Merged Q1 Coarse Spaces for Schwarz Methods in 2D and 3D

Martin J. Gander and Serge Van Criekingen

1 Introduction

Coarse spaces are a key ingredient in the development of scalable domain decomposition methods. Our goal here is to develop new 2D and 3D low-dimensional but robust coarse spaces based on the observation of the eigenmodes of the considered operator, namely here the RAS [4] iteration operator. A first step of this development was presented in [8] with the introduction, for the 2-D case, of the so-called Half_Q1 coarse space, which will here be re-named as Merged_2 coarse space in view of its 3D generalization. Our merged coarse spaces are obtained by taking a subset of the coarse space functions of the Q1 coarse space [6, 7], which is based, in 2D, on the Q1 bilinear finite element functions on rectangular elements, these elements being here the subdomains. This leads, in 2D, to four coarse grid points around each cross point of the non-overlapping decomposition (Fig. 2(a)). In 3D, the equivalent Q1 coarse space has eight coarse grid points around each cross point (Fig. 2(b)). The merged coarse spaces we consider will be described in section 2 for the 2D case and 3 for the 3D case. Our coarse spaces are implemented using PETSc [1, 2, 3] which, besides benefiting from its parallel features, enables us to compare our optimized two-level domain decomposition results, i.e., using the optimized RAS (ORAS [9, 12]) method, with multigrid options (GAMG, HYPRE). Our model problem is the Laplace problem $-\Delta u = 0$ or, adding advection with an upwind scheme, the non-symmetric problem

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

using the 5-point finite difference scheme and homogeneous boundary conditions.

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

Serge Van Criekingen
Institut du Développement et des Ressources en Informatique Scientifique (IDRIS), CNRS, Université Paris-Saclay, F-91403, Orsay, France and Université Paris-Saclay, UVSQ, CNRS, CEA, Maison de la Simulation, 91191, Gif-sur-Yvette, France, e-mail: serge.van.criekingen@idris.fr

2 Merged coarse spaces in 2D

In Fig. 1 we show the four largest eigenmodes (i.e., with eigenvalues closest to one in modulus) of the RAS iteration operator (i.e., $I - BA$ with A the stiffness matrix and B the RAS preconditioner) for the 2D Laplace model problem using a 2×2 algebraically non-overlapping decomposititon (RAS then reduces to Block Jacobi), for a global 256×256 fine mesh resolution. The first pair of eigenmodes exhibits eigenvalues equal in modulus and of opposite signs, and similarly for the second pair. Moreover, the first two modes appear to be made of piecewise Q1 functions (as was first shown in [5]): denoting by q_1, q_2, q_3, q_4 the Q1 basis functions at a cross point (with the numbering convention of Fig. 2(a)), i.e., bilinear with value 1 at the cross point corner and 0 at the other corners of the subdomain (see the squared points in Fig. 2(c)), the first mode (continuous, with positive eigenvalue) appears to be $q_1 + q_2 + q_3 + q_4$, while the second mode (discontinuous, with negative eigenvalue) appears to be $q_1 - q_2 - q_3 + q_4$. Therefore the idea of the Merged_2 coarse space (already introduced in [8] under the name Half_Q1) is to merge the four basis

