

Overlapping Schwarz Methods for Helmholtz's Equation

M. A. Casarin¹, O. B. Widlund²

Introduction

In this paper, we report on our experience in solving the scalar indefinite Helmholtz equation iteratively. We propose a modification of the standard Overlapping Schwarz Methods, allowing discontinuities across the interfaces between the subdomains, and using a Sommerfeld-type, quasi-transparent boundary condition for the local solvers; this results in a very efficient iterative method. We concentrate on the numerical performance of the accelerated version of our method and study the dependence of the rate of convergence on several important parameters, e.g. the number of subdomains, the geometry of the subdomains, the wave number, the mesh size, the size of overlap, coefficient inhomogeneities and the shape of the obstacle. We can conclude that our method, not yet fully explained theoretically, is very fast and efficient. Previous work on this project was performed in collaboration with F. W. Elliott and X-C. Cai [CCFWEW97].

We consider a Helmholtz model problem given by

¹ IMECC-UNICAMP, Caixa Postal 6065, 13081 - 970 - Campinas - SP, Brazil. URL: <http://www.ime.unicamp.br/~casarin>. Electronic mail address: casarin@ime.unicamp.br. This work was supported in part by the National Science Foundation under Grants NSF-ECS-9527169 and NSF-CCR-9732208, in part by CNPq Brazil, and in part by FAPESP - Brazil, under project 97/12389-5.

² Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, N.Y. 10012. URL: <http://cs.nyu.edu/cs/faculty/widlund/index.html>. Electronic mail address: widlund@cs.nyu.edu. This work was supported in part by the National Science Foundation under Grants NSF-CCR-9732208 and NSF-ECS-9527169, and in part by the U.S. Department of Energy under Contract DE-FG02-92ER25127.

Eleventh International Conference on Domain Decomposition Methods

Editors Choi-Hong Lai, Petter E. Børstad, Mark Cross and Olof B. Widlund

©1999 DDM.org

$$\begin{cases} -\Delta u - (k(x))^2 u = f & \text{in } \Omega \\ \frac{\partial u}{\partial n} + ik u = g_S & \text{on } \partial\Omega_S \\ \frac{\partial u}{\partial n} = g_N & \text{on } \partial\Omega_N \\ u = h & \text{on } \partial\Omega_D \end{cases}, \quad (1)$$

where Ω is a bounded two or three-dimensional region. We assume that $\partial\Omega_S$ is non empty. This equation is then uniquely solvable, and we note that the boundary condition, said to be of Sommerfeld type, is essential in the proof of this fact.

We use Green's formula, and complex conjugation of the test functions, to convert (1) into variational form: Find $u \in V$ such that,

$$\begin{aligned} b(u, v) &= \int_{\Omega} (\nabla u \cdot \nabla \bar{v} - k^2 u \bar{v}) dx + ik \int_{\partial\Omega_S} u \bar{v} ds \\ &= \int_{\Omega} f \bar{v} dx + \int_{\partial\Omega_S} g_S \bar{v} ds + \int_{\partial\Omega_N} g_N \bar{v} ds = F(v) \quad \forall v \in V, \end{aligned}$$

where V is an appropriate subspace of $H^1(\Omega)$. A finite element discretization can now be defined straightforwardly by replacing V by a suitable conforming finite element space. We concentrate here on low order finite elements.

Our interest in the application of the direct (non-mixed) finite element method to this equation has been inspired by the work of Ihlenburg and Babuška [IB95, IB97]. They have considered the well-posedness of the original problem and different finite element discretizations and proven, for a model problem in one dimension, that the basic estimate

$$|u|_{H^1} \leq Ck|F|_{H^{-1}}$$

holds. In the finite element case, an assumption of $hk < 1$ is used. The constant C is independent of p , the degree of the finite elements. Ihlenburg has also conducted extensive numerical experiments which suggest that this bound holds for problems in two or three dimensions. Error bounds of the following form are also given for $p = 1$ and kh small enough:

$$|error|_{H^1} \leq C_1\theta + C_2k\theta^2 \text{ where } \theta = \text{best } H^1\text{-error.}$$

With oscillatory solutions typical, we can expect θ to be on the order of kh . The second term, which is due to the phase error, will dominate unless k^2h is on the order of 1. In our experiments we have used a mesh fine enough to guarantee that the phase error is of the same order as the other component of the error.

