Spectral `Q1`-based coarse spaces for Schwarz methods

Martin J. Gander and Serge Van Criekingen

1 Introduction

The `Q1` coarse space [DGL$^+$12, GHS14] is based on coarse `Q1` bilinear finite element functions on rectangular elements which are here the subdomains. Hence the coarse grid points are placed (in 2-D) around each cross point of the non-overlapping decomposition. It was studied by the authors in [GV19] and [GV21], together with several of its variants, in the context of the Restricted Additive Schwarz (RAS) method [CS99] with optimized Robin-type transmission conditions [Gan06]. Encouraging numerical results were obtained in that the resulting method, implemented in `PETSc` [BAA$^+$22a, BAA$^+$22b, BGMS97], showed computing times competitive with multigrid approaches on a 2-D Laplace test case, for both symmetric and non-symmetric (i.e., with advection) problems. Among the different options invesitigated in [GV21], the so-called `Half_Q1` (see also [CGH21]) appeared most promising, in that it halves the coarse space dimension compared to `Q1` by using a selected combination of its basis functions, while causing only a moderate increase in iteration count, resulting in our best observed computing times. We therefore pursue here the investigation around this `Half_Q1` coarse space and, more generally, `Q1`-based spectral coarse spaces [CV21], that is, coarse spaces based on the study of the eigenvectors of the underlying iteration operator, in our case the RAS iteration operator. Note however that we here do not compute a spectrum specific to each problem (as for instance in [CV21]): we define our coarse spaces based on the observation of the eigenmodes of the non-overlapping symmetric Laplace test case and hope that the resulting method will succesfully apply to a broader set of problems, as was the case

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

in [GV21] adding overlap and advection. Note that we here restrict our analysis to homogeneous Dirichlet boundary conditions.

As already pointed out in [GV21], the largest two eigenvalues (i.e. closest to 1 in modulus) of the RAS iteration operator for the 2-D symmetric Laplace model problem appear to be equal in modulus and of opposite signs, while the corresponding eigenvectors appear to be one continuous (for the positive eigenvalue) and one discontinuous (for the negative one). We display these eigenmodes[1] in Fig. 1 for various square domain decompositions in the algebraically non-overlapping case (RAS then reduces to Block Jacobi), as obtained using the `SLEPc` [HRV05] eigenproblem companion package to `PETSc`, with the traditional 5-point finite difference discretization. These modes appear to be piecewise `Q1` functions.

The `Q1` basis functions at a cross point will be denoted by q_1, q_2, q_3, q_4 (i.e., bilinear with value 1 at the cross point and 0 at the other corners of the subdomain - see Fig. 2a), the four of them building up a "hat" around the cross point. (Note that we do not need to solve eigenproblems to use these `Q1` functions.) The `Half_Q1` coarse space is based on the observation of the 2×2 eigenmodes (Fig. 1a and 1b): these modes appear to be particular combinations of the `Q1` basis functions at the cross point, namely $q_1 + q_2 + q_3 + q_4$ (the "hat" itself) and $q_1 - q_2 + q_3 - q_4$. 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. This is equivalent to taking the combinations $q_1 + q_4$ and $q_2 + q_3$ at each cross point.

By construction, the `Half_Q1` space contains the first two eigenmodes of the non-overlapping RAS iteration operator in the case of a 2×2 decomposition. In turn, taking one of these first two RAS eigenmodes as initial guess of a coarse corrected (i.e., two-level) RAS iteration process, we obtain convergence at iteration 1 using the `Half_Q1` coarse space (with square subdomains and a non-overlapping decomposition). For more than 2×2 subdomains, convergence at iteration 1 does not hold, but it still holds (with the same restrictions) for square decompositions using the `Q1` coarse space, and it is moreover possible to define a `Half_Q1`$^+$ coarse space, larger than `Half_Q1` but smaller than `Q1`, so as to include the first two modes, i.e., so that this convergence at iteration 1 is verified. This will be described in section 2.

Another new `Q1`-based coarse space, named `Checkerboard`, is introduced in section 3. Based on the first two modes of the decomposition considered (not only the 2×2 one), it can be applied to non-square decompositions.

2 The `Half_Q1`$^+$ coarse space

The `Half_Q1`$^+$ coarse space is built by adding a minimal number of extra basis functions to the `Half_Q1` coarse space so as to contain the first two eigenmodes of the RAS iteration operator. It is meant to be smaller than the `Q1` coarse space.

