Restrictions on the use of sweeping type preconditioners for Helmholtz problems

Martin J. Gander and Hui Zhang

1 Introduction

Helmholtz problems, and time harmonic problems in general like Maxwell's equations, are notoriously difficult to solve numerically. The first problem is that they require very fine discretizations to avoid the so called pollution effect [1], and then the discretized systems are so large that one needs to solve them iteratively, and none of the classical iterative methods are suitable for this task [10]. Over the past decade, several new ideas arrived for the iterative solution of Helmholtz problems, among them the shifted Laplace preconditioner [9]. Unfortunately in this preconditioner, one has to choose the shift small enough (at most $O(k)$ where k is the wave number) for the preconditioner to be close to the underlying operator to give provable wave number independent convergence [12], and large enough (at least $O(k^2)$) for the preconditioner to be easily invertible by multigrid independently of the wave number [5, 6]. In practice, a compromise has to be chosen, which can lead to a growth of up to $O(k^2)$ in the iteration numbers of preconditioned GMRES in the multigrid case [6]; for a rigorous analysis in the case of classical domain decomposition, see [20]. The best current preconditioners are based on domain decomposition methods using special transmission conditions, and have their roots in optimized Schwarz methods [14, 13] and the AILU preconditioner [15, 16]. These algorithms use transmission conditions adapted to the underlying Helmholtz nature of the problem, and this idea is so important that it has been rediscovered independently several times over the last few years, see the sweeping preconditioner [7, 8], the source transfer method, the methods based on single layer potentials [3, 4], and most recently the method

Martin J. Gander
University of Geneva, Section of Mathematics, Rue du Lievre 2-4, CP 64, 1211 Geneva 4, e-mail: martin.gander@unige.ch

Hui Zhang (corresponding author)
Zhejiang Ocean University, Key Laboratory of Oceanographic Big Data Mining & Application of Zhejiang Province, Zhoushan 316022, China, e-mail: huiz@zjou.edu.cn

of polarized traces [24, 23]. All these methods use the same underlying mathematical algorithm, which at the continuous level is the class of optimal Schwarz methods [11], and at the discrete level the block-LU factorization, and one can prove formally that they are all basically equivalent, see the review monograph [19]. The methods use a one way decomposition of the domain into a sequence of subdomains, and between subdomains they use as transmission condition an approximation of the Dirichlet to Neumann operator. An important technique advocated by these more recently proposed algorithms is the use of perfectly matched layers (PML) in the transmission conditions; for an earlier use of PML transmission conditions in a domain decomposition setting, see [22, 21], and [2] for high order Padé transmission conditions, with [17, 18] for their relation to PML transmission conditions. While one might think intuitively that the absorption at the interfaces is the most important property, and with PML one can reach as much absorption as one wants, the truly important property for the algorithm is not absorption, but approximation of the Dirichlet to Neumann operator, which is well known from optimized Schwarz theory [11]. For a constant wave number, these two coincide, and it was therefore possible to prove for the above methods that they can be made into arbitrarily good solvers by improving the PML, but this holds only for constant wave number. We show here that like all the other iterative Helmholtz solvers so far, the performance of these methods deteriorates as soon as the approximation of the Dirichlet to Neumann operator is not perfect any more in the case of wave propagation. To do so, we use a common algorithm formulation at the discrete level from [19], and provide the algorithm without any of the technicalities related to the various inventions, so that anybody can implement and check the method for themselves.

2 Common formulation of sweeping, source transfer, single layer, polarized traces and optimal/optimized Schwarz algorithms

To illustrate the limitations of these methods, it suffices to take the Helmholtz equation in a layered medium,

$$(\Delta + k(x)^2)u = f, \quad \text{in } \Omega := (0,1)^2, \tag{1}$$

with suitable boundary conditions for well posedness, such that after discretization by a standard five point finite difference method, the piecewise constant wave speeds are aligned with the block tridiagonal matrix structure

$$A\mathbf{u} := \begin{bmatrix} D_1 & L & & & \\ L & D_2 & L & & \\ & \ddots & \ddots & \ddots & \\ & & L & D_{J-1} & L \\ & & & L & D_J \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_{J-1} \\ \mathbf{u}_J \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_{J-1} \\ \mathbf{f}_J \end{bmatrix} =: \mathbf{f}. \tag{2}$$