Overlapping Method – Continuous and Discontinuous

In the past few years, much work has been done on the iterative solution of the Helmholtz's equation that governs the time-harmonic scattering of sound waves by an object. Among many others, see, e.g., [FML98, Kim98] we mention, in particular, the thesis of Després [Des91]. We consider an overlapping version of his method (which was

briefly considered in his work in a particular case). Our basic multiplicative, one-level overlapping Schwarz method can be described as follows: let $\{\Omega_j\}$ be a set of open subregions that covers the given open region Ω . Each subregion Ω_j can have many disconnected components; it is often profitable to color the subregions of an original overlapping decomposition of Ω using different colors for any pair of subregions that intersect. The original set of subregions can then be partitioned into sets of subregions, one of each color, effectively reducing the number of subregions. This decreases the number of fractional steps of our Schwarz methods and helps make the algorithm parallel. The number of colors is denoted by J .

For $j = 1, \dots, J$, given a finite element function $u^{n+\frac{i-1}{J}}$, defined on Ω , we compute $u_{\Omega_j}^{n+\frac{i}{J}}$, the restriction of $u^{n+\frac{i}{J}}$ to Ω_j , by solving the variational form of the following local problem:

$$\left\{ \begin{array}{ll} -\Delta u_{\Omega_j}^{n+\frac{i}{J}} - (k(x))^2 u_{\Omega_j}^{n+\frac{i}{J}} & = f & \text{in } \Omega \\ \frac{\partial u_{\Omega_j}^{n+\frac{i}{J}}}{\partial n_{\Omega_j}} + ik u_{\Omega_j}^{n+\frac{i}{J}} & = -\frac{\partial u_{\Omega_j^c}^{n+\frac{i-1}{J}}}{\partial n_{\Omega_j^c}} + ik u_{\Omega_j^c}^{n+\frac{i-1}{J}} & \text{on } \partial\Omega_j \cap \Omega \\ \frac{\partial u_{\Omega_j}^{n+\frac{i}{J}}}{\partial n} + ik u_{\Omega_j}^{n+\frac{i}{J}} & = g_S & \text{on } \partial\Omega_j \cap \partial\Omega_S \\ \frac{\partial u_{\Omega_j}^{n+\frac{i}{J}}}{\partial n} & = g_N & \text{on } \partial\Omega_j \cap \partial\Omega_N \\ u_{\Omega_j}^{n+\frac{i}{J}} & = h & \text{on } \partial\Omega_j \cap \partial\Omega_D \end{array} \right. , \quad (2)$$

Here, Ω_j^c refers to the complement of Ω_j , and n is the outward normal. The boundary of the original region Ω has been partitioned into three parts, two of which are possibly empty: $\partial\Omega = \partial\Omega_S \cup \partial\Omega_N \cup \partial\Omega_D$. The Sommerfeld condition is imposed on $\partial\Omega_j \cap \Omega$ to guarantee solvability and stability of this problem; it also has a dramatic impact on the convergence rate, if compared with the use of standard Dirichlet conditions for this problem, even in the case when the latter is enough to guarantee stability of the problems defined on the Ω_j ; see [CCFWEW97].

An important detail, somewhat hidden in the above formulation of the local problems is that at step j , represented by (2), the values computed for u^{n+1} , at the mesh points on $\partial\Omega_j$, might differ from those of u^n at the same points. This forces us to consider two alternatives:

1. the new values replace the old ones, and the resulting finite element function is continuous, with possibly large gradients just outside $\partial\Omega_j$;
2. the new *and* the old values are kept.

