

# FROSch: A Fast And Robust Overlapping Schwarz Domain Decomposition Preconditioner Based on Xpetra in Trilinos

Alexander Heinlein, Axel Klawonn, Sivasankaran Rajamanickam, and Oliver Rheinbach

## 1 Introduction

This article describes a parallel implementation of a two-level overlapping Schwarz preconditioner with the GDSW (Generalized Dryja–Smith–Widlund) coarse space described in previous work [12, 10, 15] into the Trilinos framework; cf. [16]. The software is a significant improvement of a previous implementation [12]; see Sec. 4 for results on the improved performance.

In the software, now named FROSch (Fast and Robust Overlapping Schwarz), efforts were made for the seamless integration into the open-source Trilinos framework, and to allow the use of heterogeneous architectures, such as those with NVIDIA accelerators. These goals were achieved in the following way:

1. The GDSW preconditioner, i.e., the FROSch library, is now part of Trilinos as a subpackage of the package ShyLU. The ShyLU package provides distributed-memory parallel domain decomposition solvers, and node-level direct solvers for the subdomains. Currently, ShyLU has two other domain decomposition solvers, i.e., a Schur complement solver [19] and an implementation of the BDDC method by Clark Dohrmann, and node-level (in)complete LU factorizations (`basker` [2]), `fastilu` [18]), Cholesky factorization (`tacho` [17]) and triangular solves (`hts` [3]).

---

Alexander Heinlein and Axel Klawonn

Department of Mathematics and Computer Science, University of Cologne, Weyertal 86-90, 50931 Köln, Germany. E-mail: {alexander.heinlein, axel.klawonn}@uni-koeln.de.  
Center for Data and Simulation Science, University of Cologne, Germany, url: <http://www.cds.uni-koeln.de>

Sivasankaran Rajamanickam

Center for Computing Research, Scalable Algorithms Department, Sandia National Laboratories, Albuquerque, NM 87123. e-mail: [srajama@sandia.gov](mailto:srajama@sandia.gov)

Oliver Rheinbach

INMO, Technische Universität Bergakademie Freiberg, Akademiestr. 6, 09599 Freiberg, Germany. E-mail: [oliver.rheinbach@math.tu-freiberg.de](mailto:oliver.rheinbach@math.tu-freiberg.de).

2. FROSch now supports the Kokkos programming model through the use of the Tpetra stack in Trilinos. The FROSch library can therefore profit from the efforts of the Kokkos package to obtain performance portability by template meta-programming, on modern hybrid architectures with accelerators. During this process the GDSW code has been modified and improved significantly. The resulting FROSch library is now designed such that different types of Schwarz operators can be added and combined more easily. Consequently, various different Schwarz preconditioners can be constructed using the FROSch framework. Recently, FROSch has been used in a three-level GDSW implementation [13, 14] and for the solution of incompressible fluid flow problems [11].

## 2 The GDSW Preconditioner

We are concerned with finding the solution of a sparse linear system

$$Ax = b, \quad (1)$$

arising from a finite element discretization with finite element space  $V = V^h(\Omega)$  of an elliptic problem, such as, a Laplace problem, on a domain  $\Omega \subset \mathbb{R}^d$ ,  $d = 2, 3$ , with sufficient Dirichlet boundary conditions. The GDSW preconditioner [4, 5] is a two-level additive overlapping Schwarz preconditioner with exact local solvers (cf. [21]) using a coarse space constructed from energy-minimizing functions. It is meant to be used in combination with the Krylov methods from the packages Belos [1] or AztecOO. In particular, let  $\Omega$  be decomposed into  $N$  nonoverlapping subdomains  $\Omega_i$ ,  $i = 1, \dots, N$ , and overlapping subdomains  $\Omega'_i$ ,  $i = 1, \dots, N$ , respectively, and  $V_i = V^h(\Omega'_i)$ ,  $i = 1, \dots, N$ , be the corresponding local finite element spaces. Further, we define standard restriction operators  $R_i : V \rightarrow V_i$ ,  $i = 1, \dots, N$ , from the global to the local finite element spaces. Then, the Schwarz operator of the GDSW method can be written in the form