The block LU factorization of the coefficient matrix in (2) is given by

$$A = \begin{bmatrix} T_1 & & & & \\ L & T_2 & & & \\ & \ddots & \ddots & & \\ & & L & T_{J-1} & \\ & & & L & T_J \end{bmatrix} \begin{bmatrix} I & T_1^{-1}L & & & \\ & I & T_2^{-1}L & & \\ & & \ddots & \ddots & \\ & & & I & T_{J-1}^{-1}L \\ & & & & I \end{bmatrix}, \tag{3}$$

where the T_j are the Schur complements that satisfy the recurrence relation

$$T_1 = D_1, \quad T_j = D_j - LT_{j-1}^{-1}L \text{ for } j \geq 2, \tag{4}$$

as one can verify by a direct calculation. The underlying system (2) can then be solved by a forward block substitution, followed by a backward block substitution, which corresponds to the sweeping over the domain back and forth, the source transfer from layer to layer, or the alternating solution over subdomains in the optimized Schwarz setting, see [19]. In the constant wave number case, the dense blocks T_j can be implemented using PML to arbitrary precision[1], and then all these sweeping type methods can be made arbitrarily close to being direct solvers, which explains their excellent performance in the constant wave number case. In the variable wave number case however, the best a PML can do is to be perfectly absorbing for the neighboring medium, assuming it to be constant up to infinity. To get such a perfect absorption for our model problem directly algebraically, without PML techniques, we consider for each wave number block D_i the constant coefficient problem

$$A^i\mathbf{u}^i := \begin{bmatrix} D_i & L & & & \\ L & D_i & L & & \\ & \ddots & \ddots & \ddots & \\ & & L & D_i & L \\ & & & L & D_i \end{bmatrix} \begin{bmatrix} \mathbf{u}_1^i \\ \mathbf{u}_2^i \\ \vdots \\ \mathbf{u}_{J-1}^i \\ \mathbf{u}_J \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_{J-1} \\ \mathbf{f}_J \end{bmatrix} =: \mathbf{f}, \tag{5}$$

with factorization

$$A^i = \begin{bmatrix} T_1^i & & & & \\ L & T_2^i & & & \\ & \ddots & \ddots & & \\ & & L & T_{J-1}^i & \\ & & & L & T_J^i \end{bmatrix} \begin{bmatrix} I & (T_1^i)^{-1}L & & & \\ & I & (T_2^i)^{-1}L & & \\ & & \ddots & \ddots & \\ & & & I & (T_{J-1}^i)^{-1}L \\ & & & & I \end{bmatrix}, \tag{6}$$

where T_j^i are the Schur complements that satisfy now the recurrence relation

$$T_1^i = D_i, \quad T_j^i = D_i - L(T_{j-1}^i)^{-1}L \text{ for } j \geq 2. \tag{7}$$

[1] provided the domain has indeed an open end or such a high order PML on the side where the sweeping begins.

Then the approximate factorization using this best possible approximation a PML technique could provide[2] is

$$\tilde{A} = \begin{bmatrix} \tilde{T}_1 & & & & \\ L & \tilde{T}_2 & & & \\ & \ddots & \ddots & & \\ & & L & \tilde{T}_{J-1} & \\ & & & L & \tilde{T}_J \end{bmatrix} \begin{bmatrix} I & \tilde{T}_1^{-1}L & & & \\ & I & \tilde{T}_2^{-1}L & & \\ & & \ddots & \ddots & \\ & & & I & \tilde{T}_{J-1}^{-1}L \\ & & & & I \end{bmatrix}, \tag{8}$$

where $\tilde{T}_j$ are the Schur complements using the exact Schur complements of the neighboring constant wave number case, namely

$$\tilde{T}_1 = D_1, \quad \tilde{T}_j = D_j - L(T_{j-1}^{j-1})^{-1}L \text{ for } j \geq 2. \tag{9}$$

Note that this best possible information a PML could provide is not necessarily a good approximation to the Dirichlet to Neumann operator which is represented by the exact blocks T_j, and thus contains information about all the reflections that will be created by all the layers outside the present subdomain. We will test now how much variation in the wave number this approximation can tolerate before the sweeping type algorithms loose their effectiveness, and how this depends on the source term and the boundary conditions of the underlying problem.

