Coarse Corrections for Schwarz methods for Symmetric and Non-symmetric Problems

Martin J. Gander and Serge Van Criekingen

1 Introduction

As is well known, domain decomposition methods applied to elliptic problems require in most cases a coarse correction to be scalable (for exceptions, see [5, 6]), the choice of the coarse space being critical to achieve good performance. We present here four new coarse spaces for the Restricted Additive Schwarz (RAS) method of Cai and Sarkis [4], both for symmetric and non-symmetric problems, and implement them in the PETSc library [1, 2, 3]. We compare them to a coarse space named `Q1` here from [10], originating from [7] and [9], and more classical coarse spaces. In particular, we introduce the new *adapted* coarse spaces `Q1_adapt` and `Q1_inner_adapt` using basis functions that locally solve the problem considered also with advection and turn out to be more robust for strong advection. We also introduce the `Half_Q1` coarse space that halves the coarse space dimension compared to `Q1` by using a selected combination of its basis functions and turns out to be the fastest, and the new enriched coarse space `Enriched_Q1` which leads to the lowest iteration counts. We further present results of the optimized method ORAS obtained by introducing optimized transmission conditions at subdomain interfaces [8, 16, 7].

Throughout the paper, our model problem for the symmetric case is the Laplace problem, while for the non-symmetric case we consider

$$-\Delta u + \mathbf{a} \cdot \nabla u = 0 \tag{1}$$

with an upwind scheme on the unit interval (in 1-D) or unit square (in 2-D) using the 5-point finite difference discretization and homogeneous boundary conditions.

Martin J. Gander
University of Geneva, e-mail: martin.gander@unige.ch

Serge Van Criekingen
CNRS/IDRIS and Maison de la Simulation, e-mail: serge.van.criekingen@idris.fr

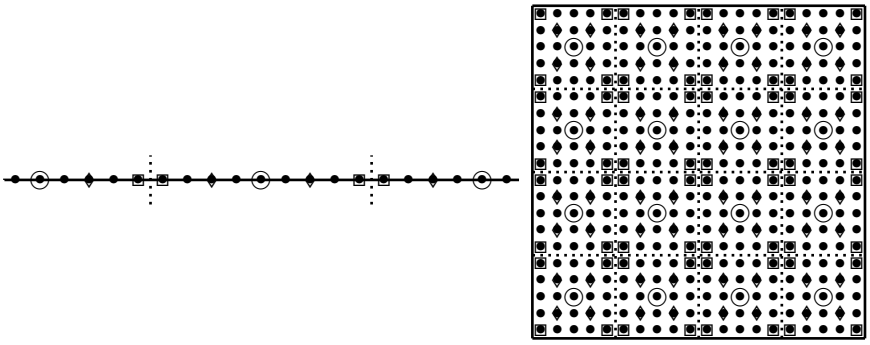

Fig. 1: Coarse grid point choice in 1-D and 2-D for the `Q1` (squares), `Q1_fair` (diamonds) and `Middle` (circles) options.

2 Two-level RAS with classical and new coarse spaces

We consider the solution of $A\mathbf{x} = \mathbf{b}$ on a domain Ω decomposed into a set of possibly overlapping subdomains Ω_j and introduce a restriction operator R_j onto each Ω_j. We also introduce a partition of Ω into non-overlapping subdomains $\tilde{\Omega}_j$ as well as the corresponding restriction operators $\tilde{R}_j$ for RAS. Obtaining a two-level method through coarse correction requires a restriction operator R_c to a coarse space, such that the resulting coarse system matrix reads $A_c = R_c\, A\, R_c^T$. The two-level coarse corrected RAS method with multiplicative coarse correction (denoted RAS2 in what follows) can then be written as

$$\mathbf{x}^{n+1/2} = \mathbf{x}^n + \sum_{j=1}^{J} \tilde{R}_j^T\, A_j^{-1}\, R_j\, (\mathbf{b} - A\mathbf{x}^n), \tag{2}$$

$$\mathbf{x}^{n+1} = \mathbf{x}^{n+1/2} + R_c^T\, A_c^{-1}\, R_c\, (\mathbf{b} - A\mathbf{x}^{n+1/2}), \tag{3}$$

where the first half iteration is the RAS method as defined by Cai and Sarkis [4].

The definition of the coarse space is critical to obtain an efficient two-level method. We consider here the following classical and new coarse spaces:

"`MidBasic`": The classical `MidBasic` coarse space, also called Nicolaides coarse space, defined by using a constant coarse basis function in each subdomain.

