Numerical Study of the Almost Nested Case in a Multilevel Method Based on Non-nested Meshes

Thomas Dickopf and Rolf Krause

University of Lugano, Institute of Computational Science,
Via G. Buffi 13, 6904 Lugano, Switzerland
thomas.dickopf@usi.ch, rolf.krause@usi.ch

Summary. Partial differential equations in complex domains are very flexibly discretized by finite elements with unstructured meshes. For such problems, the challenging task to construct coarse level spaces for efficient multilevel preconditioners can in many cases be solved by a semi-geometric approach, which is based on a hierarchy of non-nested meshes. In this paper, we investigate the connection between the resulting semi-geometric multigrid methods and the truly geometric variant more closely. This is done by considering a sufficiently simple computational domain and treating the geometric multigrid method as a special case in a family of almost nested settings. We study perturbations of the meshes and analyze how efficiency and robustness depend on a truncation of the interlevel transfer. This gives a precise idea of which results can be achieved in the general unstructured case.

1 Introduction

This paper is about multilevel methods for an efficient solution of partial differential equations in complicated domains. Our particular purpose is to provide additional insight into the design of coarse spaces in case of unstructured finite element meshes. We study an approach of semi-geometric preconditioning based on non-nested mesh hierarchies motivated by Cai [2], Chan et al. [3, 4], Griebel and Schweitzer [6], Toselli and Widlund [8], and Xu [9]. This is a concept with rather weak requirements (yet still in a variational setting) compared with other geometry-based methods. The main contribution of the present paper is a numerical study of the almost nested case, which establishes a connection between the multilevel methods based on non-nested meshes and the standard variant. Combined with our investigations of mesh perturbations, this allows for the determination of a suitable truncation parameter for the interlevel transfer. As a result, the efficiency of the completely nested case is in large part retained.

R. Bank et al. (eds.), *Domain Decomposition Methods in Science and Engineering XX*, Lecture Notes in Computational Science and Engineering 91,
DOI 10.1007/978-3-642-35275-1_65,

2 Multilevel Preconditioners Based on Non-nested Meshes

This section aims at a semi-geometric preconditioning framework. We introduce a multiplicative multilevel preconditioner based on a hierarchy of non-nested meshes. This is done in a way which allows for a powerful convergence analysis as well as an efficient implementation.

Let $\Omega \subset \mathbb{R}^d$ be a Lipschitz domain of dimension $d \in \{2,3\}$. For a right hand side $\mathscr{F} \in H^{-1}(\Omega)$ and a positive function $\alpha \in L^\infty(\Omega)$ bounded away from zero, we consider the variational model problem

$$u \in H_0^1(\Omega): \quad a(u,v) := (\alpha \nabla u, \nabla v)_{L^2(\Omega)} = \mathscr{F}(v), \quad \forall\, v \in H_0^1(\Omega). \tag{1}$$

For a Galerkin discretization of problem (1), let $(\mathscr{T}_\ell)_{\ell\in\mathbb{N}}$ be a family of *non-nested* shape regular meshes of domains $(\Omega_\ell)_{\ell\in\mathbb{N}}$. We denote the set of nodes of $\mathscr{T}_\ell$ by $\mathscr{N}_\ell$ and abbreviate $n_\ell := |\mathscr{N}_\ell|$. At each level ℓ, we consider the space X_ℓ of Lagrange conforming finite elements of first order and denote its nodal basis as $\Lambda_\ell = (\lambda_p^\ell)_{p\in\mathscr{N}_\ell}$ with $\lambda_p^\ell(q) = \delta_{pq}$, $p,q \in \mathscr{N}_\ell$. For simplicity, we assume that $\Omega_L = \Omega$ and $X_L \subset H_0^1(\Omega)$ for a fixed finest level $L \geq 2$. In addition, let $\Omega_\ell \supset \Omega$ for all $\ell \in \{0,\dots,L-1\}$. The basic idea how the setting can be chosen is exemplarily illustrated in Fig. 1 (left) for an unstructured fine mesh with structured coarse meshes.