---

[1] The problem solved here in thus $Gx = \lambda x$ where $G := I - \left(\sum_{j=1}^{J} \tilde{R}_j^T A_j^{-1} \tilde{R}_j \right) A$, and R_j are restriction operators to the J non-overlapping subdomains decomposing the global domain.
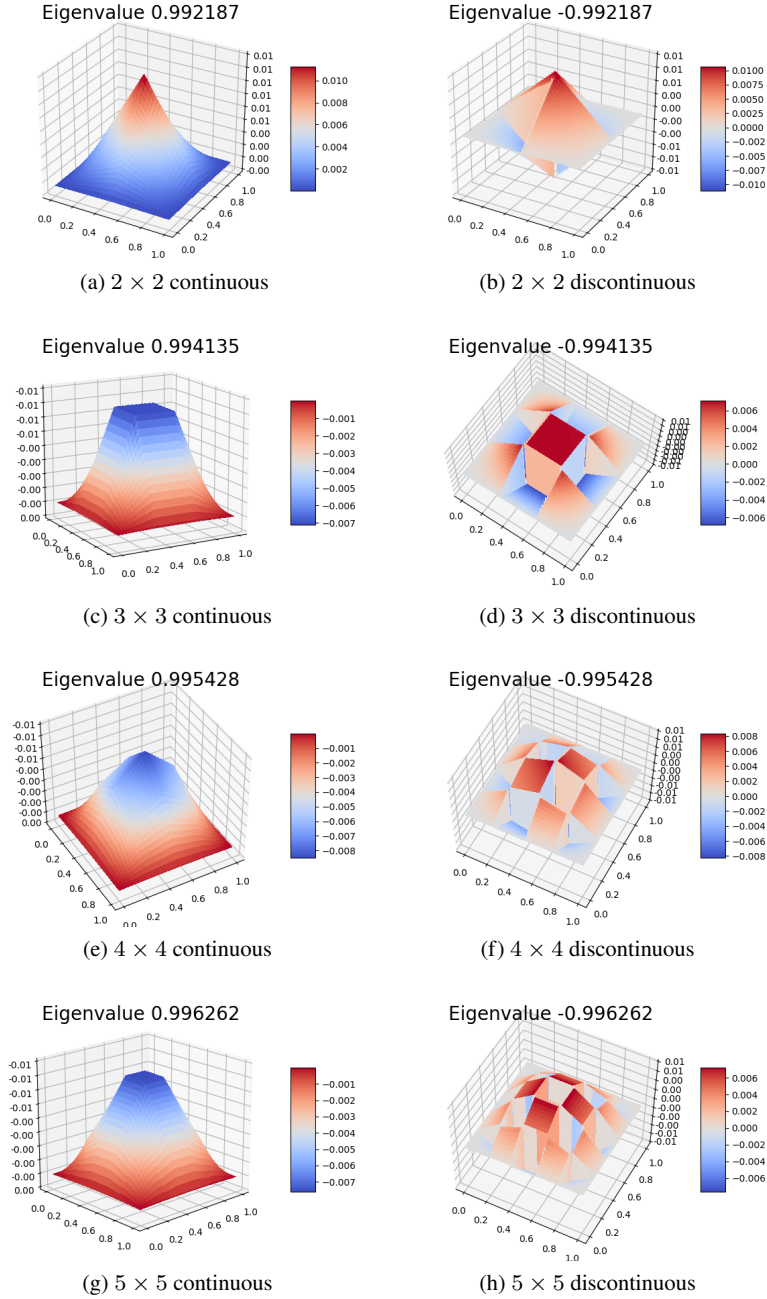
Eigenvalue 0.992187

