Parallel Overlapping Schwarz with an Energy-Minimizing Coarse Space

Alexander Heinlein, Axel Klawonn, and Oliver Rheinbach

1 Introduction and Description of the Method

The GDSW preconditioner is a two-level overlapping Schwarz preconditioner introduced in Dohrmann et al. (2008a) with a proven condition number bound for the general case of John domains for scalar elliptic and linear elasticity model problems. It is algebraic in the sense that it can be constructed from the assembled system matrix. However, compared to FETI-DP (see Toselli and Widlund 2005) or BDDC methods, in GDSW the standard coarse space is relatively large, especially in three dimensions. In Dohrmann and Widlund (2010), a related hybrid preconditioner with a reduced coarse problem for three-dimensional elasticity was introduced. Here, the degrees of freedom (d.o.f.) corresponding to the faces are modified.

The GDSW preconditioner is a two-level additive overlapping Schwarz preconditioner with exact local solvers; cf. Toselli and Widlund (2005). It can be written as

$$M_{\mathrm{GDSW}}^{-1} = \Phi \left(\Phi^T A \Phi\right)^{-1} \Phi^T + \sum_{i=1}^{N} R_i^T \tilde{A}_i^{-1} R_i, \tag{1}$$

cf. Dohrmann et al. (2008b). The matrix Φ is the essential ingredient of the GDSW preconditioner. It is composed of coarse space functions which are discrete

A. Heinlein (✉) • A. Klawonn
Mathematisches Institut, Universität zu Köln, Weyertal 86-90, 50931 Köln, Germany
e-mail: alexander.heinlein@uni-koeln.de; axel.klawonn@uni-koeln.de

O. Rheinbach
Institut für Numerische Mathematik und Optimierung, Fakultät für Mathematik und Informatik, Technische Universität Bergakademie Freiberg, Akademiestr. 6, 09596 Freiberg, Germany
e-mail: oliver.rheinbach@math.tu-freiberg.de

C.-O. Lee et al. (eds.), *Domain Decomposition Methods in Science and Engineering XXIII*, Lecture Notes in Computational Science and Engineering 116, DOI 10.1007/978-3-319-52389-7_36

harmonic extensions from the interface to the interior degrees of freedom of nonoverlapping subdomains. The values on the interface are restrictions of the nullspaces of the operator to the interface.

For $\Omega \subset \mathbb{R}^2$ being decomposed into John domains, the condition number of the GDSW preconditioner is bounded by

$$\kappa\left(M_{GDSW}^{-1}K\right) \leq C\left(1+\frac{H}{\delta}\right)\left(1+\log\left(\frac{H}{h}\right)\right)^2, \tag{2}$$

cf. Dohrmann et al. (2008a) and Dohrmann et al. (2008b). Here, H is the size of a subdomain, h is the size of a finite element, and δ is the overlap.

Implementation Our parallel implementation of the GDSW preconditioner is based on Trilinos version 12.0; cf. Heroux et al. (2005). For the mesh partitioning, we use ParMETIS, cf. Karypis et al. (2011), the problems corresponding to the local level are solved using UMFPACK, cf. Davis and Duff (1997) (version 5.3.0), and the coarse level is solved using Mumps, cf. Amestoy et al. (2001) (version 4.10.0), in parallel mode. For the finite element implementation, we use the library LifeV; see Formaggia et al. (2016) (version 3.8.8).

On the JUQUEEN BG/Q supercomputer, we use the clang compiler 4.7.2 and ESSL 5.1 when compiling Trilinos and the GDSW preconditioner implementation. On the Cray XT6m at Universität Duisburg-Essen, we use the Intel compiler 11.1 and the Cray Scientific Library (libsci) 10.4.4.

2 Model Problems

We consider model problems in two and three dimensions, i.e. $\Omega = [0,1]^2$ or $\Omega = [0,1]^3$. The domain is decomposed either in a structured way, i.e., into squares or cubes, or in an unstructured way, using ParMETIS.