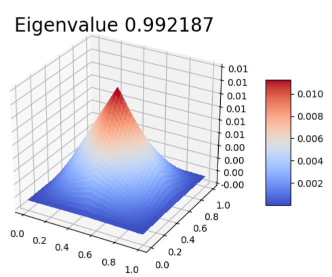

(a) mode 1

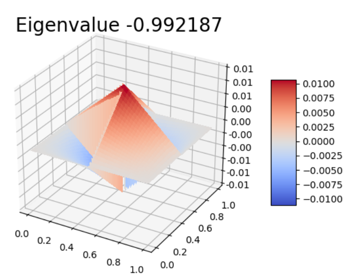

(b) mode 2

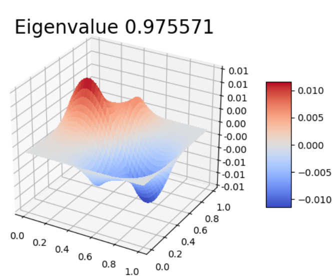

(c) mode 3

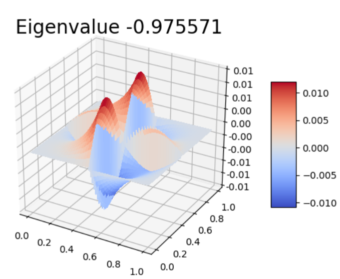

(d) mode 4

Fig. 1: Largest four modes of the RAS iteration operator for the Laplace problem using a 2×2 algebraically non-overlapping domain decomposition.

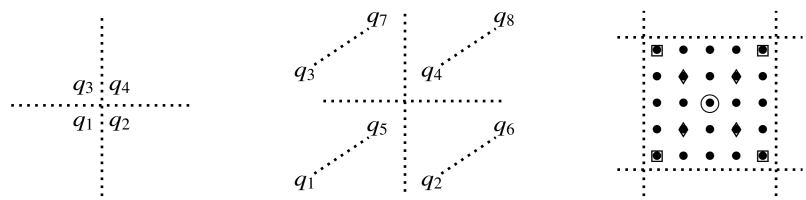

(a) 2D cross point (b) 3D cross point (c) 2D coarse point options

Fig. 2: (a) and (b): Numbering conventions of the Q1 basis functions at a 2D and 3D cross point. (c): Coarse grid point choice in 2-D for the Q1 (squares), Q1_fair (diamonds) and Middle (circles) options.

functions of the Q1 coarse space into these two combinations, yielding a coarse space which is half the size of the original Q1 one. For the two largest eigenmodes, this reduction of the coarse space size comes at no cost. Note that the Merged_2 coarse space can equivalently be defined by taking the combinations $q_1 + q_4$ and $q_2 + q_3$ at each cross point.

We can further pursue the merging idea for the case where RAS is used with GMRES acceleration. Discontinuous eigenmodes of the RAS iteration operator $I - BA$ have negative eigenvalues close to -1, see e.g. Fig. 1, which implies that in the preconditioned system BA they are close to 2 and thus can easily be handled by GMRES. This is in contrast to the continous eigenmodes which have eigenvalues close to 1, and thus in the preconditioned system they are close to 0, which is hard for GMRES, and thus these modes must be handled by another mechanism, which here is the coarse space. Thus, we define the Merged_1 coarse space as the one where only the $q_1 + q_2 + q_3 + q_4$ combination of the Q1 basis functions (i.e., the largest continuous mode) is kept, yielding a coarse space which is four times smaller than Q1 and half the size of Merged_2. Table 1(a) shows the number of iterations necessary to reach the zero solution of the model Laplace problem, when successively taking as initial guess the largest four modes of the RAS iteration operator displayed in Fig. 1, and using stationary iterations with or without coarse correction, or using GMRES acceleration. Using stationary iterations, the Q1 coarse space has the first two modes converge in one iteration. By construction, the Merged_2 coarse space also contains the first two modes and has them converge in one iteration, while the Merged_1 coarse space contains only the largest continuous mode. The last line of Table 1(a) shows that the GMRES acceleration, already without coarse correction, yields convergence in one iteration for the largest (and second largest) discontinuous mode. The Merged_1 coarse space is thus meant to be used with GMRES acceleration, avoiding an unnecessary duplication of effects, at least for the first two modes.

Mode number	1	2	3	4
Eigenvalue	0.99219	-0.99219	0.97557	-0.97557
Stationary iterations				
no coarse correction	2349	2349	745	745
`Q1`	1	1	551	417
`Merged_2`	1	1	745	745
`Merged_1`	1	2349	745	745
`GMRES` accelerated				
no coarse correction	17	1	23	1

(a) 2D case (256^2 fine mesh resolution)

Eigenvalue	0.96875	-0.96875	0.93035	-0.93035
Stationary iterations				
no coarse correction	581	581	256	256
`Q1`	1	2	175	141
`Merged_2`	1	2	256	256
`Merged_1`	1	581	256	256
`GMRES` accelerated				
no coarse correction	8	1	7	1

(b) 3D case (64^3 fine mesh resolution)

Table 1: Number of iterations to reach the zero solution of the Laplace model problem with the largest four `RAS` iteration operator eigenmodes taken as initial guess.