The second alternative produces an algorithm which proceeds to the converged, continuous solution through a sequence of discontinuous iterates. It is hereafter called the Discontinuous Overlapping Schwarz Method (OSM-D). In [CCFWEW97], we established, through a variety of numerical experiments, that the discontinuous version is actually superior in most cases considered; many details about the implementation of the method are also given in that paper in which the discontinuous method is

referred to as ALG3. We only mention here that in the finite element context the computation of normal derivatives of the form $\frac{\partial u_{\Omega_j}^n}{\partial n_{\Omega_j^c}}$ is made with the consistent use of Green's formula, and that the discontinuities force us to perform a subassembly involving elements outside Ω_j , which are next to $\partial\Omega_j$. Similarly, the computation of terms originating from the term $iku_{\Omega_j}^n$ involve a subassembly of the mass matrix corresponding to different parts of $\partial\Omega_j$. It is easy to see that in the absence of crosspoints interior to Ω , and formed by the boundaries of the internal overlapping subdomains, the continuous and discontinuous methods produce the same iterates.

In this paper, we report on an experimental study of the discontinuous method which is more detailed than the previous ones, and show that the method is indeed very fast and efficient, even in the absence of a coarse space, which is very important for the efficiency of some competing alternatives; see [FML98]. We consider several different geometries, unstructured meshes and different partitions of the domain into subdomains. The experiments were all conducted with MATLAB, and the geometry, mesh generation, and matrix assemblies were done using the PDE toolbox.

The theoretical results that we have obtained have already been outlined in [CCFWEW97]. They are satisfactory only for the continuous version of the method.

Convergence Study

We now describe the geometry and discretization used in our numerical experiments. We work with three different geometries and boundary conditions: a) a Sommerfeld square, which is a unit square with homogeneous Sommerfeld conditions on the four sides. The right hand side is chosen so that the solution is the combination of two plane waves incoming at angles 36° and 178° ; b) a waveguide problem, which is again a unit square, with homogeneous Sommerfeld conditions on the horizontal sides, homogeneous Neumann conditions on the right vertical side, and a constant Dirichlet condition identically equal to 1 on the left vertical side. The right hand side f is zero; and c) the scattering geometry, which corresponds to a cavity of diameter approximately 1 inside a circle of radius approximately 3.7; the right hand side is again zero, and Dirichlet data is imposed on the surface of the scatterer, corresponding to an incoming plane wave of wavenumber k (given in the following) and angle 36° . For this last case, the geometry of Ω and the 14 non-overlapping subdomains Ω_j is given in Figure 1 under the name GSC14-1. The scatterer is a cavity with vertices at: $(-0.52, -0.52)$, $(-0.52, 0.52)$, $(0.5, 0.5)$, $(-0.5, 0.5)$, $(-0.5, -0.5)$, $(0.5, -0.5)$. The overlap is produced by a layer of δ elements added in all directions which are interior to Ω . In all experiments the mesh is unstructured, and except for the very small test cases, the number of mesh points is certainly enough to resolve the oscillations in the solution.

Apart from the geometry, the relevant parameters for describing an experiment are:

- the total number of grid points N . We define $1/h \simeq \sqrt{N}$ when Ω is a unit square;
- k the spatial frequency, which can depend on $x \in \Omega$;
- $nsub$ the number of subregions;

- δ the number of elements across half of the overlap;

Our main goal is to study the dependence of the number of iterations of our method on the parameters listed above, as well as to evaluate how this iteration count will depend on the original boundary conditions, the variation of the coefficients, the geometrical placement of the subdomains, and the global geometry. The linear, sequential method of (2) is accelerated by the GMRES method with restart equal to 30, and the algorithm is stopped as soon as the norm of the residual is reduced by 5 orders of magnitude (10^{-5}).

When possible, we have compared our method with two competing alternatives: a) a FETI type algorithm as reported by Farhat, Macedo, and Lesoinne [FML98]; and b) the ADOP method proposed by Kim [Kim98]. This is possible because our implementation is extremely flexible with respect to the geometry and all other parameters; this comparison is also made in order to provide a fair evaluation of our scheme. We note that the comparisons are made using exactly the same problems as reported in the references given above.

Dependence on the mesh size

Our first set of experiments involves a study of the optimality with respect to the mesh size. We compare our method first with the FETI method, applied to a waveguide problem [FML98]Table 1, and then with the ADOP method, for a Sommerfeld square [Kim98]Table 6. We can see, from Tables 1 and 2, that the OSM-D method is very competitive, and that when compared with the ADOP method it has the added advantage of requiring no parameter estimation. The data shows that the iteration growth is mild with respect to the number of points. This observation is enforced by the data in the other tables; see, e.g. Tables 10 and 11. All our experiments used unstructured meshes, and that therefore we cannot exactly reproduce the discrete problem being solved by an alternative algorithm on a structured mesh.