In the following, we consider an iterative method to efficiently solve the discrete problem, namely the ill-conditioned equation

$$\boldsymbol{A}_L \boldsymbol{u}_L = \boldsymbol{F}_L \quad \text{in } \mathbb{R}^{n_L}.$$

Here, $\boldsymbol{A}_L \in \mathbb{R}^{n_L\times n_L}$ is the stiffness matrix associated with X_L, i.e., $(\boldsymbol{A}_L)_{pq} := a(\lambda_p^L, \lambda_q^L)$ for $p,q \in \mathscr{N}_L$, and the right hand side $\boldsymbol{F}_L \in \mathbb{R}^{n_L}$ is given by $(\boldsymbol{F}_L)_p := \mathscr{F}(\lambda_p^L)$ for $p \in \mathscr{N}_L$.

For the construction of an appropriate coarse space hierarchy, let the spaces $(X_\ell)_{\ell=0,\dots,L}$ be connected by the prolongation operators $(\Pi_{\ell-1}^\ell)_{\ell=1,\dots,L}$, namely

$$\Pi_{\ell-1}^\ell : X_{\ell-1} \to X_\ell, \quad \forall\, \ell \in \{1,\dots,L\}.$$

The choice of a concrete transfer concept generating a set of suitable linear operators $(\Pi_{\ell-1}^\ell)_{\ell=1,\dots,L}$ in practice is discussed in full detail in [5]. An example is nodal interpolation. Now, let $V_L := X_L$; we emphasize that the fine space will not be touched in the present framework. We construct a nested sequence of spaces $(V_\ell)_{\ell=0,\dots,L}$ via

$$V_\ell := \Pi_{L-1}^L \cdots \Pi_\ell^{\ell+1} X_\ell, \quad \forall\, \ell \in \{0,\dots,L-1\}.$$

The images of the compositions of the given operators determine the coarse spaces.

With the nodal bases $(\Lambda_\ell)_{\ell=0,\dots,L}$, matrix representations $\boldsymbol{\Pi}_{\ell-1}^\ell \in \mathbb{R}^{n_\ell\times n_{\ell-1}}$ of $\Pi_{\ell-1}^\ell$ can be computed for $\ell \in \{1,\dots,L\}$ via $\boldsymbol{\Pi}_{\ell-1}^\ell \boldsymbol{v} := \boldsymbol{\Phi}_\ell^{-1}(\Pi_{\ell-1}^\ell \boldsymbol{\Phi}_{\ell-1}(\boldsymbol{v}))$ for all $\boldsymbol{v} \in \mathbb{R}^{n_{\ell-1}}$ with the coordinate isomorphisms $\boldsymbol{\Phi}_\ell : \mathbb{R}^{n_\ell} \to X_\ell$. Assume that these matrices have full rank. Then, bases of $(V_\ell)_{\ell=0,\dots,L-1}$ can recursively be defined by

$$\widetilde{\lambda}_q^\ell := \sum_{p\in\mathcal{N}_{\ell+1}} (\boldsymbol{\Pi}_\ell^{\ell+1})_{pq}\widetilde{\lambda}_p^{\ell+1}, \quad \forall\, q\in\mathcal{N}_\ell,$$

starting with $\widetilde{\lambda}_q^L := \lambda_q^L$ for $q\in\mathcal{N}_L$. The new coordinate isomorphisms with respect to the bases $\widetilde{\Lambda}_\ell := (\widetilde{\lambda}_p^\ell)_{p\in\mathcal{N}_\ell}$, $\ell\in\{0,\dots,L\}$, will be denoted by $\widetilde{\boldsymbol{\Phi}}_\ell : \mathbb{R}^{n_\ell}\to V_\ell$. Moreover, $\boldsymbol{M}_\ell\in\mathbb{R}^{n_\ell\times n_\ell}$ is the mass matrix with respect to $\widetilde{\Lambda}_\ell$, i.e., $(\boldsymbol{M}_\ell)_{pq} := (\widetilde{\lambda}_p^\ell,\widetilde{\lambda}_q^\ell)_{L^2(\Omega)}$ for $p,q\in\mathcal{N}_\ell$, $\ell\in\{0,\dots,L\}$.