3 Merged coarse spaces in 3D

Our numbering convention of the `Q1` basis functions at a 3D cross point is given in Fig. 2(b). We define our 3D merged coarse spaces using the same symmetry as the one used in 2D. That is, the `Merged_2` coarse space will use the following combinations as basis functions: $\sum_i q_i$ and $q_1 - q_2 - q_3 + q_4 - q_5 + q_6 + q_7 - q_8$ or, equivalently, $q_1 + q_4 + q_6 + q_7$ and $q_2 + q_3 + q_5 + q_8$. As for the `Merged_1` coarse space, it will use the sum $\sum_i q_i$ as unique basis function.

We consider again the four largest eigenmodes of the `RAS` iteration operator for the Laplace model problem, this time using a $2 \times 2 \times 2$ algebraically non-overlapping decomposititon. These modes can not be plotted as in 2D but we have that, as in 2D, the first pair of eigenmodes exhibits eigenvalues equal in modulus and of opposite signs, and similarly for the second pair: the actual values are given on the second line of Table 1(b). Moreover, this table shows that the `Q1` and `Merged_2` coarse spaces contain the first two modes, while the `Merged_1` coarse space contains only the first (positive) mode[1]. This justifies the definition of our 3D merged coarse spaces. Moreover, here again the `GMRES` acceleration alone yields convergence in one iteration for

[1] The number of iterations for the mode number 2 with `Q1` and `Merged_2` coarse spaces is 2 and not 1, but this is due only to its high sensitivity to the convergence criteria: taking a convergence criteria of 1.05E-8 instead of 1.E-8 would have made the convergence achieved at iteration 1.

the largest (and second largest) negative mode, making the Merged_1 coarse space a method of choice when this acceleration is used.

4 Numerical Results

We present numerical results of a weak scaling experiment for our model problem, using square domain decompositions (2×2, 4×4, ... and similarly in 3D). Each run starts from a random initial guess and converges to the zero solution with a relative convergence tolerance of 1.e-8. Our results were obtained on the CPU partition of the Adastra machine hosted at CINES in Southern France (www.cines.fr), using one CPU per subdomain. We compare our merged coarse spaces to other coarse spaces, namely the Nicolaides coarse space [11] which uses a constant basis function on each subdomain, as well as, in 2D, to other coarse spaces using bilinear functions centered on coarse points depicted in Fig. 2(c): while the Q1 coarse space has coarse points in each corner of the subdomain, the Q1_fair coarse space has as many coarse points as the Q1 coarse space but equally distributed in space, and the Middle coarse space takes the fine mesh point in the middle of each subdomain as coarse grid point. Note that these coarse spaces were already considered in [8]. For an $N \times N$ domain decomposition, the asymptotic size of the coarse space is thus $4N^2$ for Q1 and Q1_fair , $2N^2$ for Merged_2 , and N^2 for Merged_1 , Middle and Nicolaides .

In Fig. 3(a) we show the number of iterations needed for convergence when RAS is used without GMRES acceleration on the algebraically non-overlapping Laplace problem. As expected, the Merged_1 coarse space yields very poor results, while the Merged_2 coarse space yields results which are not too far away from the Q1 ones, notably better than the Q1_fair methods which uses a twice bigger coarse space. With GMRES (Fig. 3(b)), the Merged_1 results come very close to the Merged_2 ones (they are in fact undistinguishable on the plot), as expected. Note that Q1_fair yields here the best results: it appears to take more advantage of the GMRES acceleration than the Q1 method and its merged variants.