Table 1 Waveguide problem with $k = 20$, variable mesh size: comparison between the FETI method with 5×5 subdomains, and OSM-D with 24 subdomains and $\delta = 6$.

FETI		OSM-D	
$1/h$	itcount	$1/h$	itcount
100	71	83	26
200	70	163	32
250	73		

Dependence on the wavenumber k

We also study the behavior of the method as the wavenumber k varies. The waveguide problem is solved with both the FETI and the OSM-D methods, and the data is

Table 2 Sommerfeld square problem, $k = 25$, and 16×4 subdomains: comparison between the ADOP method and OSM-D with $\delta = 5$

ADOP		OSM-D	
$1/h$	itcount	$1/h$	itcount
128	60	114	37
256	68	173	43

reported in Table 3.

Table 3 Waveguide Problem: comparison between the FETI method with $1/h = 384$, no coarse space, $nsub = 5 \times 5$, and the OSM-D method with 27649 points, 24 subdomains, and $\delta = 6$

FETI	Waveguide	OSM-D	Waveguide
k	itcount	k	itcount
20	70	20	58
40	136	40	64
60	164	60	70

We have also solved the scattering problem; the data is in Table 4, and the geometry (GSC14-1) is given in Figure 1.

Table 4 OSM-D method applied to the Scattering problem. The domain is partitioned into 14 subdomains (GSC14-1), $\delta = 6$, and there are a total of 26596 points

OSM-D	
k	itcount
12	11
24	15
36	19 ($\delta = 9$)

We can see from Tables 3 and 4 that the increase of the wavenumber does lead to a growth in the iteration count, although this effect is not as pronounced as for the FETI method without a coarse space. It is also clear that even in the scattering geometry the dependence on the wave number is quite weak. This is in contrast with a plausible intuitive interpretation of the method as producing a system of waves propagating across the subdomains; the larger the wavenumber the more complex would be the

reflection and diffraction phenomena, and one would therefore expect the performance of the algorithm to deteriorate quite rapidly.

In actual applications, the coefficients of the time-harmonic wave equation often vary, and a natural question is therefore whether the performance of the method deteriorates if the coefficients are inhomogeneous. We borrow the following coefficient functions from Kim [Kim98] equation (69):

$$\begin{aligned} k_1 &= 12 / ((2 + \sin \pi x) (2 - \sin 3\pi y)) \\ k_2 &= 12 / ((1.1 + \sin \pi x) 1 (1.1 - \sin \pi y)) . \end{aligned}$$

The Discontinuous Overlapping Method with $\delta = 4$ is then applied to the scattering geometry partitioned into 14 subdomains, and the results appear in Table 5, where we have also included a set of results for the case of constant k . It can be seen from this Table that the OSM-D appears to be very robust with respect to coefficient inhomogeneities in the wavenumber function; the actual iteration count is also very modest. Inhomogeneities in the other coefficients have also been studied, but since the behavior is very similar, we omit that particular data.

Table 5 Application of the OSM-D method to a Scattering problem with variable wavenumber. The domain is partitioned into 14 subdomains (GSC14-1), and $\delta = 4$.

	$k = 12$	k_1	k_2
N	itcount	itcount	itcount
1735	7	7	10
6746	7	10	14
26596	11	13	20

Dependence on the number of subdomains

The OSM-D method proposed here relies only on local solves. The performance of similar methods, in the positive definite case, deteriorates with an increasing number of subdomains. The same behavior occurs here; our goal is now to explore this dependence quantitatively, and to compare that with the FETI method without a coarse space. The results for the Waveguide problem are given in Table 6. The value of k is chosen so that we have approximately 3 wavelengths across the domain, and the number of mesh points is of the same order in both cases. The iteration count clearly grows with the number of subdomains but less so for the OSM-D method. At this time, we do not have a good explanation for this fact.