Eigenvalue -0.992187

(a) 2×2 continuous

(b) 2×2 discontinuous

Eigenvalue 0.994135

Eigenvalue -0.994135

(c) 3×3 continuous

(d) 3×3 discontinuous

Eigenvalue 0.995428

Eigenvalue -0.995428

(e) 4×4 continuous

(f) 4×4 discontinuous

Eigenvalue 0.996262

Eigenvalue -0.996262

(g) 5×5 continuous

(h) 5×5 discontinuous

Fig. 1: Eigenmodes of the non-overlapping RAS iteration operator corresponding to the two largest eigenvalues in modulus for the 2×2 to 5×5 decompositions, for a global 256x256 fine mesh resolution.
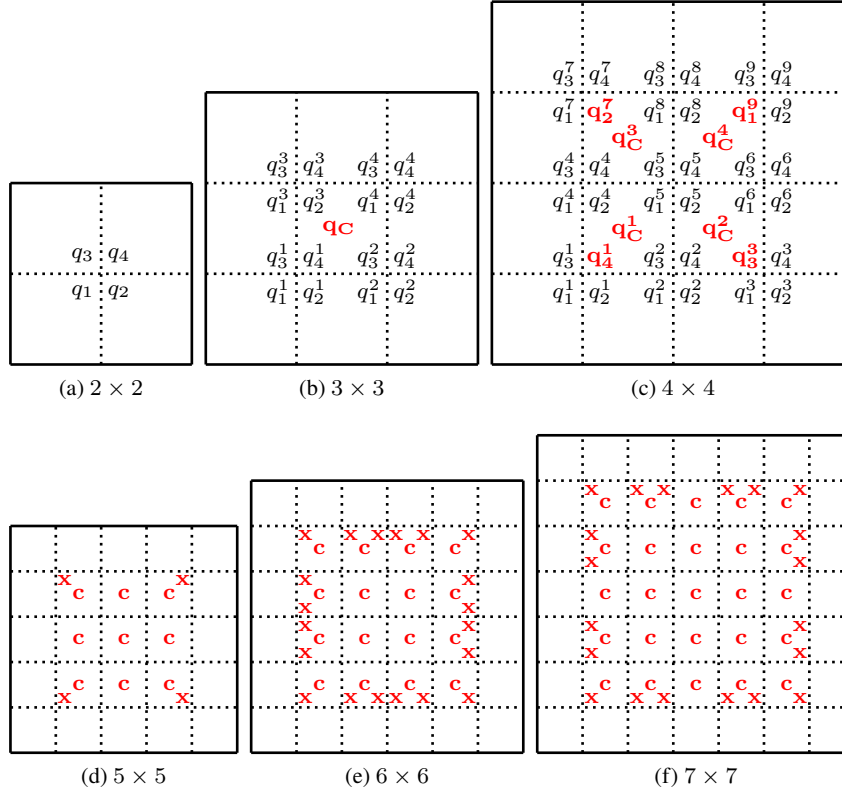
Fig. 2: In red, basis functions to be added to the `Half_Q1` coarse space to obtain the `Half_Q1`$^+$ one, for various decompositions. $q_1^i, q_2^i, q_3^i, q_4^i$ are the `Q1` basis functions at cross point i and $\mathbf{q_C}$ represents a constant function in the considered subdomain. For the 5×5 to 7×7 decompositions, only a schematic view is given, with c representing a constant and x a `Q1` basis function.

For the 2×2 decomposition, the `Half_Q1`$^+$ coarse space is the same as the `Half_Q1` one. This is not the case anymore for the 3×3 decomposition: starting from one of the first two modes of the RAS iteration operator (Figs. 1c and 1d), convergence of the `Half_Q1` coarse-corrected RAS iteration process is not obtained at iteration 1, while it is the case with `Q1`. But what is missing in `Half_Q1` to achieve convergence at iteration 1? Observing Figs. 1c and 1d, one can intuitively infer that adding a constant coarse function in the central subdomain to the `Half_Q1` coarse space will greatly improve convergence. Our numerical implementation showed that this is actually sufficient to obtain convergence at iteration 1. The `Half_Q1`$^+$ coarse space is thus obtained from `Half_Q1` by adding one single constant coarse function

in the central subdomain ($\mathbf{q_C}$ in Fig. 2b) and is of size 9 (8 for `Half_Q1` and 16 for `Q1`).

For the 4×4 decomposition, the first two RAS modes are given in Figs. 1e and 1f. In this case, the minimal function set we found to add to `Half_Q1` to resolve the first two modes is made out of one constant coarse function on each inner subdomain as well as the extra `Q1` basis functions located at the four "inner corners" one subdomain away from the boundary, namely q_4^1, q_3^3, q_2^7, q_1^9 in Fig. 2c. Thus, for this decomposition, `Half_Q1`$^+$ is of size 26 (18 for `Half_Q1` and 36 for `Q1`).

We pursued our investigations for larger $N \times N$ decompositions and observed that the extra basis functions to be added to the `Half_Q1` coarse space to build `Half_Q1`$^+$ remain of two types, namely constants on each non-boundary subdomain and extra `Q1` basis functions located one subdomain away from the boundary, as described schematically in Figs. 2d, 2e and 2f. Note that for N odd, the extra basis functions on the "middle" subdomain on each side (one subdomain away from the boundary) appear not to be needed (see Figs. 2d and 2f). However, these appear to be needed in the case $N = 11, 15, 19, 23, ...$ This was tested numericallly up to $N = 50$, i.e., 2500 subdomains. Note that, while the size of `Q1` and `Half_Q1` asymptotically grow as $4N^2$ and $2N^2$ respectively, the size of `Half_Q1`$^+$ grows as $3N^2$.