3 Numerical Study

We discretize the Helmholtz equation (1) using $n = 64$ interior mesh points, so that the mesh size is $h = 1/(n+1)$, and we use $p = 4, 8, 16$ layers. For the case of four layers, we use the wave numbers

$$k = [20\ 20\ 20\ 20] + \alpha[0\ 20\ 10\ -10], \tag{10}$$

where α is a contrast parameter, and for larger p we just repeat this structure. The resolution we chose guarantees at least ten points per wavelength resolution for this experiment. We start with the case of a wave guide in the x direction, where we used Robin radiation conditions on the left and right, and homogeneous Dirichlet conditions on top and bottom. We show in Figure 1 the solution[3] we obtain for $\alpha = 1$ with a point source at $x = 2h$, $y = \frac{1-h}{2}$ for the case of four and sixteen layers in the top row, and below for the constant source $f = 1$.

We now test the approximate factorization (8) both as an iterative solver and as a preconditioner for GMRES for varying contrast parameter α and right hand sides. We do this both for $n = 64$ interior meshpoints and the contrast profile (10), and on

[2] it is the exact Schur complement, including all boundary information, the only approximation is the constant wave number.

[3] The boundary points are not plotted, so one can not see the homogeneous Dirichlet condition.

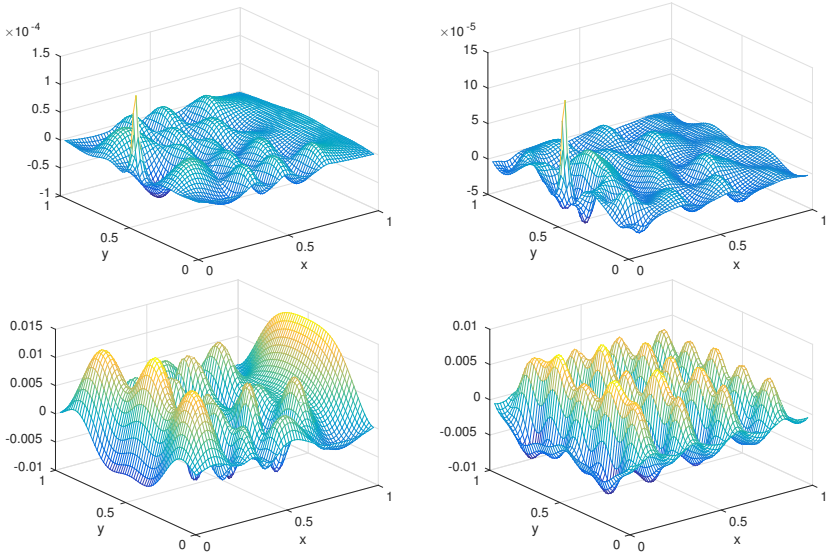

Fig. 1 Top: Solutions computed with a point source. Bottom: Solutions computed with $f = 1$. Left: 4 layers. Right: 16 layers.

a refined mesh with twice the number of interior meshpoints, $n = 128$, but also a profile with twice the size for the wave number, i.e.