Laplacian in 2D The first model problem is: find $u \in H^1(\Omega)$

$$\begin{aligned} -\Delta u &= 1 \qquad \text{in } \Omega, \\ u &= 0 \qquad \text{on } \partial\Omega. \end{aligned} \tag{3}$$

Linear Elasticity in 2D and 3D The second model problem is: find $u \in (H^1(\Omega))^2$;

$$\begin{aligned} \operatorname{div} \boldsymbol{\sigma} &= \boldsymbol{f} \qquad \text{in } \Omega, \\ \mathbf{u} &= 0 \qquad \text{on } \partial\Omega_D = \partial\Omega \cap \{x=0\} \end{aligned} \tag{4}$$

where $\boldsymbol{\sigma} = 2\mu\boldsymbol{\varepsilon} + \lambda \text{trace}(\boldsymbol{\varepsilon})I$ is the stress and $\boldsymbol{\varepsilon} = \frac{1}{2}(\nabla\boldsymbol{u} + (\nabla\boldsymbol{u})^T)$ the strain. The Lamé parameters are $\lambda = 1/2.6$ and $\mu = 0.3/0.52$.

3 Numerical Results

We first show parallel scalability results in two and three dimensions. Finally, we show an application of the preconditioner within a block preconditioner in monolithic fluid-structure interaction. The model problems are discretized using piecewise quadratic (P2) finite elements. Our default Krylov method is GMRES and will be used also for the symmetric positive definite model problems. Our stopping criterion is the relative criterion $\|r^{(k)}\|_2 / \|r^{(0)}\|_2 \le 10^{-7}$ with $r^{(0)}$ and $r^{(k)}$ being the initial and the k-th residual, respectively. In our experiments, each subdomain is assigned to one processor core.

Weak Scalability in 2D We use five different meshes with $H/h = 100$ and an increasing number of subdomains; see Tables 1 and 2. The results of weak scaling tests from 4 to 1024 processor cores for both model problems and an overlap $\delta = 1h$ or $\delta = 2h$ are presented in Figs. 1 and 2. The GDSW preconditioner

Table 1 Number of degrees of freedom of the total mesh, coarse and local space dimensions of the GDSW preconditioner for the weak scaling tests in Fig. 1

# Subdomains	4	16	64	256	1024
Total problem, P2 finite elements	160,801	641,601	2,563,201	10,246,401	40,972,801
Avg. first level, P2, overlap 1h	41,207.5	41,612.6	41,815.7	41,917.3	41,968.1
Avg. first level, P2, overlap 2h	42,020	42,837.8	43,248.7	43,454.7	43,557.8
Coarse level	5	33	161	705	2945
Avg. first level, P2, overlap 1h (ParMETIS)	41,581.5	41,841.9	42,101.8	42,225.7	42,263.1
Avg. first level, P2, overlap 2h (ParMETIS)	42,686.5	43,243.7	43,752.9	43,999.4	44,077.9
Coarse level (ParMETIS)	3	45	241	1129	4822

Table 2 Number of degrees of freedom of the total mesh, coarse and local space dimensions of the GDSW preconditioner for the weak scaling tests in Figs. 2 and 3

# Subdomains	4	16	64	256	1024
Total problem, P2	321,602	1,286,408	5,126,402	20,492,802	81,945,602
Avg. first level, P2, overlap 1h	82,415	83,225.2	83,631.3	83,834.6	83,936.3
Avg. first level, P2, overlap 2h	84,040	85,675.5	86,497.4	86,909.3	87,115.6
Coarse level	14	90	434	1890	7874
Coarse level, no rotations	10	66	322	1410	5890
Avg. first level, P2, overlap 1h (ParMETIS)	83,163	83,683.9	84,203.6	84,451.3	84,526.2
Avg. first level, P2, overlap 2h (ParMETIS)	85,373	86,487.4	87,505.8	87,998.7	88,155.9
Coarse level (ParMETIS)	9	120	633	2950	12,567
Coarse level, no rotations (ParMETIS)	6	90	482	2258	9644