$$P_{\text{GDSW}} = M_{\text{GDSW}}^{-1}A = \Phi A_0^{-1} \Phi^T A + \sum_{i=1}^N R_i^T A_i^{-1} R_i A, \quad (2)$$

where  $A_0 = \Phi^T A \Phi$  is the coarse space matrix, and the matrices  $A_i = R_i A R_i^T$ ,  $i = 1, \dots, N$ , represent the overlapping local problems; cf. [5]. The matrix  $\Phi$  is the essential ingredient of the GDSW preconditioner. It is composed of coarse space functions which are discrete harmonic extensions from the interface to the interior degrees of freedom of nonoverlapping subdomains. The values on the interface are typically chosen as restrictions of the elements of the null space of the operator  $\hat{A}$  to the edges, vertices, and faces of the decomposition, where  $\hat{A}$  is the global matrix corresponding to  $A$  but with homogeneous Neumann boundary condition. Therefore, for a scalar elliptic problem, the coarse basis functions form a partition of unity on

all subdomains that do not touch the Dirichlet boundary. The condition number of the GDSW Schwarz operator is bounded as

$$\kappa(P_{\text{GDSW}}) \leq C \left(1 + \frac{H}{\delta}\right) \left(1 + \log\left(\frac{H}{h}\right)\right)^2, \quad (3)$$

where  $h$  is the size of a finite element,  $H$  the size of a nonoverlapping subdomain, and  $\delta$  the width of the overlap; see [4, 5, 6]. The exponent of the logarithmic term can be reduced to 1 for variants of the GDSW coarse space; see, e.g., [7, 8].

However, the dimension of the standard GDSW coarse space is in the order of  $\dim(V_0) = \mathcal{O}(\dim(\text{null}(\hat{A}))(N_{\mathcal{V}} + N_{\mathcal{E}} + N_{\mathcal{F}}))$ , where  $N_{\mathcal{V}}$ ,  $N_{\mathcal{E}}$ , and  $N_{\mathcal{F}}$  are the global numbers of vertices, edges, and faces of the nonoverlapping domain decomposition, respectively. The dimension of the coarse space is fairly high. Therefore, GDSW coarse spaces of reduced dimension have very recently been introduced in [8]; see also [15] for parallel results. For general problems, the dimension of the reduced GDSW coarse spaces is  $\dim(V_0) = \mathcal{O}(\dim(\text{null}(\hat{A}))(N_{\mathcal{V}}))$ , which is, especially for unstructured decompositions, significantly smaller. Both types of GDSW coarse spaces are implemented in FROSch, and in Sec. 4, we present performance results.

### 3 Software Design of the FROSch Library


During the integration of the FROSch library into Trilinos, the code was substantially restructured. In particular, in the transition from the Trilinos Epetra (used in [12]) to the newer Xpetra sparse matrix infrastructure, it was extended to a framework of Schwarz preconditioners. Additionally, parts of the code have been improved and functionality has been added. As opposed to [12], FROSch is completely based on Xpetra.

**A Framework for Schwarz Preconditioners** As described in Sec. 2, the GDSW preconditioner is a two-level overlapping Schwarz method using a specific coarse space. The GDSW Schwarz operator is of the form

$$P_{2\text{-Lvl}} = \underbrace{\Phi A_0^{-1} \Phi^T A}_{P_0} + \sum_{i=1}^N \underbrace{R_i^T A_i^{-1} R_i A}_{P_i};$$

cf. (2); and therefore, it is the sum of local overlapping Schwarz operators  $P_i$ ,  $i = 1, \dots, N$ , and a global coarse Schwarz operator  $P_0$ . There are different ways to compose Schwarz operators  $P_i$ ,  $i = 0, \dots, N$ , e.g.:

$$\begin{aligned} \textbf{Additive:} \quad P_{\text{ad}} &= \sum_{i=0}^N P_i \\ \textbf{Multiplicative:} \quad P_{\text{mu}} &= I - (I - P_N)(I - P_{N-1}) \cdots (I - P_0) \\ P_{\text{mu-sym}} &= I - \prod_{i=0}^N (I - P_i) \prod_{i=0}^{N-1} (I - P_{N-1-i}) \\ \textbf{Hybrid:} \quad P_{\text{hy-1}} &= I - (I - P_0) \left( I - \sum_{i=0}^N P_i \right) (I - P_0) \\ P_{\text{hy-2}} &= \alpha P_0 + I - (I - P_N) \cdots (I - P_1); \end{aligned}$$