"`Middle`": The classical `Middle` coarse space taking the fine mesh points in the middle of each subdomain as coarse grid points, along with linear (bilinear in 2-D) basis functions centered on these points. This is illustrated in Fig. 1 in 1- and 2-D.

"`Q1`": The `Q1` coarse space [9, 7] based on linear basis functions with coarse grid points chosen as illustrated in Fig. 1, namely placed on each side of the subdomain interfaces (in 1-D) or around each cross point (in 2-D) of the non-overlapping decomposition. It was shown in [9] that, for the Laplace equation, the `Q1` coarse

correction yields convergence in two iterations in 1-D (or at iteration 1 in PETSc, where iteration count starts at 0).

"Q1_fair": This coarse space uses linear basis functions and the same number of coarse mesh points as Q1, but equally distributed as illustrated in Fig. 1. It is introduced for a fair comparison with Q1 in terms of coarse space dimensions.

"Q1_adapt": The new Q1_adapt coarse space using the same coarse points as Q1, but *computed* ("adapted") basis functions that solve the homogeneous equation considered in each subdomain. In the Laplace case, the Q1_adapt basis functions are thus the same as the Q1 (i.e., linear) functions, while with advection, the basis functions are different. In 1-D, Q1_adapt gives convergence of the two-level method at iteration 1 in PETSc, even in the non-symmetric case when advection is present, like Q1 for the Laplace problem in [9].

In 2-D, the Q1_adapt basis functions are computed in two steps, first on the edges with a 1-D stencil obtained by lumping (i.e., summing up) the system matrix coefficients in the perpendicular direction, then inside each subdomain using the computed edge functions as boundary conditions, a bit like in MsFEM.

"Q1_inner_adapt": Defined in 2-D only, this new coarse space differs from Q1_adapt in that the coarse basis functions are "adapted" only inside each subdomain: the first of the two steps in Q1_adapt is skipped, and linear edge functions are used as boundary conditions to compute the basis functions within each subdomain.

"Half_Q1": The new Half_Q1 coarse space is motivated by the eigenmodes of the RAS iteration matrix corresponding to its eigenvalues closest to 1 in modulus. In Fig. 2a, we computed them with SLEPc [13] (https://slepc.upv.es), for the Laplace test case and a 2×2 subdomain decomposition using minimal overlap (no algebraic overlap, i.e., block Jacobi). If q_1, q_2, q_3, q_4 are the Q1 basis functions at a cross point, it can be observed that these modes appear to be $q_1 + q_2 + q_3 + q_4$ and $q_1 - q_2 + q_3 - q_4$, respectively. The Half_Q1 coarse space is therefore obtained by taking these 2 combinations as basis functions, thus with 2 basis functions per cross point instead of 4 in the Q1 case. (To add to Fig. 2a which gives only the first two eigenvalues, note that the next eigenvalues are .975571 (double), −.975571 (double), −.969651 and .969651).

With minimal overlap, we observed that the property of having the two largest eigenmodes in modulus corresponding to one continuous and one discontinuous mode remains verified when increasing the number of subdomains. We also observed this property when introducing various types of advection. This is illustrated in Fig. 2b for the case with 25 subdomains on our model problem (1) with rotating fluid advection $a_x = -10y$ and $a_y = 10x$ (- for this case, the next eigenvalues are complex: $-0.989110 \pm 0.001513i$, $0.989110 \pm 0.001513i$ and $0.983268 \pm 0.002304i$).

With more than minimal overlap (i.e., non-zero algebraic overlap), even if we observed exceptions (typically when using more than 4 subdomains and a relatively low fine mesh resolution), the largest two modes tend to remain one continuous and one discontinuous one, but the corresponding eigenvalues are then different in modulus, with a difference that increases when increasing the overlap. We illustrate

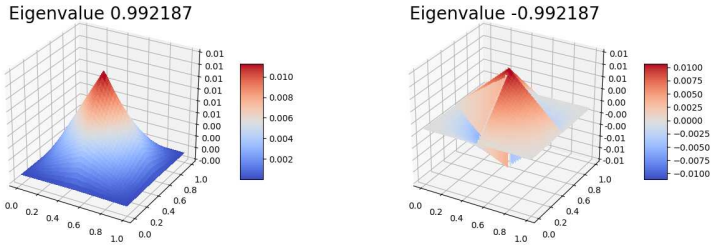

(a) 2×2 subdomains, no advection, minimal overlap.