Note that the mapping $\Pi_{\ell-1}^\ell$ between the given spaces $X_{\ell-1}$ and X_ℓ usually does not act on $V_{\ell-1}$ directly. Still, the matrix $\boldsymbol{\Pi}_{\ell-1}^\ell$ determines a linear transfer operator $\widetilde{\Pi}_{\ell-1}^\ell : V_{\ell-1}\to V_\ell$ by

$$v\mapsto \widetilde{\Pi}_{\ell-1}^\ell v := \widetilde{\boldsymbol{\Phi}}_\ell(\boldsymbol{\Pi}_{\ell-1}^\ell\widetilde{\boldsymbol{\Phi}}_{\ell-1}^{-1}(v)), \quad \forall\, v\in V_{\ell-1}, \quad \forall\,\ell\in\{1,\dots,L\}.$$

One can easily see that $\widetilde{\Pi}_{\ell-1}^\ell$ is the natural embedding because it interpolates the respective basis exactly. Thus, we can regard the matrix $\boldsymbol{\Pi}_{\ell-1}^\ell$ as an algebraic representation of the natural embedding of $V_{\ell-1}$ into V_ℓ. Consequently, the L^2-projection from V_ℓ to $V_{\ell-1}$ is represented by the matrix $\boldsymbol{M}_{\ell-1}^{-1}(\boldsymbol{\Pi}_{\ell-1}^\ell)^T\boldsymbol{M}_\ell\in\mathbb{R}^{n_{\ell-1}\times n_\ell}$. This holds true for any imaginable set of operators between the original non-nested spaces $(X_\ell)_{\ell=0,\dots,L}$; no special structure is required.

With this information we can summarize our efforts as follows. From the completely unrelated finite element spaces $(X_\ell)_{\ell=0,\dots,L}$ we have constructed a sequence of nested spaces $(V_\ell)_{\ell=0,\dots,L}$ such that the given prolongation operators $(\Pi_{\ell-1}^\ell)_{\ell=1,\dots,L}$ induce the natural embeddings $(V_{\ell-1}\hookrightarrow V_\ell)_{\ell=1,\dots,L}$ by their matrix representations $(\boldsymbol{\Pi}_{\ell-1}^\ell)_{\ell=1,\dots,L}$ with respect to the original bases $(\Lambda_\ell)_{\ell=0,\dots,L}$. In particular, the coarse level matrices for the nested spaces with the respective bases $\widetilde{\Lambda}_\ell$, as customary in a variational approach, can be written as

$$\boldsymbol{A}_{\ell-1} = (\boldsymbol{\Pi}_{\ell-1}^\ell)^T\boldsymbol{A}_\ell\,\boldsymbol{\Pi}_{\ell-1}^\ell\in\mathbb{R}^{n_{\ell-1}\times n_{\ell-1}}, \quad \forall\,\ell\in\{1,\dots,L\}. \tag{2}$$

If $\boldsymbol{A}_L$ is symmetric positive definite and if $\boldsymbol{\Pi}_{\ell-1}^\ell$ has full rank for all $\ell\in\{1,\dots,L\}$, the respective coarse level matrices $(\boldsymbol{A}_\ell)_{\ell=0,\dots,L-1}$ are symmetric positive definite, too. Note that the bandwidth of the coarse matrices depends on the transfer concept employed to obtain the prolongation operators.