$$k = [40\ 40\ 40\ 40] + \alpha[0\ 40\ 20\ -20], \tag{11}$$

so that we still have at least ten points per wavelength resolution. We show in Table 1 the number of iterations the methods took, where we stopped the iterative version of the algorithm at the relative error tolerance $1e-6$, and GMRES when the residual was reduced by $1e-6$, and we started with a zero initial guess. The three columns within each 'Iterative' or 'GMRES' column correspond to the point source f, constant $f = 1$ throughout the domain, and also a random f. The top part is for the smaller wave number experiment (10), and the bottom part is for the larger wave number experiment (11). We first see that for $\alpha = 0$, i.e. in the constant wave number case, the factorization is exact, both the iterative version and GMRES converge in one iteration step, and the contraction factor ρ (the spectral radius) of the iterative version equals numerically zero. As soon as we have however a non-constant wave number, already for $\alpha = 0.001$, the factorization is not exact any more. Nevertheless the methods still converge well, up to $\alpha = 0.01$ in the smaller wave number case in the top half of the table, i.e. a one percent variation in the wave number k. Here the contraction factor is $\rho = 0.2460$ for $p = 4$ subdomains, and grows when the number of subdomains p is increasing. For larger contrast, the iterative version of the algorithm can not be used any more, $\rho > 1$, and GMRES deteriorates now rapidly, for example if the contrast is at a factor of two, i.e. $\alpha = 1$, GMRES iteration numbers double when the number of subdomains doubles, the sweeping type methods are not

	$p=4$						$p=8$						$p=16$					
α	Iterative			ρ	GMRES		Iterative			ρ	GMRES		Iterative			ρ	GMRES	
0	1	1	1	5.8e-15	1	1	1	1	1	6.1e-15	1	1	1	1	1	4.6e-15	1	1	1
0.001	4	3	4	0.0250	3	3	4	5	4	5	0.0738	4	3	4	6	4	6	0.0979	5	4	5
0.005	6	4	6	0.1250	4	3	5	13	5	14	0.4031	7	5	7	12	8	11	0.3155	8	6	8
0.01	9	4	9	0.2460	5	4	5	32	7	34	0.6877	9	6	9	25	13	28	0.6244	11	7	11
0.05	-	7	-	1.6072	8	6	8	-	-	-	11.135	15	11	15	-	-	-	20.593	21	15	21
0.1	28	11	26	0.6887	9	7	9	-	-	-	3.0238	17	13	18	-	-	-	2.7604	25	17	26
1	-	-	-	2.4141	18	12	19	-	-	-	173.66	35	29	37	-	-	-	7.0979	62	44	67
0	1	1	1	5.9e-15	1	1	1	1	1	1	5.3e-15	1	1	1	1	1	1	6.5e-15	1	1	1
0.001	4	4	4	2.49e-2	4	3	4	5	5	5	0.1055	5	4	5	8	5	8	0.1991	6	5	6
0.005	7	7	7	0.1428	5	5	5	84	9	90	0.8824	9	6	9	26	11	27	0.6328	12	8	12
0.01	12	12	12	0.3300	6	5	6	-	18	-	1.9386	12	8	12	-	-	-	1.1614	19	11	19
0.05	-	-	-	4.5040	13	9	13	-	-	-	8.1397	23	17	22	-	-	-	1408.4	43	34	44
0.1	-	-	-	2.2412	14	11	15	-	-	-	20.614	20	14	19	-	-	-	2515.4	43	38	40
1	-	-	-	8.7091	31	20	33	-	-	-	6.9288	61	46	66	-	-	-	4.079e5	67	99	83

Table 1 Iteration numbers in the wave guide setting.

robust any more[4]. In the higher wave number case in the bottom part of the table, the methods start having problems already at $\alpha = 0.005$, variations of the wave speed of half a percent, and they deteriorate even more rapidly for higher contrast. We can also see comparing the last two lines of the top and bottom half of the table that doubling the wave number leads to twice the iteration numbers with GMRES as soon as the contrast is large enough, and GMRES failed to converge in less than hundred iterations at the bottom right. We also measured that in certain cases, the relative residual reduction of $1e-6$ for GMRES does not lead to a relative error of the same size. This is notably the case for $\alpha = 1$ in the smaller wave number case when $p = 8$ with point or random source (relative error $1.83e-4$ and $1.26e-4$ only), and in the larger wave number case when $p = 16$ with point or random source, (relative error 0.27317 and 0.52128 only !). So the corresponding GMRES iteration numbers (67 and 83) would need to be substantially higher to reach the same level of accuracy of $1e-6$ as for the other results in the table: we measured 129 instead of 67 to reach $1.8607e-6$ and 139 instead of 83 to reach $2.9641e-6$ respectively.