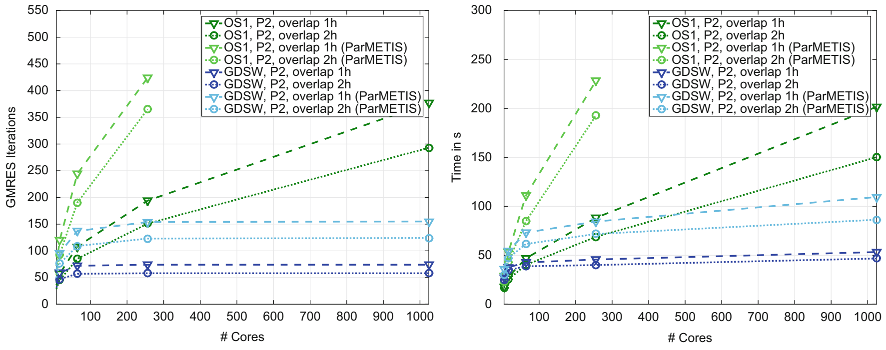

Fig. 1 Weak scaling for the Laplacian model problem in 2D, cf. (3), using P2 finite elements: number of iterations (*left*), runtimes (*right*). For the structured and the unstructured decomposition (ParMETIS), we have approximately 40,000 d.o.f. per subdomain

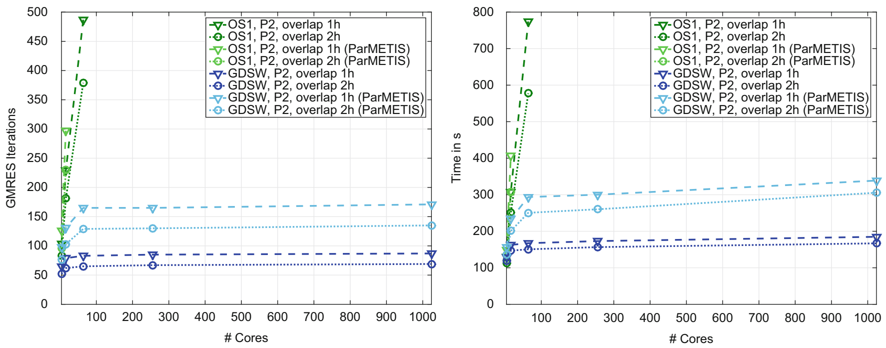

Fig. 2 Weak scaling for the linear elastic model problem in 2D, cf. (4), using P2 finite elements: number of iterations (*left*), runtimes (*right*). For the structured and the unstructured decomposition (ParMETIS), we have approximately 80,000 d.o.f. per subdomain

is numerically and parallel scalable, i.e., the number of iterations is bounded, both, for structured and unstructured decompositions, and the time to solution grows only slowly. The one-level preconditioner (OS1) does not scale numerically, and the number of iterations grows very fast. Indeed, for the unstructured decomposition, no convergence is obtained for OS1 within 500 iterations for more than 256 subdomains for the scalar problem and for more that 16 subdomains for elasticity. This is, of course, also due to the comparably small overlap. As a result of the better constant in (2), for the GDSW preconditioner, we observe better convergence for structured decompositions. Note that for the case of four subdomains the overlapping subdomains are significantly smaller.

A detailed analysis of different phases of the method is presented for linear elasticity in 2D in Fig. 3. We consider the standard full GDSW coarse space and the GDSW coarse space without rotations, i.e., the rotations are omitted from the