Once defined, the `Half_Q1`$^+$ coarse space can be used in a general context: weak scaling experiment results for RAS with overlap 1 (in the `PETSc` sense, i.e, algebraic overlap of 2) are given in Fig. 3a. Starting from a random initial guess, the number of iterations and computing times necessary to bring the relative tolerance below 1.e-8 are given. In terms of iterations, `Half_Q1`$^+$ tends to behave asymptotically like `Q1`, while using only $3N^2$ coarse functions instead of $4N^2$. In terms of computing time, `Half_Q1`$^+$ yields scalable timings very close to `Q1` and better than `Half_Q1`.

Note that the `Half_Q1`$^+$ coarse space is not defined in the general rectangular case since then even the `Q1` coarse space does not resolve the first two non-overlapping RAS eigenmodes. This is the case for instance for the 3×2 subdomain case whose eigenmodes are depicted in Figs. 4a and 4b. A close observation of these plots reveals that what appears as a horizontal edge at $y = .5$ (assuming three subdomains in x and two in y) is in fact slightly curved, giving an intuitive explanation to the non-inclusion of these modes into the `Q1` coarse space.

3 The `Checkerboard` coarse space

The `Checkerboard` coarse space is based on the first two modes of the decomposition considered, not only the 2×2 one as in the `Half_Q1` case. It is defined for square and rectangular decompositions, and contains 2 modes. For the 3×2 case and as illustrated in Fig. 3d these two modes are $q_1^1 + q_4^1 + q_3^2 + q_2^2$ and $q_3^1 + q_2^1 + q_1^2 + q_4^2$. This definition comes from the observation of Fig. 4a and 4b. Starting from one of these two modes, we observed that the `Checkerboard` coarse space gives the exact same iterates as `Half_Q1` but with 2 coarse functions instead of 4.
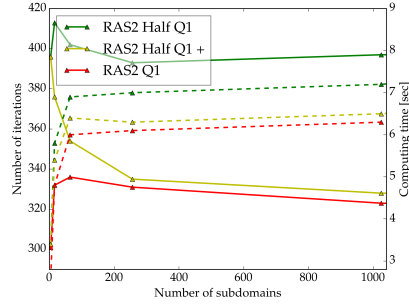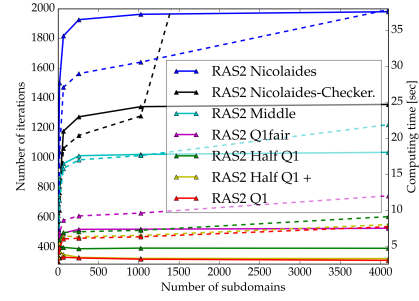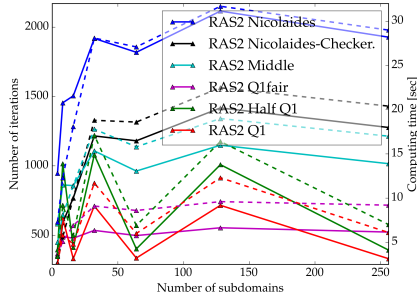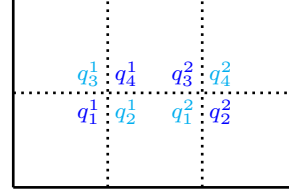
(a) Square decompositions up to 32×32.

(b) Square decompositions up to 64×64.

(c) With 4×2, 8×4 and 16×8 decompositions.

(d) 3×2 `Checkerboard`

Fig. 3: (a) to (c): Weak scaling experiment for overlapping RAS2 (256x256 fine mesh per subdomain) for various decompositions. Solid: number of iterations, dashed: computing times. (d): 3×2 `Checkerboard` coarse basis function definition.
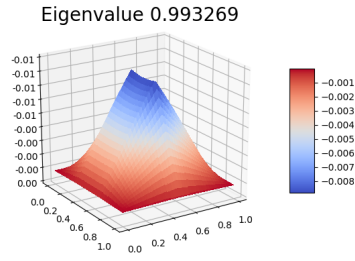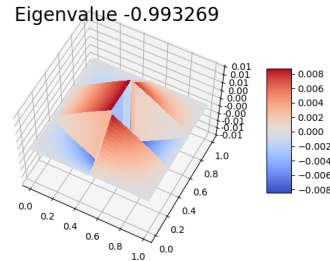

(a) 3×2 continuous

(b) 3×2 discontinuous

Fig. 4: Eigenmodes of the non-overlapping RAS iteration operator corresponding to the two largest eigenvalues in modulus in the case of 3×2 subdomains, for a global 256x256 fine mesh resolution.