We next perform the same set of experiments, but now using Robin boundary conditions all around the domain, see Table 2. We see that the outer Robin boundary conditions are better than the wave guide setting for the sweeping type algorithms, they work now in the iterative version up to about a 10 percent variation of the wave number in this specific experiment. As soon as however there is a variation as large

[4] There are also two interesting apparent anomalies: in the smaller wavenumber case, for $p = 4$ and $\alpha = 0.05$ (and also one in the larger wave number case), the spectral radius is bigger than one, but for the source term $f = 1$ we observe convergence. We iterated in this case however further, and then the iterations also start to diverge, it is only that the divergent modes are not stimulated at the beginning by the source term $f = 1$ and zero initial guess, a typical phenomenon known from power iterations, which explains in the table the general observation that the problem with $f = 1$ is easier to solve than with the other sources, also for GMRES. For the same $p = 4$ and $\alpha = 0.1$, we then get surprisingly a spectral radius again smaller than 1, which is a lucky configuration and not observed for more subdomains or different α.

	$p=4$				$p=8$				$p=16$			
α	Iterative			ρ	GMRES			Iterative			ρ	GMRES		
0	1	1	1	3.6e-15	1	1	1	1	1	1	4.5e-15	1	1	1	1	1	1	2.8e-15	1	1	1
0.001	2	3	3	1.28e-3	2	2	2	3	3	3	3.40e-3	3	3	3	3	3	3	3.74e-3	3	3	3
0.005	3	3	3	6.58e-3	3	3	3	4	4	4	1.69e-2	4	3	3	4	4	4	1.92e-2	4	4	4
0.01	4	4	4	1.36e-2	3	3	3	4	4	4	3.35e-2	4	4	4	5	4	4	3.79e-2	4	4	4
0.05	6	6	6	8.25e-2	5	5	5	7	7	7	0.1446	6	6	6	10	9	9	0.2403	7	7	7
0.1	8	8	8	0.1677	6	5	6	9	9	9	0.2202	7	7	7	15	16	16	0.4182	9	9	10
1	80	80	80	0.8471	13	10	13	-	-	-	2.8446	24	19	25	-	-	-	3.1188	39	30	38
0	1	1	1	3.6e-15	1	1	1	1	1	1	4.0e-15	1	1	1	1	1	1	4.4e-15	1	1	1
0.001	3	3	3	1.91e-3	2	3	3	3	3	3	5.57e-3	3	3	3	4	4	4	1.29e-2	3	3	3
0.005	3	3	3	9.63e-3	3	3	3	4	4	4	2.73e-2	4	4	4	5	5	5	6.58e-2	5	5	5
0.01	4	4	4	1.97e-2	4	4	4	5	5	5	5.29e-2	5	5	5	7	7	7	0.1343	6	6	6
0.05	6	6	6	0.1006	5	5	5	11	11	11	0.2771	8	8	8	21	22	22	0.5287	10	10	11
0.1	10	9	9	0.2353	7	7	7	14	13	13	0.3796	9	9	9	41	44	43	0.7344	12	11	12
1	-	-	-	1.4684	19	14	19	-	-	-	2.9234	36	25	35	-	-	-	36.193	76	65	80

Table 2 Iteration numbers for a domain with Robin conditions all around.

		$p=4$	$p=8$	$p=16$	$p=4$		$p=8$	$p=16$
L/h	C	GMRES	GMRES	GMRES	Iterative	GMRES	GMRES	GMRES
5	4π	18 13 19	36 29 38	62 42 63	20 24 28	12 11 12	23 24 25	32 32 34
10	8π	18 13 19	36 29 38	61 41 62	18 24 25	11 11 12	20 20 22	28 25 29
5	4π	28 17 21	61 46 64	86 86 95	- - -	14 13 13	25 24 24	58 60 61
10	8π	28 17 19	61 46 64	87 81 90	39 47 46	11 11 12	22 21 22	46 45 48