The multiplicative Schwarz method studied in this paper is the symmetric multigrid $\mathcal{V}$-cycle in the novel space hierarchy $(V_\ell)_{\ell=0,\dots,L}$, which combines (Gauß–Seidel) smoothing and coarse level correction in the standard way. Naturally, only multiplications with the matrices $(\boldsymbol{\Pi}_{\ell-1}^\ell)_{\ell=1,\dots,L}$ and their transposes appear in the interlevel transfer of the algorithm; no mass matrices need to be inverted. Given the meshes $(\mathcal{T}_\ell)_{\ell=0,\dots,L}$ and a suitable transfer concept, we can compute all auxiliary matrices in a setup phase.

For a complete convergence analysis of this class of algorithms, which puts the semi-geometric approach into the well-known context of [1], we refer to [5]. There, we carefully distinguish between the generally different domains $(\Omega_\ell)_{\ell=0,\dots,L}$ and elaborate requirements for the meshes and the interlevel transfer to obtain a quasi-optimal result.

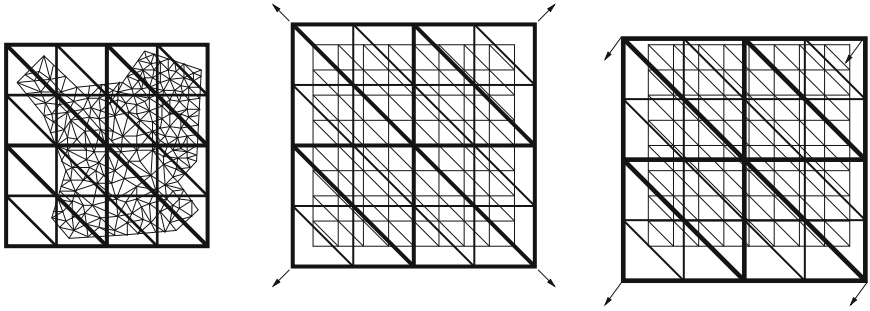

Fig. 1. Simplified sketch in $d = 2$. Basic idea of the coarse space construction based on non-nested meshes (*left*). Concerning the experiments: scaling (*center*) and translation (*right*) of the coarse meshes keeping the respective fine mesh fixed. We emphasize that all computations are in $d = 3$

The geometric nature of the construction usually requires some modifications of the meshes and operators, e.g., to ensure full rank. Moreover, a prevalent technique to keep the operator complexity $\mathscr{C}_{\mathrm{op}} := \sum_{\ell=0}^{L} n_\ell^A / n_L^A$ small, where n_ℓ^A is the number of non-zero entries of $\boldsymbol{A}_\ell$, is truncation of the prolongation operators by deleting the entries of $(\boldsymbol{\Pi}_{\ell-1}^{\ell})_{\ell=1,\dots,L}$ which are less than a truncation parameter $\varepsilon_{\mathrm{tr}} > 0$ times the maximal entry in the respective row. Afterwards, the modified rows are rescaled such that the row totals remain unchanged; see [7]. All this is done in the setup before the computation of the respective Galerkin products (2). In this paper, we choose $\Pi_{\ell-1}^{\ell}$ as standard nodal interpolation in X_ℓ for $\ell \in \{1,\dots,L\}$, namely $\Pi_{\ell-1}^{\ell} v := \sum_{p \in \mathscr{N}_\ell} v(p) \lambda_p^\ell$ for all $v \in X_{\ell-1}$, and refer to [5] for a detailed discussion.

3 Numerical Studies

3.1 The Almost Nested Limiting Case

We consider a hierarchy of four nested meshes $(\mathscr{T}_\ell)_{\ell=0,\dots,3}$ of the unit cube in $\mathbb{R}^3$ where the coarsest mesh consists of 768 elements with 189 nodes. Throughout the study, we keep the finest mesh $\mathscr{T}_L = \mathscr{T}_3$ with 393,216 elements and 68,705 nodes fixed. In contrast, the coarse domains $(\Omega_\ell)_{\ell<3}$ and the corresponding coarse meshes $(\mathscr{T}_\ell)_{\ell<3}$ are scaled around the center with a different factor between 0.95 and 1.05 for each set of tests; see Fig. 1 (center).