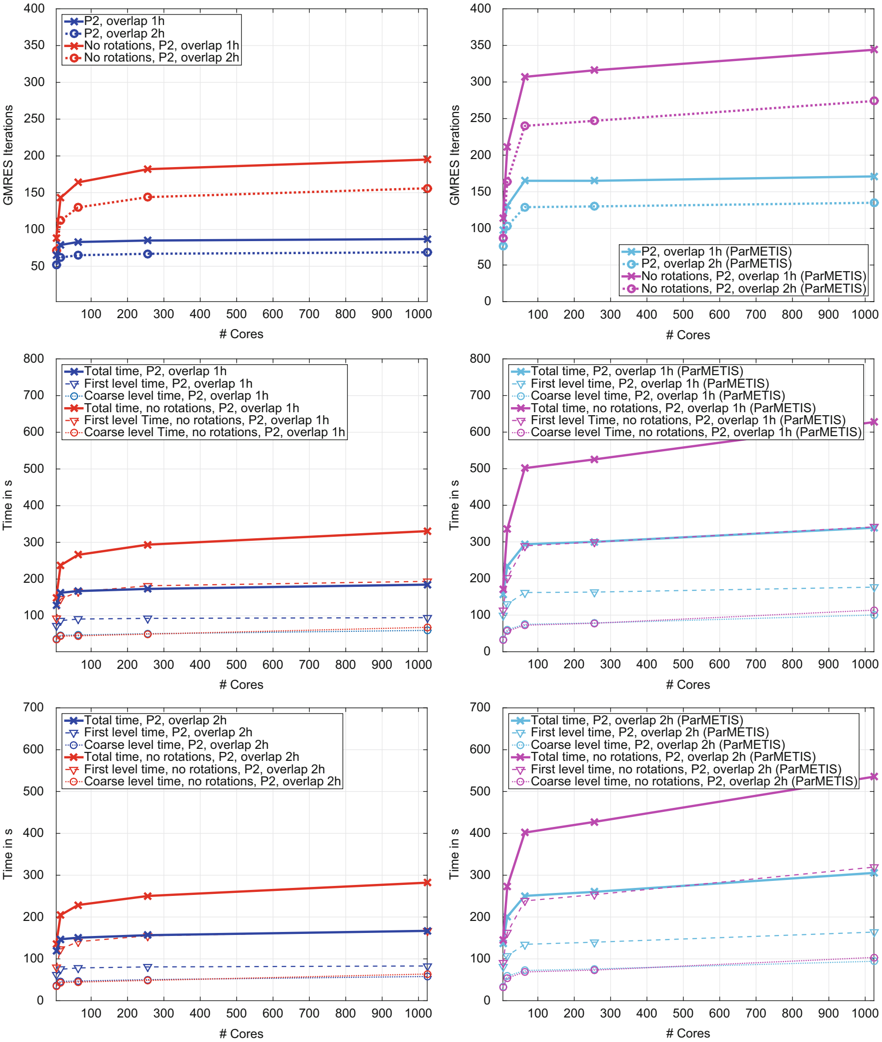

Fig. 3 Weak parallel scalability using the GDSW preconditioner for the model problem of linear elasticity in 2D, cf. (4): structured (*left*) and unstructured decomposition (*right*); number of iterations (*top*), timings for overlap $\delta = 1\,h$ (*middle*), and timings for overlap $\delta = 2\,h$ (*bottom*). For the structured and the unstructured decomposition (ParMETIS) we use a subdomain size of roughly 40,000 degrees of freedom

coarse space. This latter case is not covered by the bound (2), but the results indicate numerical and parallel scalability.

Strong Scalability in 2D Results for strong parallel scaling tests are shown in Fig. 4 for linear elasticity in 2D. We observe very good strong scalability for structured and unstructured domain decompositions. Note that the number of d.o.f.

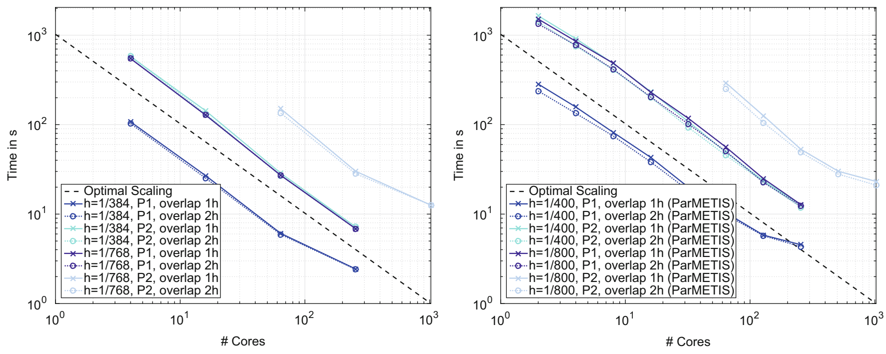

Fig. 4 Strong parallel scalability using the GDSW preconditioner for the model problem of linear elasticity in 2D, cf. (4): structured decomposition (*left*), ParMETIS decomposition (*right*)

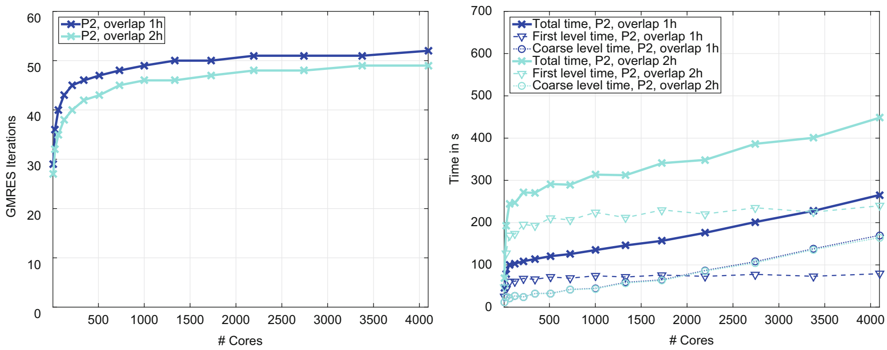

Fig. 5 Weak parallel scalability using the GDSW preconditioner for the problem of linear elasticity in 3D: number of iterations (*left*), timings (*right*). We use a subdomain size of $H/h = 6$ and P2 finite elements

per subdomain decreases when increasing the number of processor cores, and, to a certain extent, we thus benefit from an increasing speed of the local sparse direct solvers.