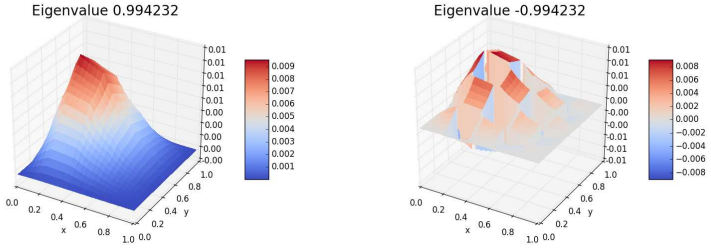

(b) 5×5 subdomains, advection $a_x = -10y$ and $a_y = 10x$, minimal overlap.

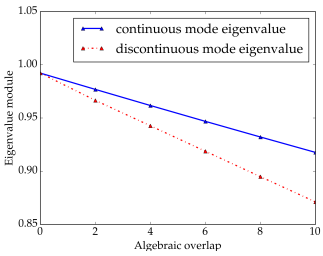

(c) 2×2 subdomains, no advection: evolution of the two largest eigenvalues in modulus with the overlap.

Fig. 2: In (a) and (b), eigenmodes of the RAS iteration operator corresponding to the two largest eigenvalues in modulus using a 256×256 fine mesh resolution; continuous modes on the left and discontinuous modes on the right. In (c), evolution of the two largest eigenvalues in modulus.

this for the 2×2 subdomain decomposition by displaying in Fig. 2c the evolution of the two largest eigenvalues in modulus when increasing the overlap.

"`Enriched_Q1`": This new coarse space is obtained by adding extra linear basis functions to the `Q1` coarse space, namely (in 2-D) with one extra coarse point placed in the middle of each edge and corresponding extra linear basis function. The goal is to come a step closer to the 2D grid representing a complete coarse space, leading

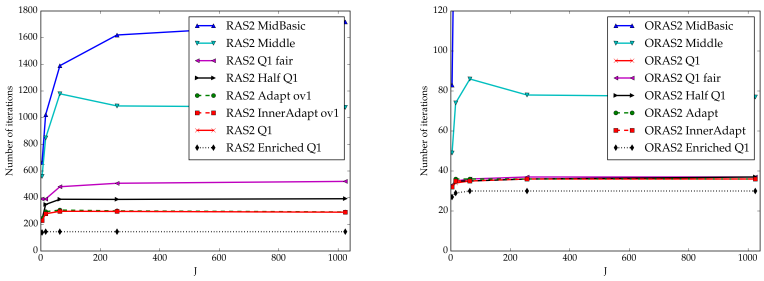

(a) RAS2 **(b)** ORAS2

Fig. 3: Results with advection $a_x = -10y$ and $a_y = 10x$.

to convergence in two iterations [9, Fig. 8]. This coarse space is thus twice as big as as `Q1`.

3 Numerical Results

Fig. 3 shows the iteration count for a weak scalability analysis on our non-symmetric model problem (1) with rotating fluid advection $a_x = -10y$ and $a_y = 10x$. This analysis consists in increasing the size of the problem while maintaining constant the workload per subdomain. The subdomain decomposition ranges from 2×2 to 32×32, each subdomain having a 256×256 fine mesh and being handled by one CPU core. The number of cores J ranges thus from 4 to 1024 here, and the coarse space dimension is J for `MidBasic` and `Middle`, $4J$ for `Q1`, `Q1_fair`, `Q1_adapt` and `Q1_inner_adapt`, $2J$ for `Half_Q1` and $8J$ for `Enriched_Q1` . An algebraic overlap of 2 is considered, which means one extra mesh layer for both subdomains at an interface and corresponds to an overlap of 1 in the PETSc sense. The corresponding Laplace results are very similar, we thus only show them in Table 1 for comparison.

While Fig. 3a displays the result for the RAS2 method, Fig. 3b displays the results for the optimized ORAS2 method obtained by modifying the local A_j matrices in the RAS2 iterations (2)-(3) to express Robin interface conditions [16], with a first-order accurate discretization of the normal derivative and two-level optimized coefficients (determined for the symmetric case) as defined in [7].