Table 3 Iteration numbers in the presence of outer PMLs. Left: waveguide. Right: PMLs all around. Bottom part has doubled wavenumber and half the mesh size like in Tables 1 and 2.

as a factor of two, the method is not an effective solver any more, the iterative version diverges because $\rho > 1$, and GMRES iteration numbers deteriorate when the number of subdomains increases, like in the previous case: we still observe a doubling of the GMRES iteration count when the number of subdomains doubles, and also when the wave number is multiplied by 2. With Robin conditions all around, there is less loss of accuracy compared to the residual tolerance than in the wave guide case: only in the high wave number case for $\alpha = 1$ and $p = 16$, the relative error reached $1.6463e-05$ for the point source and $1.2333e-05$ for the random source instead of the $1e-6$ asked for in the relative residual, all other results had the required level also in the relative error.

Finally, we use a complex stretching PML instead of the outer Robin boundary condition. For example, we extend the right boundary from 1 to $1+L$ and perform in the extended region in (1) the transform $\partial_x \to s\partial_x$, $s = \frac{1}{1-\mathbf{i}C(x-1)^2/(L^3k(1,y))}$, $\mathbf{i} = \sqrt{-1}$, and similarly on the other boundaries. We increase L and C to get more absorption in the PMLs, and check how this affects the results for $\alpha = 1$ in Table 1 and Table 2, see Table 3. The iterative version diverges in most cases except when $p = 4$ for the PML-all-around problem. Absorption helps GMRES marginally for the waveguide problem but remarkably for the PML-all-around problem. Note that, however, the iteration count still doubles along with the number of subdomains and when dou-

bling the wave number for many subdomains. We also tested the case of a fixed wave number profile, namely the one in Table 3 at the bottom right with 16 layers: for $p = 4$ we obtain for GMRES the iteration numbers $\boxed{16\ 20\ 21}$, and for $p = 8$ $\boxed{59\ 69\ 70}$. This indicates that also for a fixed difficulty, i.e. fixed number of layers, iteration numbers grow when subdomain numbers are increasing. We observe however also when comparing with $p = 16$ at the bottom right of Table 3 the interesting phenomenon that once layers are all aligned with subdomains, the problem becomes apparently a bit easier. We are currently studying this phenomenon theoretically. Note that if too many PMLs are used, the 2-norm of the residuals may be dominated by the residuals in the PMLs, and one should use a more reliable metric for the stopping criterion.

4 Conclusion

We presented the simplest common form of the fundamental algorithm underlying the new type Helmholtz (and Maxwell) solvers based on sweeping. These solvers are among the best currently available solvers for such type of problems, and they can be made robust in the wave number by increasing the accuracy of the PML, provided the wave number is constant. If the wave number is not constant however, the PML is not the right approximation of the Dirichlet to Neumann operator or the Schur complement any more, which is the essential ingredient for these algorithms to be effective. We showed by a simple set of numerical experiments which is easy to reproduce that in a layered medium with contrast of only one percent, these algorithms already perform substantially less well if the layers are not aligned with the sweeping direction, and when the contrast is as large as a factor of two, the methods do not work any more as stationary iterations, and preconditioned GMRES iteration numbers start to grow drastically: they increase linearly in the number of subdomains and the wave number in our experiments. One must therefore investigate an approximation different from PML for the Dirichlet to Neumann operator in the case of non-constant wave numbers.

References

1. Ivo M. Babuska and Stefan A. Sauter. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? *SIAM Journal on numerical analysis*, 34(6):2392–2423, 1997.
2. Yassine Boubendir, Xavier Antoine, and Christophe Geuzaine. A quasi-optimal non-overlapping domain decomposition algorithm for the Helmholtz equation. *Journal of Computational Physics*, 231(2):262–280, 2012.
3. Zhiming Chen and Xueshuang Xiang. A source transfer domain decomposition method for Helmholtz equations in unbounded domain. *SIAM J. Numer. Anal.*, 51:2331–2356, 2013.