In the semi-geometric framework, it is absolutely necessary to perform a truncation procedure to retain the optimality of the algorithms. Otherwise, one can in general not prevent the appearance of very small and thus irrelevant entries in the prolongation matrices. We study the complexity of the constructed space hierarchy and the convergence of the semi-geometric multigrid method (stand-alone or in a preconditioned conjugate gradient method) for a variety of values for the parameter $\varepsilon_{\mathrm{tr}}$ in $[0.01, 0.49]$. Note that, for linear finite elements associated with simplicial meshes, it does generally not make sense to choose $\varepsilon_{\mathrm{tr}}$ greater than or equal to 0.5.

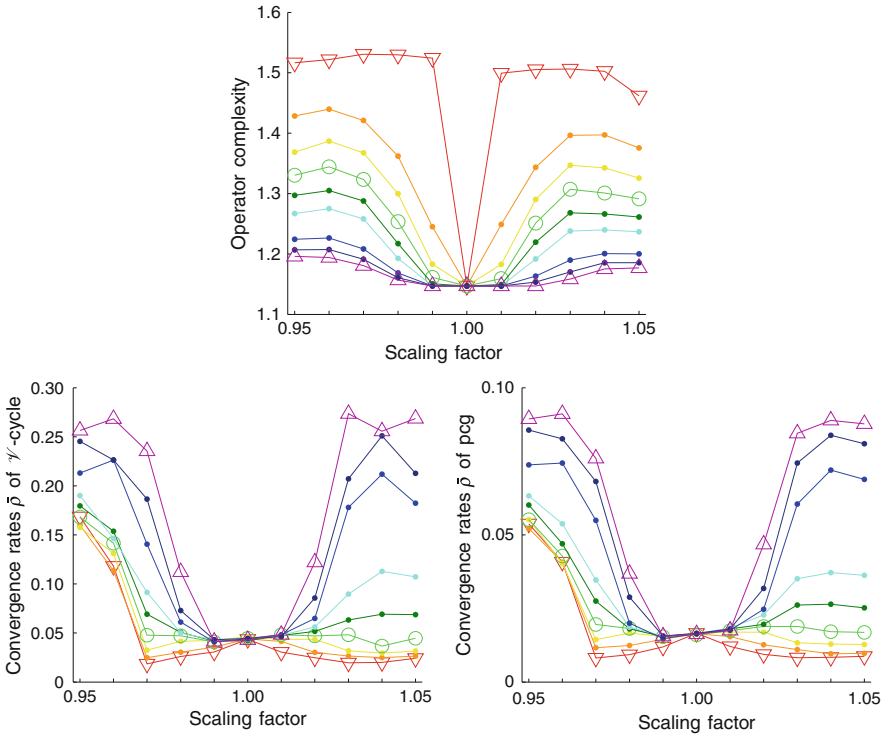

Fig. 2. The complexity measure $\mathscr{C}_{\text{op}}$ (*top*) and the convergence rates $\bar{\rho}_{\mathscr{V}(2,2)}$ (*left*) and $\bar{\rho}^{\text{pcg}}_{\mathscr{V}(2,2)}$ (*right*) of a semi-geometric multigrid method, plotted versus the scale of the coarse meshes. Each line represents a different parameter $\varepsilon_{\text{tr}} \in [0.01, 0.49]$. The *marked lines* correspond to the values 0.01 (▽), 0.20 (○) and 0.49 (△), respectively

This is because such a choice would result in deleting entries even in case of perfectly nested meshes, leaving nodes without direct coupling to the next coarser level.