We also provide the iteration count for a scattering problem, with the domain partitioned into 5, 14, and 98 subdomains, respectively, all with good aspect ratios. The value of k is set to 20, and $\delta = 4$. The same trend appears here, although the actual number of iterations is low. The development and implementation of computationally

Table 6 Waveguide problem, $k = 20$, 3 wavelengths across the domain. Comparison between the FETI method with approximately 40,000 points in the mesh, and $\delta = 6$.

FETI		OSM-D	
$nsub$	itcount	$nsub$	itcount
2×2	15	2×2	12
4×4	46	4×4	26
5×5	70	24	32

efficient coarse spaces is currently being considered, in order to treat cases with possibly many more subdomains.

Table 7 OSM-D, Scattering $k = 20$, $\delta = 4$

N	$nsub$	itcount
27649	5	8
26596	14	14
15392	98	27

Dependence on the original boundary conditions

We next compare two sets of runs of the OSM-D on the same square region. The main difference from one run to the other is the boundary condition: the first is a Waveguide square, partitioned into 24 subdomains, and the second is a Sommerfeld square, partitioned into 5×5 subdomains. Again, the subdomain partition is not regular, and was made by hand, in order to produce non-square subregions with relatively good aspect ratios. We use $k = 20$ and $\delta = 4$ in both cases and the results are reported in Table 8. It is immediately apparent that the boundary condition imposed on $\partial\Omega$ has a considerable impact on the iteration count. This is surprising if compared with similar experiments for the standard OSM method in the positive definite case, with Dirichlet local problems. It appears that the waveguide boundary condition excites modes that are not damped as efficiently by the algorithm, and therefore the iteration count is larger for this case.

Dependence on the subdomain geometry

As previously mentioned, our method could be viewed intuitively in terms of waves propagating across the subdomains: the Sommerfeld-like condition for the local problems could be considered a better choice than the standard Dirichlet condition in this respect, and that would be a possible explanation for its much better performance;

Table 8 Study of the dependence on the original boundary conditions for the OSM-D method. Comparison between the Waveguide square partitioned into 24 subdomains, and the Sommerfeld square partitioned into 25 subdomains. In both cases, $k = 20$ and $\delta = 4$.

Waveguide		Sommerfeld	
$1/h$	itcount	$1/h$	
83	26	121	9
163	32	161	13

see [CCFWEW97]. We note that such heuristics is at the core of some methods for the Helmholtz's equation which look for more transparent boundary conditions for the local problems; see [Gha96]. Inspired by this heuristics, we explore, in this section, how the algorithm depends on the particular choice and configuration of the subdomains.

Our first test case involves the scattering geometry GSC14-1. We use the OSM-D method, accelerated with GMRES, for various values of k , δ , and N (the total number of mesh points). The results are reported in Table 9. The geometry of this and the next cases are represented in Figure 1.

Table 9 OSM-D, Scattering problem, geometry GSC14-1, full opening, 14 subdomains, and $\delta = 6$, unless otherwise noted.

N	$k = 12$	$k = 24$	$k = 36, \delta = 9$	$k = 12, \delta = 4$
4438	46	78	50	>60
17404	72	>200	> 100	> 60

The tips at the entrance of the cavity are possibly points where our method would have difficulty converging to the exact solution. In order to check that, we change the subdomain partition from the geometry GSC14-1 to GSC14-2 (see Figure 1). The number of subdomains and the other parameters are kept the same, and the results are reported in Table 10. The iteration count decreases in all cases considered, at times by a factor of 6, showing that the configuration of the subdomains (and not only their number and aspect ratio) is indeed of a considerable importance for the iterative solution of this problem.

Continuing the exploration of the heuristics mentioned above, we now question whether the aperture of the cavity has a deleterious impact on iteration count. We keep roughly the same subdomain configuration as in GSC14-2, except that we change the cavity into one with a tighter aperture, respectively $1/2$ and $1/3$ of the opening, for Tables 11 and 12, corresponding to the geometries GSC14-3 and GSC14-4. The iteration count increases from GSC14-3 to GSC14-4, but the changes are not very dramatic.