**Fig. 1:** Heuristic reconstruction of the domain decomposition interface: uniquely distributed map (left); extension of the uniquely distributed map by one layer of elements resulting in an overlapping map, where the overlap contains the interface (middle); by selection, using the lower subdomain ID, the interface is defined (right).

cf. [21]. Using the FROSch library, it is simple to construct the different variants once the ingredients are set up. Let us explain this based on the example of the class `GDSWPreconditioner` in FROSch, which is derived from the abstract class `SchwarzPreconditioner` and contains an implementation of the GDSW preconditioner: in FROSch, the `SumOperator` is used to combine Schwarz operators in an additive way. The additive first level is implemented in the class `AlgebraicOverlappingOperator` and the coarse level of the GDSW preconditioner in the class `GDSWCoarseOperator`. Therefore, the `GDSWPreconditioner` is basically just the following composition of Schwarz operators:

```
GDSWPreconditioner = SumOperator( AlgebraicOverlappingOperator,
                                  GDSWCoarseOperator )
```

By replacing the `SumOperator` by a `ProductOperator`, the levels can be coupled in a multiplicative way. The different classes for Schwarz operators are all derived from an abstract `SchwarzOperator`, and the classes `SchwarzOperator` and `SchwarzPreconditioner` are both derived from the abstract `Xpetra::Operator`. **Transition from Epetra to Xpetra** To facilitate the use of FROSch on novel architectures, the code was ported completely from Epetra data structures to Xpetra. As Xpetra provides a lightweight interface to Epetra as well as Tpetra, FROSch can now profit from the computational kernels from Kokkos, while maintaining compatibility to older Epetra-based software such as LifeV [9].

**Improvement of the Code & Additional Functionality** The efficiency of the code was improved and new functionality was added as part of this redesign. In particular, the routines for the computation of local-to-global mappings and the identification of the interface components have been rewritten and therefore improved with respect to their performance; see Sec. 4 for the numerical results.

Two important features have been added. First, we have introduced the possibility to reconstruct a domain decomposition interface algebraically based on a unique distribution of the degrees of freedom into subdomains and the nonzero pattern of the matrix; cf. Fig. 1. This works particularly well for scalar elliptic problems and piecewise linear elements. In general, the best performance is obtained when a `RepeatedMap` is provided by the user; cf. Fig. 2. This map corresponds to the

**Previous implementation from [12]:**

```

Teuchos::RCP<SOS::SOS> M_SOS(new SOS(numVectors, numSubdomainsPerProcess,
    M_DomainMap, M_RangeMap));
Teuchos::RCP<SOS::SOSSetup> M_SOSSetup(new SOS::SOSSetup(
    numSubdomainsPerProcess, dimension, dofs, M_rowMatrixTeuchos, M_DomainMap));
M_SOSSetup->FirstLevel(M_ProcessMapOverlap);
M_SOSSetup->SecondLevel(M_ProcessMapNodes, M_ProcessMap, SOS::LifeVOrdering,
    M_LocalDirichletBoundaryDofs, "Mumps", useRotations, M_LocalNodeList);
M_SOSSetup->SetUpPreconditioner(M_SOS, "Mumps", secondLevelSolverParameterList,
    Type);

```

**Current implementation Shylu/FROSch:**

```

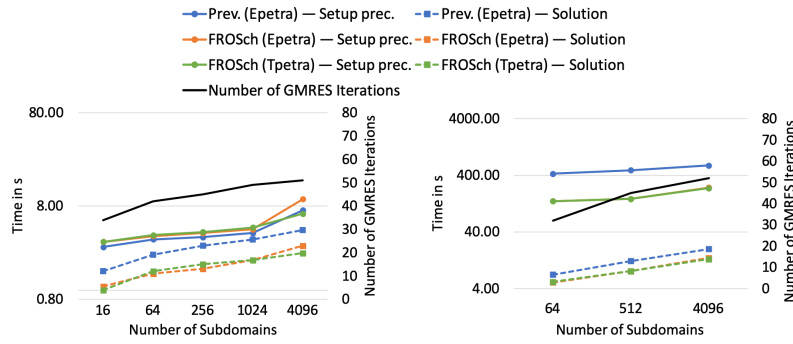
Teuchos::RCP<FROSch::GDSWPreconditioner<SC,LO,GO,NO> > FROSchGDSW(new FROSch::
    GDSWPreconditioner<SC,LO,GO,NO>(A, ParameterList);
FROSchGDSW->initialize(Dimension, Overlap, RepeatedMap);
FROSchGDSW->compute();