The results of the experiments with scaled $(\Omega_\ell)_{\ell<3}$ are illustrated in Fig. 2. Each single line represents either the complexity $\mathscr{C}_{\text{op}}$ or one of the asymptotic convergence rates $\bar{\rho}_{\mathscr{V}(2,2)}$ and $\bar{\rho}^{\text{pcg}}_{\mathscr{V}(2,2)}$ for a fixed parameter ε_{tr} plotted versus the scale of the coarse meshes. The lines corresponding to the extreme ε_{tr}-values 0.01 and 0.49 are marked by downward and upward triangles, respectively; an intermediate value of 0.20 is marked by circles. Table 1 contains the numbers for these three values. We stop with the scales 0.95 and 1.05, respectively. For smaller factors, the convergence rates further increase quite fast as less and less of the computational domain $\Omega = \Omega_L$ is covered by the coarse meshes; the complexity measures do not change much in this case. For larger factors, the convergence rates slowly increase whereas the complexity measures decrease. This is due to the fact that more and more elements of the coarse meshes lie completely outside the computational domain.

scale	$\mathscr{C}_{\mathrm{op}}$	$\bar{\rho}_{\mathscr{V}(2,2)}$	$\bar{\rho}^{\mathrm{pcg}}_{\mathscr{V}(2,2)}$	$\mathscr{C}_{\mathrm{op}}$	$\bar{\rho}_{\mathscr{V}(2,2)}$	$\bar{\rho}^{\mathrm{pcg}}_{\mathscr{V}(2,2)}$	$\mathscr{C}_{\mathrm{op}}$	$\bar{\rho}_{\mathscr{V}(2,2)}$	$\bar{\rho}^{\mathrm{pcg}}_{\mathscr{V}(2,2)}$
0.95	1.52	0.169	0.054	1.33	0.168	0.055	1.20	0.256	0.089
0.96	1.52	0.118	0.041	1.34	0.142	0.043	1.19	0.268	0.091
0.97	1.53	0.018	0.008	1.32	0.048	0.020	1.18	0.235	0.076
0.98	1.53	0.026	0.009	1.25	0.047	0.018	1.16	0.112	0.037
0.99	1.52	0.031	0.012	1.16	0.041	0.015	1.15	0.041	0.016
1.00	1.15	0.044	0.016	1.15	0.044	0.016	1.15	0.044	0.016
1.01	1.50	0.031	0.012	1.16	0.048	0.017	1.15	0.048	0.018
1.02	1.51	0.025	0.009	1.25	0.047	0.019	1.15	0.122	0.047
1.03	1.51	0.020	0.008	1.31	0.048	0.019	1.16	0.273	0.085
1.04	1.50	0.020	0.008	1.30	0.037	0.017	1.18	0.256	0.089
1.05	1.46	0.024	0.009	1.29	0.045	0.017	1.18	0.269	0.088
	$\varepsilon_{\mathrm{tr}} = 0.01$			$\varepsilon_{\mathrm{tr}} = 0.20$			$\varepsilon_{\mathrm{tr}} = 0.49$		

Table 1. Studying the convergence behavior for a family of almost nested meshes associated with the unit cube. The middle row (scale 1.00) corresponds to the completely nested case in which the approach coincides with the standard geometric multigrid method

3.2 Robustness of the Coarse Level Hierarchy

The second experiment is to further investigate the influence of perturbations of the meshes on the coarse level hierarchy and the multigrid performance. Here, we consider different translations of the coarse meshes associated with the cube of scale 1.05 in direction of the unit vector $(\frac{2}{3}, \frac{2}{3}, \frac{1}{3})^T \in \mathbb{R}^3$ by sizes up to 0.12. In this case, the computational domain $\Omega = \Omega_L$ is covered by the domains $(\Omega_\ell)_{\ell<L}$ for almost the entire range of translations; see Fig. 1 (right). Basic robustness of the semi-geometric construction is demonstrated by the results in Fig. 3 where the parameter $\varepsilon_{\mathrm{tr}}$ again varies in the interval $[0.01, 0.49]$.