For the 3×3 case, the two `Checkerboard` modes are defined to be (using the numbering in Fig. 2b and not including the constants) $q_1^1 + q_2^2 + q_3^3 + q_4^4 + q_4^1 + q_3^2 + q_2^3 + q_1^4$ for the first mode and the sum of the 8 other q_j^i for the second mode. This comes from the observation of Figs. 1c and 1d. It again produces the same iterates as `Half_Q1` but with 2 coarse functions instead of 8.

For the 4×4 case, the observation of Figs. 1e and 1f leads us to define the first `Checkerboard` coarse basis functions as (using the numbering in Fig. 2c and grouping the q_j^i functions by subdomain) $q_1^1 + (q_2^2 + q_1^3) + (q_4^1 + q_3^2 + q_2^4 + q_1^5) + (q_3^3 + q_2^6) + (q_4^4 + q_1^7) + (q_4^5 + q_3^6 + q_2^8 + q_1^9) + (q_4^7 + q_3^8) + q_4^9$ and the second one as made out of the other q_j^i. Here these two `Checkerboard` coarse functions do not produce the same iterates as the `Half_Q1` coarse functions (18 in this case). This is not surprising since no scalability can be achieved with only two coarse functions.

Nevertheless, it is still possible to obtain a scalable - at least in terms of iterations - two-level method based on the two `Checkerboard` functions, by adding the constant function in each subdomain, yielding a coarse space of size $N^2 + 2$ that will be named `Nicolaides-Checkerboard` since it is the same as the Nicolaides coarse space [Nic87] but with two extra basis functions. Fig. 3b presents the same weak scaling experiment as Fig. 3a, but extended up to 4096 cores and also to other coarse spaces defined in [GV21], namely `Middle` (classical coarse space with one coarse point in the middle of each subdomain) and `Q1_fair` (same number of coarse mesh points as `Q1`, but equally distributed in space). The two extra `Checkerboard` functions yield a major improvement to the Nicolaides coarse space in terms of number of iterations, and this improvement is scalable in that it remains as effective when increasing the number of subdomains. In terms of computing time, the new coarse space appears not scalable above 1024 cores: the coarse solve (performed here with a parallel direct solver) remains a challenge, the two extra functions implying the whole domain.

Fig. 3c includes non-square decompositions up to 16 subdomains in one direction. These appear to require more iterations (and computing time) than their square counterparts. For the `Half_Q1` coarse space, this can be related to the absence of affine modes for non-square subdomains pointed out in [CGH21].

4 Conclusions

We introduced two new `Q1`-based coarse spaces. Firstly, the `Half_Q1`$^+$ coarse space is built from `Half_Q1` (thus from the first two RAS modes of the 2×2 decomposition) so as to contain the first two RAS modes of the considered (square) decomposition while using a minimal set of coarse functions in order to remain smaller than `Q1`. It was shown to behave asymptotically like `Q1` in terms of number of iterations, but using $3N^2$ coarse functions instead of $4N^2$. Secondly, the `Checkerboard` coarse space is built as the first two RAS modes of the decomposition considered and can be defined for square and rectangular decompositions. Combined with Nicolaides into the `Nicolaides-Checkerboard` coarse space, it yields a significant improvement in terms of number of iterations. Its scalability in time is still under investigation.

References

[BAA⁺22a] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. M. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, J. Faibussowitsch, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. Curfman McInnes, R. Tran Mills, L. Mitchell, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang. PETSc Web page. https://petsc.org/, 2022.

[BAA⁺22b] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. M. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, J. Faibussowitsch, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. Curfman McInnes, R. Tran Mills, L. Mitchell, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang. PETSc/TAO users manual. Technical Report ANL-21/39 - Revision 3.18, Argonne National Laboratory, 2022.

[BGMS97] 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.

[CGH21] F. Cuvelier, M. J. Gander, and L. Halpern. 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, pages 39–50. Springer-Verlag, 2021.

[CS99] 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.

[CV21] G. Ciaramella and T. Vanzan. On the asymptotic optimality of spectral coarse spaces. In Domain Decomposition Methods in Science and Engineering XXVI, Lecture Notes in Computational Science and Engineering, pages 181–188. Springer-Verlag, 2021.

[DGL⁺12] 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.

[Gan06] Martin J. Gander. Optimized Schwarz methods. SIAM J. Numer. Anal., 44(2):669–731, 2006.

[GHS14] 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.

[GV19] 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.

[GV21] M.J. Gander and S. Van Criekingen. 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, pages 589–596. Springer-Verlag, 2021.

[HRV05] 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.

[Nic87] R. A. Nicolaides. Deflation of conjugate gradients with applications to boundary value problems. SIAM Journal on Numerical Analysis, 24:355–365, 1987.