4. Zhiming Chen and Xueshuang Xiang. A source transfer domain decomposition method for Helmholtz equations in unbounded domain Part II: Extensions. *Numer. Math. Theor. Meth. Appl.*, 6:538–555, 2013.
5. Pierre-Henri Cocquet and Martin J. Gander. On the minimal shift in the shifted laplacian preconditioner for multigrid to work. In *Domain Decomposition Methods in Science and Engineering XXII*, pages 137–145. Springer, 2016.
6. Pierre-Henri Cocquet and Martin J. Gander. How large a shift is needed in the shifted Helmholtz preconditioner for its effective inversion by multigrid? *SIAM Journal on Scientific Computing*, 39(2):A438–A478, 2017.
7. Bjorn Engquist and Lexing Ying. Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation. *Comm. Pure Appl. Math.*, LXIV:0697–0735, 2011.
8. Bjorn Engquist and Lexing Ying. Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers. *Multiscale Model. Sim.*, 9:686–710, 2011.
9. Yogi A. Erlangga, Cornelis Vuik, and Cornelis Willebrordus Oosterlee. On a class of preconditioners for solving the helmholtz equation. *Applied Numerical Mathematics*, 50(3-4):409–425, 2004.
10. Oliver G. Ernst and Martin J. Gander. Why it is difficult to solve Helmholtz problems with classical iterative methods. In *Numerical analysis of multiscale problems*, pages 325–363. Springer, 2012.
11. Martin J. Gander. Optimized Schwarz methods. *SIAM J. Numer. Anal.*, 44:699–731, 2006.
12. Martin J. Gander, Ivan G. Graham, and Euan A. Spence. Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed? *Numer. Math.*, 131:567–614, 2015.
13. Martin J. Gander, Laurence Halpern, and Frederic Magoules. An optimized Schwarz method with two-sided Robin transmission conditions for the Helmholtz equation. *Int. J. Numer. Meth. Fluids*, 55:163–175, 2007.
14. Martin J. Gander, Frederic Magoules, and Frederic Nataf. Optimized Schwarz methods without overlap for the Helmholtz equation. *SIAM J. Sci. Comput.*, 24:38–60, 2002.
15. Martin J. Gander and Frederic Nataf. AILU: a preconditioner based on the analytic factorization of the elliptic operator. *Numer. Linear Algebra Appl.*, 7:505–526, 2000.
16. Martin J. Gander and Frederic Nataf. An incomplete LU preconditioner for problems in acoustics. *J. Comput. Acoust.*, 13:455–476, 2005.
17. Martin J. Gander and Achim Schädle. The pole condition: a Padé approximation of the Dirichlet to Neumann operator. In *Domain Decomposition Methods in Science and Engineering XIX*, pages 125–132. Springer, 2011.
18. Martin J. Gander and Achim Schädle. On the relationship between the pole condition, absorbing boundary conditions and perfectly matched layers. 2018. in preparation.
19. Martin J. Gander and Hui Zhang. Iterative solvers for the Helmholtz equation: factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods, 2016.
20. Ivan G. Graham, Euan A. Spence, and Eero Vainikko. Domain decomposition preconditioning for high-frequency Helmholtz problems using absorption. *ArXiv e-prints*, 2015.
21. Achim Schädle and Lin Zschiedrich. Additive Schwarz method for scattering problems using the PML method at interfaces. In *Domain Decomposition Methods in Science and Engineering XVI*, pages 205–212. Springer, 2007.
22. Andrea Toselli. Some results on overlapping Schwarz methods for the Helmholtz equation employing perfectly matched layers. In *Domain Decomposition Methods in Sciences and Engineering: Eleventh International Conference London, UK*, pages 539–545, 1998.
23. Leonardo Zepeda-Núñez and Laurent Demanet. The method of polarized traces for the 2D Helmholtz equation. *J. Comput. Phys.*, 308:347–388, 2016.
24. Leonardo Zepeda-Núñez, Russel J. Hewett, and Laurent Demanet. Preconditioning the 2D Helmholtz equation with polarized traces. In *SEG Technical Program Expanded Abstracts 2014*, pages 3465–3470. SEG, 2014.