4 Discussion of the Results

As expected and observed in the vast majority of experiments, the convergence rates principally increase with increasing truncation parameter, which indicates that the constructed coarse spaces have adequate approximation power. Note that the deterioration of the convergence behavior is usually rather slow, though. It is evident that the semi-geometric methods, which leave the coarse meshes flexible, coincide with the standard geometric variants in the special case of nested meshes. In addition, an important observation from Sect. 3.1 is that both the complexities $\mathscr{C}_{\mathrm{op}}$ and the convergence rates of the geometric multigrid methods are retained in case the meshes are almost nested if a suitable parameter $\varepsilon_{\mathrm{tr}}$ is applied; see the discussion below. This also indicates that our construction is robust in the sense that the coarse level hierarchy (and with it the multigrid convergence) only varies slightly if the coarse meshes themselves change slightly. Perturbations of the meshes are irrelevant for the

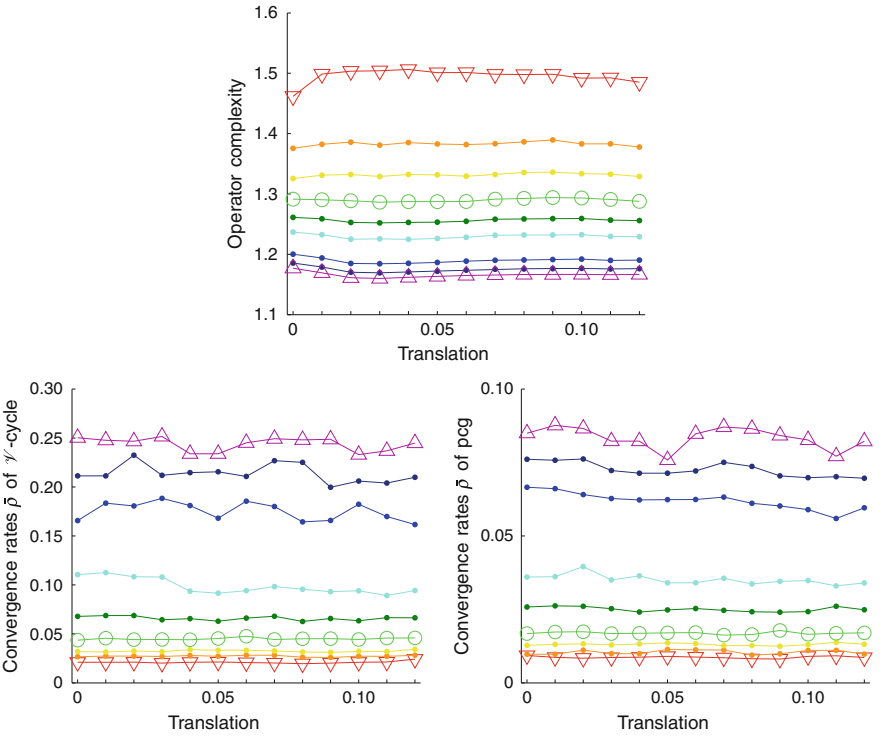

Fig. 3. The numbers $\mathscr{C}_{\text{op}}$ (*top*), $\bar{\rho}_{\mathscr{V}(2,2)}$ (*left*), and $\bar{\rho}^{\text{pcg}}_{\mathscr{V}(2,2)}$ (*right*). Each line represents a different parameter $\varepsilon_{\text{tr}} \in [0.01, 0.49]$ plotted versus the size of the coarse mesh translation

efficiency of the methods. This can also be seen clearly in the experiments described in Sect. 3.2.

As a general rule, we observe the following effects in Sect. 3.1. The larger the parameter ε_{tr} the less sensitive is the complexity $\mathscr{C}_{\text{op}}$ to changes of the coarse meshes. The smaller ε_{tr} the less sensitive are the convergence rates to changes of the coarse meshes. In our examples, the convergence actually improves in case of small perturbations for sufficiently small ε_{tr}. This is of course accompanied by a rapid increase of $\mathscr{C}_{\text{op}}$. The choice $\varepsilon_{\text{tr}} = 0.20$ (which is, interestingly enough, a standard value in many algebraic multigrid algorithms) is a reasonable attempt to achieve the two competing goals. It manages to keep the convergence rates almost constant for a rather broad range of different problem sizes while leading to an only moderate increase of $\mathscr{C}_{\text{op}}$.