```


**Fig. 2:** Comparison of the user-interface for the previous implementation of the GDSW solver (top) and the current implementation in FROSch (bottom). The setup is split into the `initialize` and `compute` phases instead of the two levels.

nonoverlapping domain decomposition and is replicated in the interface degrees of freedom. Secondly, we have introduced a function that identifies Dirichlet boundary conditions based on matrix rows with only diagonal entries. This is important because the coarse basis functions are zero on the Dirichlet boundary.

**User Interface** The user-interface of the FROSch library has been completely re-designed. Compared to the previous implementation, where the setup of the preconditioner was split into the first and second level, it is now split into the phases `initialize` and `compute`, also reducing the number of required lines of code to construct the GDSW preconditioner; cf. Fig. 2. In the `initialize` phase, all data structures that correspond to the structure of the problem are built, i.e., the index sets of the overlapping subdomains and the interface are identified and the interface values of the coarse basis are computed. In the `compute` phase, all computations related to the values of the matrix  $A$  are performed, i.e., the overlapping problems are factorized, the interior values of the coarse basis are computed, and the coarse problem is assembled and factorized. Therefore, the `initialize` and `compute` phases can be seen as the symbolic and the numerical factorizations of a direct solver: if only the values in the matrix  $A$  change, the preconditioner can be updated using `compute`, and if the structure of the problem is changed, `initialize` has to be called to update the preconditioner. Also, FROSch provides a `Stratimikos` interface for easier use in applications; `Stratimikos` provides a unified framework for solvers and preconditioners in Trilinos.




**Fig. 3:** Weak scalability of the two-level Schwarz preconditioner using the GDSW coarse space for the Poisson model problem: (Left) in two dimensions with overlap  $\delta = 5h$  and  $H/h = 100$  (approximately 50k degrees of freedom per sudomain); (Right) in two dimensions with overlap  $\delta = 2h$  and  $H/h = 14$  (approximately 50k degrees of freedom per sudomain). Comparison of the previous implementation (blue) and the current implementation in FROSch, i.e., the Epetra (orange) and the Tpetra (green) versions available through the Xpetra interface. The number of iterations (black) are identical for all versions.



**Fig. 4:** Weak scalability of the two-level Schwarz preconditioner with overlap  $\delta = 1h$  for the Poisson model problem in three dimensions with  $H/h = 14$  (approximately 35k degrees of freedom per subdomain): comparison of the GDSW and the RGDSW coarse space using the Tpetra version of the FROSch implementation.

### 4 Performance of the New FROSch Software

Here, the performance of the new software is compared against the previous implementation. We consider a Poisson model problem on  $\Omega \subset \mathbb{R}^d$ ,  $d = 2, 3$ , with full Dirichlet boundary condition, discretized by piecewise quadratic finite elements. We compare the performance of the previous implementation, which is based on Epetra, and the current implementation in FROSch. In particular, the Epetra and the Tpetra version of the current implementation, which are both available through the Xpetra interface, are compared. As a Krylov-solver GMRES from Belos [1] is used with a relative tolerance of  $10^{-7}$  for the unpreconditioned residual. For the local and coarse problems, the native direct solver in Trilinos, KLU, is used; only in Fig. 5, Mumps is used as the direct solver. We always use one subdomain per processor core. The computations were performed on the magnitUDE supercomputer at Universität Duisburg-Essen, which has 15k cores (Intel Xeon E5-2650v4, 12C, 2.2GHz) and