Weak Scalability for Linear Elasticity in 3D We present results of weak scalability runs for a linear elastic model problem in 3D from 8 to 4096 cores. We consider a structured decomposition of a cube and use the full GDSW coarse space in 3D. In Fig. 5, we present the number of iterations and the timings using P2 elements using an overlap δ of one or two elements. The number of iterations seems to be bounded by a constant number, whereas the solution times increases, i.e., the cost of the (parallel) sparse direct solver used for the coarse problem is noticeable in 3D.

Application in Fluid-Structure Interaction (FSI) We consider time-dependent monolithic FSI as in Balzani et al. (2015) but using a fully implicit scheme as in Deparis et al. (2015) and Heinlein et al. (2015). We apply a monolithic

Dirichlet-Neumann preconditioner applying the GDSW preconditioner for the structural block; see Balzani et al. (2015) and Heinlein et al. (2015) and the references therein. We use a pressure wave inflow condition for a tube using Mesh #1 from Heinlein et al. (2015). We consider a Neo-Hookean material for the tube; as opposed to Heinlein et al. (2015), we here use a fixed time step of 0.0005 s and show the runtimes during the simulation.

In Fig. 6, the runtimes of ten time steps using 128 cores of the Cray XT6m at Universität Duisburg-Essen are shown. We compare IFPACK, a one-level algebraic overlapping Schwarz preconditioner from Trilinos, our geometric one-level Schwarz preconditioner (OS1), the GDSW preconditioner without rotations (GDSW-nr), and the standard GDSW preconditioner for the structural block. We see that, although the computing times vary over the simulation time, the combination of the geometric overlap and a sufficiently large coarse space consistently reduces the runtime of the fully coupled monolithic FSI simulation by a factor of about two compared to the baseline given by IFPACK. Figure 7 shows the pressure and the deformation at $t = 0.007$ s where we have the largest computation time per timestep, cf. Fig. 6.

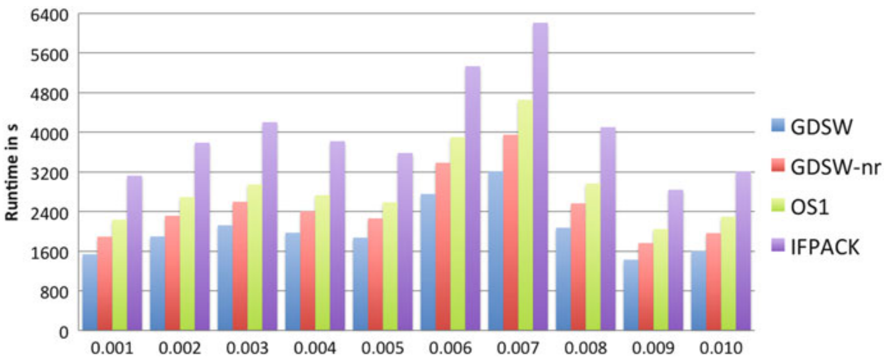

Fig. 6 Runtimes for the monolithic FSI simulation. For clarity, the runtimes of two subsequent time steps of size $\Delta t = 0.0005$ s are combined. The monolithic system has approximately 1.2 million d.o.f. We use a Neo-Hookean material. "OS1" is the one-level Schwarz preconditioner, "GDSW-nr" is the GDSW preconditioner without rotations, and "GDSW" is the GDSW preconditioner with full coarse space

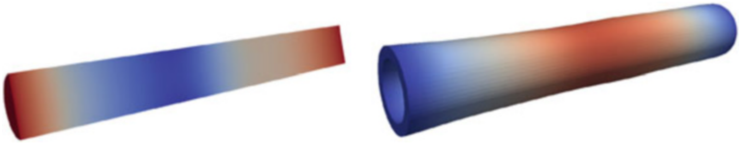

Fig. 7 Pressure and deformation at time $t = 0.007$ s. The deformation is magnified by a factor of 10

Acknowledgements The authors acknowledge the use of the JUQUEEN BG/Q supercomputer (Stephan and Docter, 2015) at JSC Jülich, the use of the Cray XT6m at Universität Duisburg-Essen and the financial support by the German Science Foundation (DFG), project no. KL2094/3 and RH122/4.