In Figs. 3(c) and 3(d) we show the number of iterations and solution times for the model problem (1) with rotating advection $a_x = -10y, a_y = 10x$, and an algebraic overlap of 2, using the optimized RAS (ORAS) method based on optimized transmission conditions at subdomain interfaces [9, 12]. The ORAS method is implemented with a first-order accurate discretization of the normal derivative, which implies modifying only the diagonal components of the local system matrix, and the Robin parameters are the ones derived for the (symmetric) Laplace problem [6]. The number of iterations of the Q1, Merged_2 and Merged_1 options are close in Fig. 3(c), here with Q1 and Merged_2 undistinguishable and Merged_1 slightly above. Note that the number of iterations with 128^2 (= 16,384) subdomains (not on the plot) are the same as the ones shown here at 64^2 (= 4,096) subdomains.

In Fig. 3(d) we show solution times corresponding to Fig. 3(c), as well as coarse correction times per iteration. As expected, the smaller size of the merged coarse

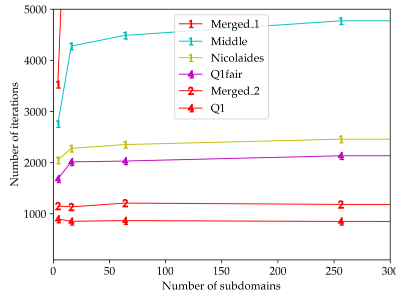

(a) Iterations with `RAS` (no `GMRES`).

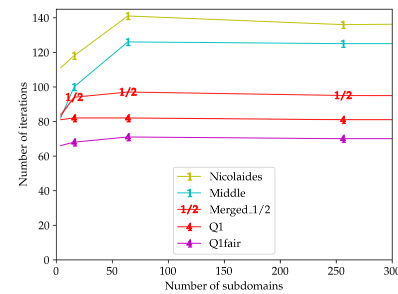

(b) Iterations with `RAS` using `GMRES`.

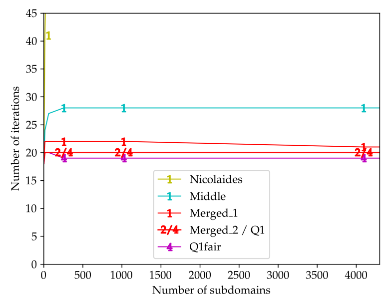

(c) Iterations with `ORAS` using `GMRES`, with overlap and advection.

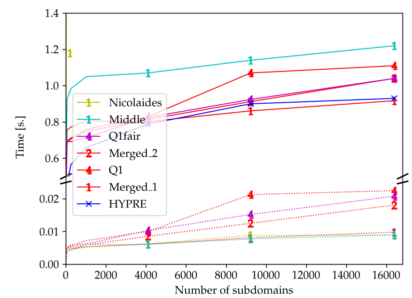

(d) Solution (solid) and coarse correction times per iteration (dotted) with `ORAS` and multigrid, using `GMRES`, with overlap and advection.

Fig. 3: Number of iterations (a)(b)(c) and solution times (d) for the model problem (1) in 2D using a 256×256 local mesh. In (a) and (b) $a_x = a_y = 0$ and the algebraic overlap is 0. In (c) and (d) $a_x = -10y,\ a_y = 10x$, and the algebraic overlap is 2.

spaces reduces the coarse correction time per iteration compared to the original `Q1` coarse space (`Merged_1`, `Middle` and `Nicolaides` are nearly undistiguishable, in accordance to their same asymptotic size) and, since the merging process comes at a very moderate price in terms of iterations, it results in a total time to solution gain. Moreover, our `Merged_1` results appear comparable to the ones obtained with the algebraic multigrid `HYPRE/BoomerAMG` option [10] available through `PETSc` (with optimized tuning form [13]). However, if one includes setup costs, the multigrid option keeps its advantage: for instance at 128^2 cores, the `Merged_1` setup cost (which notably includes the local LU factorizations) is .93s, while it is only .03s for the `HYPRE` option.