Finally, let us compare to the general semi-geometric case. For an unstructured mesh with similar size (64,833 nodes) approximating a ball, the measured rates, $\bar{\rho}_{\mathscr{V}(2,2)} = 0.060$ and $\bar{\rho}^{\text{pcg}}_{\mathscr{V}(2,2)} = 0.024$, are not much worse than the ones produced by the geometric method on the cube with completely nested meshes, $\bar{\rho}_{\mathscr{V}(2,2)} = 0.044$ and $\bar{\rho}^{\text{pcg}}_{\mathscr{V}(2,2)} = 0.016$. However, for unstructured meshes without natural coarse level hierarchy, it seems impossible to achieve this fast convergence with an operator complexity as small as 1.15 which is easily obtained in the structured case.

For comparison, we have $\mathscr{C}_{\mathrm{op}} = 1.38$ for the ball. A whole series of experiments studying the asymptotics of the semi-geometric preconditioners can be found in [5].

5 Conclusion

In this paper, we reported on numerical studies of a class of preconditioners based on non-nested meshes. Considering the almost nested case, we determined a truncation parameter $\varepsilon_{\mathrm{tr}} = 0.20$ of the interlevel transfer to be reasonable in order to ensure that the efficiency of the completely nested case is in large part retained. Moreover, perturbations of the meshes turned out to be irrelevant for the efficiency of the methods.

Our results also show that, in the variational coarse space construction, it is appropriate to choose auxiliary meshes mimicking geometric coarsening, which leads to particularly small hierarchical overhead (less than 40 %). This is in contrast to the non-variational variant of the auxiliary space method [9] where both analysis and experiments indicate that the sizes of the original space and of the auxiliary space need to be comparable in a quite restrictive sense such that $\mathscr{C}_{\mathrm{op}}$ is usually clearly larger than two.

Acknowledgments This work was supported by the Bonn International Graduate School in Mathematics and by the Iniziativa Ticino in Rete.

Bibliography

[1] J.H. Bramble, J.E. Pasciak, J. Wang, and J. Xu. Convergence estimates for multigrid algorithms without regularity assumptions. *Math. Comput.*, 57(195):23–45, 1991.

[2] X. Cai. The use of pointwise interpolation in domain decomposition methods with non-nested meshes. *SIAM J. Sci. Comput.*, 16(1):250–256, 1995.

[3] T. Chan, B. Smith, and J. Zou. Overlapping Schwarz methods on unstructured meshes using non-matching coarse grids. *Numer. Math.*, 73(2):149–167, 1996.

[4] T. Chan, J. Xu, and L. Zikatanov. An agglomeration multigrid method for unstructured grids. In J. Mandel et al., editor, *Domain Decomposition Methods 10*, volume 218 of *Contemp. Math.*, pages 67–81. AMS: Providence, RI, 1998.

[5] T. Dickopf. *Multilevel Methods Based on Non-Nested Meshes*. PhD thesis, University of Bonn, 2010. http://hss.ulb.uni-bonn.de/2010/2365.

[6] M. Griebel and M.A. Schweitzer. A particle-partition of unity method. Part III: A multilevel solver. *SIAM J. Sci. Comput.*, 24(2):377–409, 2002.

[7] K. Stüben. An introduction to algebraic multigrid. In U. Trottenberg et al., editor, *Multigrid*, pages 413–532. Academic Press, London, 2001.

[8] A. Toselli and O. Widlund. *Domain Decomposition Methods – Algorithms and Theory*, volume 34 of *Springer Ser. Comput. Math.* Springer, 2005.

[9] J. Xu. The auxiliary space method and optimal multigrid preconditioning techniques for unstructured grids. *Computing*, 56(3):215–235, 1996.