Table 10 OSM-D, Scattering problem, geometry GSC14-2, full opening, 14 subdomains, $\delta = 6$ unless otherwise noted. The subdomain configuration is slightly different from that of the previous table.

N	$k = 12$	$k = 24$	$k = 36, \delta = 9$	$k = 12, \delta = 4$
6746	7	14	29	11
26596	11	15	19	13
105608				13

Table 11 OSM-D, Scattering problem, geometry GSC14-3, 1/2 the opening, 14 subdomains, $\delta = 6$ unless otherwise noted. The subdomain configuration is slightly different from that of the previous table.

N	$k = 12$	$k = 24$	$k = 36, \delta = 9$	$k = 12, \delta = 4$
7389	11	58	39	15
28966	10	20	44	13

Table 12 OSM-D, Scattering problem, geometry GSC14-4, 14 subdomains, $\delta = 6$ unless otherwise noted. The subdomain configuration is slightly different from that of the previous table.

N	$k = 12$	$k = 24$	$k = 36, \delta = 9$	$k = 12, \delta = 4$
9552	13	89	54	24
37576	15	33	74	17

In these last four tables, one can also observe the effect of increasing the total number of points, and/or increasing the wavenumber k . The behavior follows the trends identified in our previous tables. In Table 11, for example, we can also observe that increasing the size of the overlap δ does improve the method, both for $k = 12$ and $k = 36$. So far, we have not been able to explain the extraneous behavior showed by the method for $k = 24$, namely that increasing the number of points improves the method, and also that the iteration count is relatively large.

Conclusion

We have developed and tested the OSM-D method in a variety of situations, and showed that it compares favorably with some alternative algorithms. While efforts to

develop coarse spaces appropriate to our method are under way, we have here used only a method without a global component even for cases with relatively many subdomains. The iteration count does deteriorate with the number of subdomains, but the actual numbers are small. The algorithm is robust in regards to the wavenumber; this is a salient feature of our method. An interesting new issue considered is the effect of the configuration of the subdomains on the performance of the scheme; this has been studied in subsection 19, where we have established a basic trend; a more refined study is also in progress. At last we want to point out that for a non-symmetric, indefinite problem like the Helmholtz's equation considered here, the boundary conditions on $\partial\Omega$ can have a significant impact on the iterative method; see subsection 19.

REFERENCES

- [CCFWEW97] Cai X.-C., Casarin M. A., Frank W. Elliott J., and Widlund O. B. (1997) Overlapping Schwarz methods for solving Helmholtz's equation. In Mandel J., Farhat C., and Cai X.-C. (eds) *Domain Decomposition Methods 10*, volume 218 of *Contemporary Mathematics*, pages 391–399. AMS.
- [Des91] Després B. (October 1991) *Méthodes de Décomposition de Domaine pour les Problèmes de Propagation d'Ondes en Régime Harmonique*. PhD thesis, Paris IX Dauphine.
- [FML98] Farhat C., Macedo A., and Lesoinne M. (March 1998) A two-level domain decomposition method for the iterative solution of high frequency Helmholtz problems. Technical Report CU-CAS-98-06, Center for Aerospace Structures, University of Colorado at Boulder. Submitted to *Numerische Mathematik*.
- [Gha96] Ghanemi S. (January 1996) *Méthode de Décomposition de Domaine avec Conditions de Transmissions Non Locales pour des Problèmes de Propagation d'Ondes*. PhD thesis, Paris IX Dauphine.
- [IB95] Ihlenburg F. and Babuška I. (1995) Finite element solution of the Helmholtz equation with high wave number, Part I: The h-version of the FEM. *Computers Math. Applic.* 30(9): 9–37.
- [IB97] Ihlenburg F. and Babuška I. (1997) Finite element solution of the Helmholtz equation with high wave number, Part II: The h-p-version of the FEM. *SIAM J. Numer. Anal.* 34(1): 315–358.
- [Kim98] Kim S. (1998) Domain decomposition iterative procedures for solving scalar waves in the frequency domain. *Numer. Math.* 79: 231–259.