In Figs. 4(a) and 4(b) we show the number of iterations for the Laplace problem in 3D, with and without `GMRES` acceleration, using an algebraic overlap of 2 and a 64^3 local mesh. As expected, the `Merged_1` coarse space performs poorly (even if not as bad as in 2D) without `GMRES`, but yields results very similar to `Merged_2` with `GMRES`.

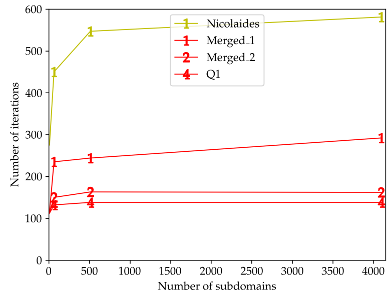

(a) Iterations with `RAS` (no `GMRES`).

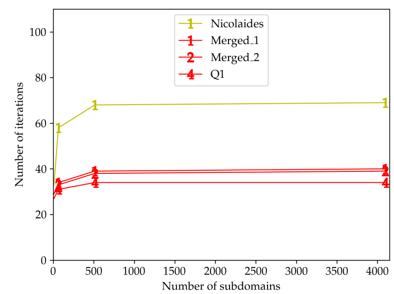

(b) Iterations with `RAS` using `GMRES`.

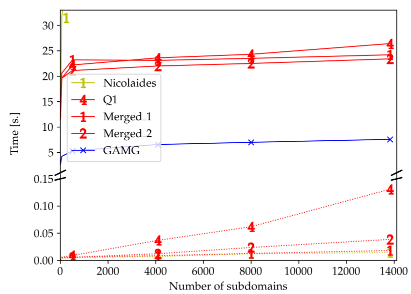

(c) Solution time with `ORAS` and multigrid, using `GMRES` (64^3 local mesh).

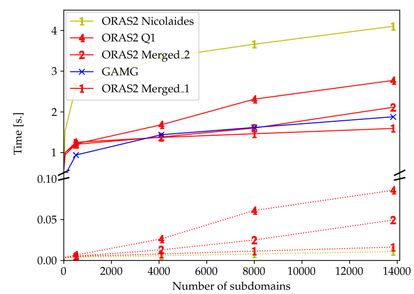

(d) Solution time with `ORAS` and multigrid, using `GMRES` (32^3 local mesh).

Fig. 4: Number of iterations (a)(b) and solution times (c)(d) for the Laplace problem in 3D using a 64^3 (a)(b)(c) or 32^3 (d) local mesh and an algebraic overlap of 2.

Moreover, the `GMRES` `Merged_1` and `Merged_2` results are close to the `Q1` result, showing again that the decrease in coarse space size resulting from the merging process comes at a low cost in terms of iterations. In Fig. 4(c) we show time to solution results for the same Laplace problem but with `ORAS`, as well as the corresponding coarse correction times per iteration. As in 2D and as expected, the `Merged_1` coarse space has its coarse correction time very close to the `Nicolaides` one and smaller than `Merged_2`, itself smaller than `Q1`. In turn, the merged coarse spaces tend to perform better than the `Q1` one even if, at 24^3 (= 13,824) cores, the `Merged_2` coarse space remains a bit better than the `Merged_1`. The solution time with `PETSc`'s native algebraic multigrid preconditioner `GAMG` (with smoothed aggregation and CG eigenvalue estimator [2]) is also given, with here a clearly better performance. However, the algebraic multigrid solution time performance (outside of setup costs[2]) is approached more closely by our merged coarse space if one considers a smaller local mesh, for instance 32^3, as shown in Fig. 4(d) (with the expected hierarchy:

[2] At 24^3 cores: 1.26s. (`ORAS` with `Merged_1`) vs. 0.02s. (`GAMG`) with the 32^3 local mesh, 23.19s. vs. 0.03s. with the 64^3 local mesh.

`Merged_1` faster than `Merged_2`, itself faster than `Q1`). The number of subdomains, thus the number of cross points and, in turn, our merged coarse space size, is then relatively larger with respect to the global number of fine mesh points. This suggests that the efficiency of the `Merged_Q1` coarse space compared to multigrid options could be improved by using more than one subdomain per `CPU` core.