**Fig. 5:** Weak scalability for the Poisson model problem in two dimensions with  $H/h = 200$  (approximately 195k degrees of freedom per subdomain): comparison of FROSch using the GDSW coarse space and the one-level overlapping Schwarz preconditioner Ifpack with overlap  $\delta = 20h$ ; numbers of GMRES iterations (left) and total solver times (right). Using Mumps for all direct solves. For 1 024 subdomains, Ifpack did not converge within 500 GMRES iterations.

a total memory of 36 096 GB. Here, we do not exploit any node parallelism when using Tpetra. We consider the setup phase and the solution phase and include the identification of the interface components in the setup phase. This part does not scale very well and can take a significant amount of time for a large number of processes; cf. [12]. In Fig. 3 (left), we present numerical results for the GDSW preconditioner in two dimensions. We observe that, in the solution phase, the new implementation is always faster than the previous implementation. The time for the setup phase is comparable. The results in Fig. 3 (right), where we compare the preconditioners in three dimensions, are more interesting. Again, we observe that the solution phase is faster by a similar factor. However, in three dimensions, the setup phase in the FROSch implementation is much faster compared to the previous implementation. We also observe that the Tpetra version is always slightly faster than the Epetra version of the new code. In Fig. 4, the GDSW and the RGDSW coarse spaces are compared for the Tpetra version of the FROSch implementation. We observe that, due to the increasing dimension of the coarse space, the computation time can be improved when using reduced dimension coarse spaces. This effect becomes stronger when the number of subdomains is increased; cf. [15]. Finally, we present a comparison of FROSch using the GDSW coarse space and Ifpack [20], i.e., a one-level overlapping Schwarz preconditioner, in Fig. 5. We observe that Ifpack does not scale as it lacks a second level. Already for 64 subdomains, FROSch converges much faster, and for 1 024 subdomains, Ifpack does not converge within a maximum number of 500 GMRES iterations.

**Conclusion** We presented the new Trilinos library FROSch that allows the flexible construction of different overlapping Schwarz methods. The FROSch implementation of the GDSW preconditioner is significantly faster than the previous one.

**Acknowledgements** The authors gratefully acknowledge the computing time granted by the Center for Computational Sciences and Simulation (CCSS) at Universität Duisburg-Essen and provided on the supercomputer magnitUDE (DFG grants INST 20876/209-1 FUGG, INST 20876/243-1 FUGG) at Zentrum für Informations- und Mediendienste (ZIM). Sandia National Laboratories is a multitechnology laboratory managed and operated by National Technology and Engineering Solutions