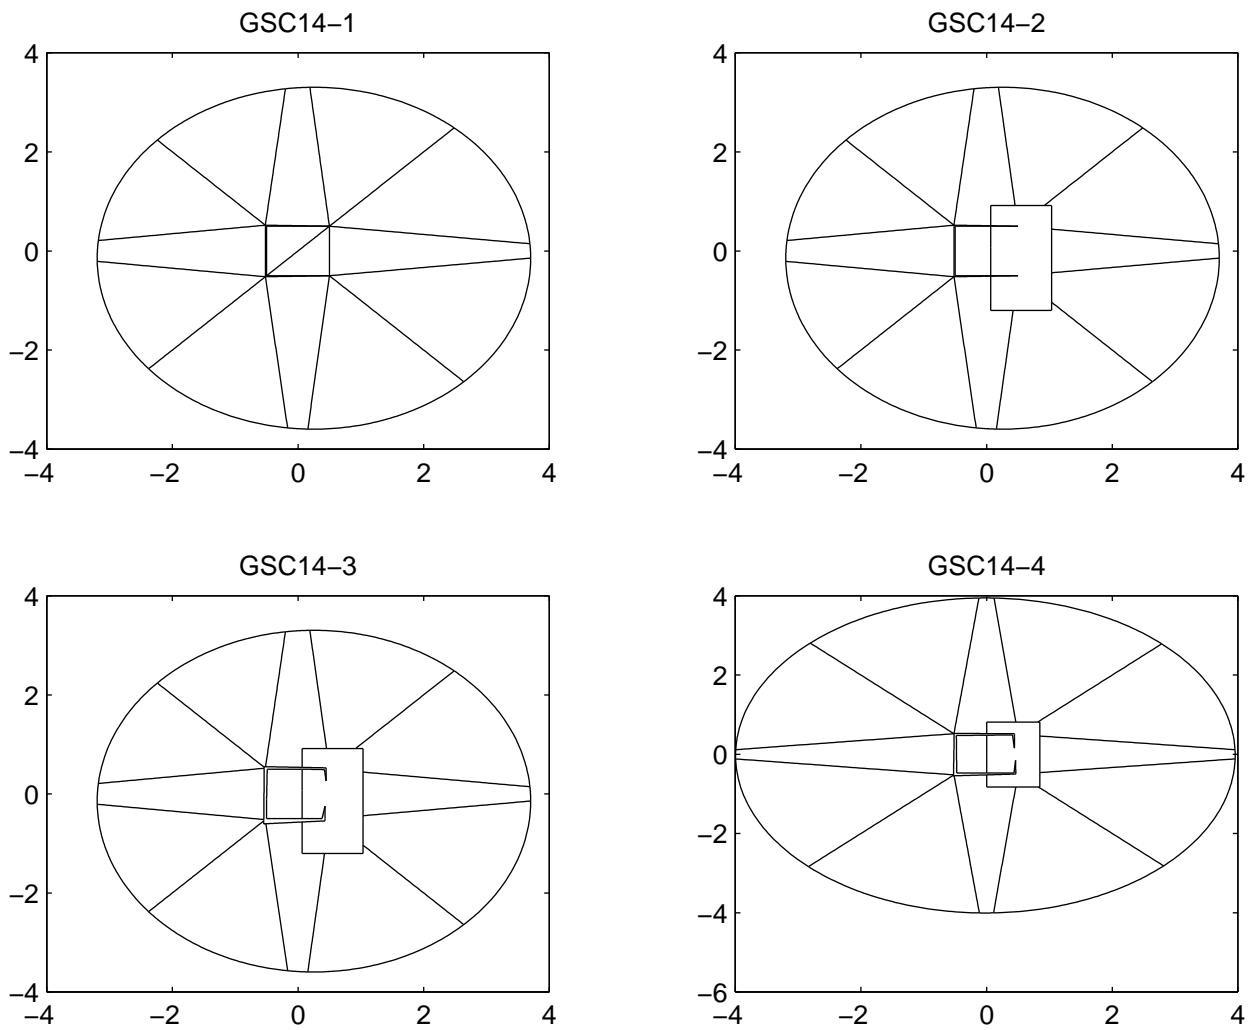


Figure 1 Scattering geometries, with several configurations of the subdomains