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

References

1. Balay, S., Abhyankar, S., Adams, M.F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E.M., Dalcin, L., Dener, A., Eijkhout, V., Faibussowitsch, J., Gropp, W.D., Hapla, V., Isaac, T., Jolivet, P., Karpeev, D., Kaushik, D., Knepley, M.G., Kong, F., Kruger, S., May, D.A., McInnes, L.C., Mills, R.T., Mitchell, L., Munson, T., Roman, J.E., Rupp, K., Sanan, P., Sarich, J., Smith, B.F., Zampini, S., Zhang, H., Zhang, H., Zhang, J.: PETSc Web page. `https://petsc.org/` (2024). URL `https://petsc.org/`
2. Balay, S., Abhyankar, S., Adams, M.F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E.M., Dalcin, L., Dener, A., Eijkhout, V., Faibussowitsch, J., Gropp, W.D., Hapla, V., Isaac, T., Jolivet, P., Karpeev, D., Kaushik, D., Knepley, M.G., Kong, F., Kruger, S., May, D.A., McInnes, L.C., Mills, R.T., Mitchell, L., Munson, T., Roman, J.E., Rupp, K., Sanan, P., Sarich, J., Smith, B.F., Zampini, S., Zhang, H., Zhang, H., Zhang, J.: PETSc/TAO users manual. Tech. Rep. ANL-21/39 - Revision 3.21, Argonne National Laboratory (2024)
3. Balay, S., Gropp, W., McInnes, L.C., Smith, B.: Efficient management of parallelism in object oriented numerical software libraries. In: E. Arge, A.M. Bruaset, H.P. Langtangen (eds.) Modern Software Tools in Scientific Computing, pp. 163–202. Birkhäuser Press (1997)
4. Cai, X.C., Sarkis, M.: A restricted additive Schwarz preconditioner for general sparse linear systems. SIAM J. Sci. Comp. **21**(2), 239–247 (1999)
5. Cuvelier, F., Gander, M.J., Halpern, L.: Fundamental coarse space components for schwarz methods with crosspoints. In: Domain Decomposition Methods in Science and Engineering XXVI, Lecture Notes in Computational Science and Engineering, pp. 39–50. Springer-Verlag (2021)
6. Dubois, O., Gander, M., Loisel, S., St-Cyr, A., Szyld, D.: The optimized Schwarz methods with a coarse grid correction. SIAM J. Sci. Comp. **34**(1), A421–A458 (2012)
7. Gander, M., Halpern, L., Santugini, K.: A new coarse grid correction for RAS/AS. In: Domain Decomposition Methods in Science and Engineering XXI, Lecture Notes in Computational Science and Engineering, pp. 275–284. Springer-Verlag (2014)
8. Gander, M., Van Criekingen, S.: Coarse corrections for Schwarz methods for symmetric and non-symmetric problems. In: Domain Decomposition Methods in Science and Engineering XXVI, Lecture Notes in Computational Science and Engineering, pp. 589–596. Springer-Verlag (2021)
9. Gander, M.J.: Optimized Schwarz methods. SIAM J. Numer. Anal. **44**(2), 669–731 (2006)
10. Henson, V., Yang, U.: Boomeramg: a parallel algebraic multigrid solver and preconditioner. Applied Numerical Mathematics **41**, 155–177 (2002)
11. Nicolaides, R.A.: Deflation of conjugate gradients with applications to boundary value problems. SIAM Journal on Numerical Analysis **24**, 355–365 (1987)
12. St-Cyr, A., Gander, M., Thomas, S.: Optimized multiplicative, additive, and restricted additive Schwarz preconditioning. SIAM J. Sci. Comp. **29**(6), 2402–2425 (2007)
13. Vassilevski, P.S., Yang, U.M.: Reducing communication in algebraic multigrid using additive variants. Numer. Linear Algebra Appl. **21 (2)**, 275–296 (2014)