We observe that, except for the larger `Enriched_Q1` coarse space, the `Q1` coarse space gives the lowest iteration count when used with (non-optimized) RAS. Using adapted basis functions (i.e., `Q1_adapt` or `Q1_inner_adapt`) does not reduce the iteration count in the present case. However, these adapted coarse spaces appear more robust than `Q1` when increasing the advection strength, as can be seen in Table 1 with a five times larger advection: some of the stationary iterations appear to

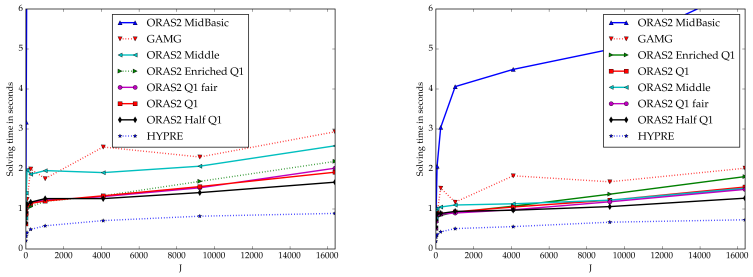

(a) Without GMRES acceleration **(b)** With GMRES acceleration

Fig. 4: Computation times (s.) for the weak scaling experiment for the non-symmetric model problem with $a_x = -10y$ and $a_y = 10x$.

diverge using the `Q1` and/or `Half_Q1` coarse spaces with a rotating fluid advection of magnitude 50, while this is not the case with magnitude 10 (Fig. 3).

We also observe from Fig. 3 that `Q1_fair` and `Half_Q1` take more advantage of the application of the optimized ORAS method than `Q1`, since their iteration counts then become all quasi-identical.

Timing results are presented in Fig. 4 for our weak scalability analysis, this time using up to $128 \times 128 = 16,384$ CPU cores (one per subdomain) of the CPU partition of the `Jean Zay` supercomputer at the Institute for Development and Resources in Intensive Scientific Computing (CNRS/IDRIS). A relative tolerance of 1.e-8 is used as convergence criteria. Note that PETSc's native direct solver is used for the local serial subdomain solves, while the coarse solve is performed in parallel with the MUMPS direct solver, after agglomeration of the coarse unknowns on a subset of the processors (here maximum 64) using PETSc's "Telescope" tool [14]. Beside the results obtained with the various coarse corrections introduced above, timings obtained with two algebraic multigrid options available through PETSc are also presented, namely `HYPRE/BoomerAMG` [12] (with tuning form [17]) and PETSc's native algebraic multigrid preconditioner `GAMG` (with smoothed aggregation and CG eigenvalue estimator [2]).

J	4	16	64	256	1024	4	16	64	256	1024
	RAS2					ORAS2				
`Q1_fair`	179(420)	357(424)	481(479)	506(509)	521(522)	80(34)	39(35)	37(36)	37(37)	37(37)
`Q1`	Div(255)	Div(295)	259(303)	281(298)	288(291)	68(35)	45(35)	35(35)	35(36)	36(36)
`Half_Q1`	Div(257)	Div(348)	Div(380)	494(385)	409(391)	Div(32)	Div(34)	172(35)	41(36)	37(37)
`Q1_adapt`	145	237	291	310	313	66	35	37	37	36
`Q1_inner_ad.`	145	213	261	282	289	70	36	35	35	36

Table 1: Number of RAS2 and ORAS2 stationary iterations with advection $a_x = -50y$, $a_y = 50x$, where "Div" means that the iterations are diverging. Laplace results are in parentheses, with "adapted" results then the same as `Q1`.

As already observed in [10] for the symmetric case, we see here that results with the ORAS2 method can be competitive with the multigrid options also in the non-symmetric case when using one of the `Q1`, `Q1_fair`, `Half_Q1` or `Enriched_Q1` coarse spaces (or even `Middle` with GMRES acceleration). Among the various coarse spaces considered, `Half_Q1` exhibits the fastest computational times, most presumably thanks to its lower dimensionality that does not significantly impacts the iteration count (as observed in Fig. 3 up to 1024 cores and as can be verified up to 16,384 cores). This remains true when plotting not only the solving times as in Fig. 4, but the total timings including the setup/assembly phase.

