SIAM J. Numer. Anal. **8**, 722 (1971).
17. R. Guyan, AIAA J. **3**, 380 (1965).
18. W. Hurty, AIAA J. **3**, 678 (1965).
19. P. Bjorstad and O. Widlund, SIAM J. Numer. Anal. **23**, 1097 (1986).
20. S. Selleri, O. Farle, G. Guarnieri, M. Losch, G. Pelosi, and R. Dyczij-Edlinger, in *Proceedings of IEEE Antennas and Propagation Society International Symposium* (IEEE, 2008).
21. M. Hammerschmidt, D. Lockau, L. Zschiedrich, and F. Schmidt, Proc. SPIE **8980**, 898007 (2014).