References

P.R. Amestoy, I.S. Duff, J.-Y. L'Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling. SIAM J. Matrix Anal. Appl. **23**(1), 15–41 (2001)

D. Balzani, S. Deparis, S. Fausten, D. Forti, A. Heinlein, A. Klawonn, A. Quarteroni, O. Rheinbach, J. Schröder, Numerical modeling of fluid-structure interaction in arteries with anisotropic polyconvex hyperelastic and anisotropic viscoelastic material models at finite strains. Int. J. Numer. Methods Biomed. Eng. (2015). ISSN 2040-7947. http://dx.doi.org/10.1002/cnm.2756.

T.A. Davis, I.S. Duff, An unsymmetric-pattern multifrontal method for sparse LU factorization. SIAM J. Matrix Anal. Appl. **18**(1), 140–158 (1997)

S. Deparis, D. Forti, G. Grandperrin, A. Quarteroni, FaCSI: a block parallel preconditioner for fluid-structure interaction in hemodynamics, Technical Report 13, MATHICSE, EPFL, Lausanne, 2015

C.R. Dohrmann, O.B. Widlund, Hybrid domain decomposition algorithms for compressible and almost incompressible elasticity. Int. J. Numer. Methods Eng. **82**(2), 157–183 (2010)

C.R. Dohrmann, A. Klawonn, O.B. Widlund, Domain decomposition for less regular subdomains: overlapping Schwarz in two dimensions. SIAM J. Numer. Anal. **46**(4), 2153–2168 (2008a). ISSN 0036-1429

C.R. Dohrmann, A. Klawonn, O.B. Widlund, A family of energy minimizing coarse spaces for overlapping Schwarz preconditioners, in *Domain Decomposition Methods in Science and Engineering XVII*. Lecture Notes in Computational Science and Engineering, vol. 60 (Springer, Berlin, 2008b), pp. 247–254

L. Formaggia, M. Fernandez, A. Gauthier, J.F. Gerbeau, C. Prud'homme, A. Veneziani, The LifeV Project. Web. http://www.lifev.org (2016)

A. Heinlein, A. Klawonn, O. Rheinbach, Parallel two-level overlapping Schwarz methods in fluid-structure interaction, in *Proceedings of the European Conference on Numerical Mathematics and Advanced Applications* (*ENUMATH*), Ankara, September, 2015. Springer Lecture Notes on Computational Science and Engineering, vol. 112 (2016), pp. 521–530. TUBAF Preprint 15/2015: http://tu-freiberg.de/fakult1/forschung/preprints

M.A. Heroux, R.A. Bartlett, V.E. Howle, R.J. Hoekstra, J.J. Hu, T.G. Kolda, R.B. Lehoucq, K.R. Long, R.P. Pawlowski, E.T. Phipps, A.G. Salinger, H.K. Thornquist, R.S. Tuminaro, J.M. Willenbring, A. Williams, K.S. Stanley, An overview of the Trilinos project. ACM Trans. Math. Softw. **31**(3), 397–423 (2005)

G. Karypis, K. Schloegel, V. Kumar, ParMETIS - Parallel graph partitioning and sparse matrix ordering. Version 3.2, Technical Report, University of Minnesota, Department of Computer Science and Engineering, April 2011

M. Stephan, J. Docter, JUQUEEN: IBM Blue Gene/Q® Supercomputer System at the Jülich Supercomputing Centre. J. Large-Scale Res. Facil. **1**, A1 (2015). ISSN 2364-091X. doi:10.17815/jlsrf-1-18. http://dx.doi.org/10.17815/jlsrf-1-18

A. Toselli, O. Widlund, *Domain Decomposition Methods—Algorithms and Theory*. Springer Series in Computational Mathematics, vol. 34 (Springer, Berlin, 2005). ISBN 3-540-20696-5