4 Conclusions

We considered several coarse space options for the two-level RAS method applicable to non-symmetric problems and implemented them in the PETSc library. The `Q1` option, that enables a solution in two iterations on a 1-D Laplace test case, shows good performance on our 2-D non-symmetric model problem as well (using coarse points placed around the cross points), in that it has a better iteration count than the `Q1_fair` option (which uses as many but equally distributed coarse points). The new `Q1_adapt` and `Q1_inner_adapt` coarse spaces enable a solution in two iterations for a non-symmetric 1-D advection-diffusion test case, as in the Laplace case in [10]. Despite this promising feature, iteration counts on our 2-D model problem did not show improvements compared to the `Q1` option for moderate advection, but increased robustness was observed for strong advection. The `Enriched_Q1` coarse space, with its higher dimensionality, yields lower iteration counts but appears not to improve the overall computation time. Finally, the new `Half_Q1` coarse space shows promising performance in that the increase in iteration count due to its lower dimensionality appears very moderate and virtually disappears if optimized transmission conditions are introduced (ORAS method). In turn, this option provided the best computational time results in our weak scaling analysis, of the same order of magnitude as multigrid options. Other harmonic coarse spaces like GenEO [15] and GDSW/RGDSW [11] that target improving condition number estimates of Additive Schwarz, in contrast to accelerating low frequency continuous and discontinuous modes of RAS like our new coarse spaces, are also intrinsically based on MsFEM techniques. A more extensive comparison of all these coarse spaces will appear elsewhere.

Acknowledgements This work was performed using HPC resources from GENCI-IDRIS.

References

1. S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, , D. Karpeyev, D. Kaushik, M.G. Knepley, D.A. May, L.Curfman McInnes, R. Tran Mills, T. Munson, K. Rupp, P. Sanan, B.F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Web page. `http://www.mcs.anl.gov/petsc`, 2019.
2. S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, , D. Karpeyev, D. Kaushik, M.G. Knepley, D.A. May, L.Curfman McInnes, R. Tran Mills, T. Munson, K. Rupp, P. Sanan, B.F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.14, Argonne National Laboratory, 2020.
3. S. Balay, W.D. Gropp, L. Curfman McInnes, and B.F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.
4. X.-C. Cai and M. Sarkis. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM J. Sci. Comp.*, 21(2):239–247, 1999.
5. G. Ciaramella and M.J. Gander. Analysis of the parallel Schwarz method for growing chains of fixed-sized subdomains: Part I. *SIAM Journal on Numerical Analysis*, 55(3):1330–1356, 2017.
6. G. Ciaramella and M.J. Gander. Analysis of the parallel Schwarz method for growing chains of fixed-sized subdomains: Part II. *SIAM Journal on Numerical Analysis*, 56(3):1498–1524, 2018.
7. O. Dubois, M.J. Gander, S. Loisel, A. St-Cyr, and D.B. Szyld. The optimized Schwarz methods with a coarse grid correction. *SIAM J. Sci. Comp.*, 34(1):A421–A458, 2012.
8. M.J. Gander. Optimized Schwarz methods. *SIAM J. Numer. Anal.*, 44(2):669–731, 2006.
9. M.J. Gander, L. Halpern, and K. Santugini. A new coarse grid correction for RAS/AS. In *Domain Decomposition Methods in Science and Engineering XXI*, Lecture Notes in Computational Science and Engineering, pages 275–284. Springer-Verlag, 2014.
10. M.J. Gander and S. Van Criekingen. New coarse corrections for restricted additive Schwarz using PETSc. In *Domain Decomposition Methods in Science and Engineering XXV*, Lecture Notes in Computational Science and Engineering, pages 483–490. Springer-Verlag, 2019.
11. A. Heinlein, C. Hochmuth, and A. Klawonn. Reduced dimension GDSW coarse spaces for monolithic Schwarz domain decomposition methods for incompressible fluid flow problems. *International Journal for Numerical Methods in Engineering*, 121(6):1101–1119, 2020.
12. V.E. Henson and U.M. Yang. Boomeramg: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, 41:155–177, 2002.
13. V. Hernandez, J.E. Roman, and V. Vidal. SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362, 2005.
14. A. May, P. Sanan, K. Rupp, M.G. Knepley, , and B.F. Smith. Extreme-scale multigrid components within petsc. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, 2016.
15. N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl. Abstract robust coarse spaces for systems of pdes via generalized eigenproblems in the overlaps. *Numer. Math.*, 126(4):741–770, 2014.
16. A. St-Cyr, M.J. Gander, and S.J. Thomas. Optimized multiplicative, additive, and restricted additive Schwarz preconditioning. *SIAM J. Sci. Comp.*, 29(6):2402–2425, 2007.
17. P. S. Vassilevski and U. M. Yang. Reducing communication in algebraic multigrid using additive variants. *Numer. Linear Algebra Appl.*, 21 (2):275–296, 2014.