of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

## References

1. Bavier, E., Hoemmen, M., Rajamanickam, S., Thornquist, H.: Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems. *Scientific Programming* **20**(3), 241–255 (2012)
2. Booth, J.D., Ellingwood, N.D., Thornquist, H.K., Rajamanickam, S.: Basker: Parallel sparse lu factorization utilizing hierarchical parallelism and data layouts. *Parallel Computing* **68**, 17–31 (2017). DOI:<https://doi.org/10.1016/j.parco.2017.06.003>. URL <http://www.sciencedirect.com/science/article/pii/S0167819117300868>
3. Bradley, A.M.: A hybrid multithreaded direct sparse triangular solver. In: *Proceedings of CSC16*, pp. 13–22. SIAM (2016)
4. Dohrmann, C.R., Klawonn, A., Widlund, O.B.: Domain decomposition for less regular subdomains: overlapping Schwarz in two dimensions. *SIAM J. Numer. Anal.* **46**(4), 2153–2168 (2008)
5. Dohrmann, C.R., Klawonn, A., Widlund, O.B.: A family of energy minimizing coarse spaces for overlapping Schwarz preconditioners. In: *Domain decomposition methods in science and engineering XVII, LNCSE*, vol. 60, pp. 247–254. Springer, Berlin (2008)
6. Dohrmann, C.R., Widlund, O.B.: Hybrid domain decomposition algorithms for compressible and almost incompressible elasticity. *IJNME* **82**(2), 157–183 (2010)
7. Dohrmann, C.R., Widlund, O.B.: An alternative coarse space for irregular subdomains and an overlapping Schwarz algorithm for scalar elliptic problems in the plane. *SIAM J. Numer. Anal.* **50**(5), 2522–2537 (2012). DOI: [10.1137/110853959](https://doi.org/10.1137/110853959)
8. Dohrmann, C.R., Widlund, O.B.: On the design of small coarse spaces for domain decomposition algorithms. *SIAM J. Sci. Comput.* **39**(4), A1466–A1488 (2017). DOI: [10.1137/17M1114272](https://doi.org/10.1137/17M1114272)
9. Formaggia, L., Fernandez, M., Gauthier, A., Gerbeau, J.F., Prud'homme, C., Veneziani, A.: The LifeV Project. Web. [Http://www.lifev.org](http://www.lifev.org)
10. Heinlein, A.: Parallel overlapping Schwarz preconditioners and multiscale discretizations with applications to fluid-structure interaction and highly heterogeneous problems. PhD thesis, Universität zu Köln, Germany (2016)
11. Heinlein, A., Hochmuth, C., Klawonn, A.: Monolithic overlapping Schwarz domain decomposition methods with GDSW coarse spaces for incompressible fluid flow problems. *SIAM Journal on Scientific Computing* **41**(4), C291–C316 (2019). DOI: [10.1137/18M1184047](https://doi.org/10.1137/18M1184047). URL <https://doi.org/10.1137/18M1184047>
12. Heinlein, A., Klawonn, A., Rheinbach, O.: A parallel implementation of a two-level overlapping Schwarz method with energy-minimizing coarse space based on Trilinos. *SIAM J. Sci. Comput.* **38**(6), C713–C747 (2016). DOI: [10.1137/16M1062843](https://doi.org/10.1137/16M1062843)
13. Heinlein, A., Klawonn, A., Rheinbach, O., Röver, F.: A Three-Level Extension of the GDSW Overlapping Schwarz Preconditioner in Two Dimensions, pp. 187–204. Springer International Publishing, Cham (2019). DOI: [10.1007/978-3-030-14244-5\\_10](https://doi.org/10.1007/978-3-030-14244-5_10). URL [https://doi.org/10.1007/978-3-030-14244-5\\_10](https://doi.org/10.1007/978-3-030-14244-5_10)
14. Heinlein, A., Klawonn, A., Rheinbach, O., Röver, F.: A three-level extension of the GDSW overlapping Schwarz preconditioner in three dimensions. Tech. rep. (March 2019). Accepted for publication to LNCSE.
15. Heinlein, A., Klawonn, A., Rheinbach, O., Widlund, O.B.: Improving the parallel performance of overlapping Schwarz methods by using a smaller energy minimizing coarse space. In: *Domain Decomposition Methods in Science and Engineering XXIV*, pp. 383–392. Springer International Publishing, Cham (2018)
16. Heroux, M.A., Bartlett, R.A., Howle, V.E., Hoekstra, R.J., Hu, J.J., Kolda, T.G., Lehoucq, R.B., Long, K.R., Pawlowski, R.P., Phipps, E.T., Salinger, A.G., Thornquist, H.K., Tuminaro,



- R.S., Willenbring, J.M., Williams, A., Stanley, K.S.: An overview of the Trilinos project. *ACM Trans. Math. Softw.* **31**(3), 397–423 (2005). DOI:<http://doi.acm.org/10.1145/1089014.1089021>
17. Kim, K., Edwards, H.C., Rajamanickam, S.: Tacho: Memory-scalable task parallel sparse cholesky factorization. In: 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 550–559. IEEE (2018)
  18. Patel, A., Boman, E., Rajamanickam, S., Chow, E.: Cross platform fine grained ILU and ILDL factorizations using Kokkos. In: Proceedings of the CCR (2015)
  19. Rajamanickam, S., Boman, E.G., Heroux, M.A.: ShyLU: A hybrid-hybrid solver for multicore platforms. In: 2012 IEEE 26th International Parallel and Distributed Processing Symposium, pp. 631–643 (2012). DOI:[10.1109/IPDPS.2012.64](https://doi.org/10.1109/IPDPS.2012.64)
  20. Sala, M., Heroux, M.: Robust algebraic preconditioners with IFPACK 3.0. Tech. Rep. SAND-0662, Sandia National Laboratories (2005)
  21. Toselli, A., Widlund, O.: Domain decomposition methods—algorithms and theory, *Springer Series in Computational Mathematics*, vol. 34. Springer-Verlag, Berlin (2005)