A Study of Prolongation Operators Between Non-nested Meshes

Thomas Dickopf[1]* and Rolf Krause[2]

[1] Institute for Numerical Simulation, University of Bonn,
53115 Bonn, Germany, dickopf@ins.uni-bonn.de

[2] Institute of Computational Science, University of Lugano,
6904 Lugano, Switzerland, rolf.krause@usi.ch

Summary. For a class of multilevel preconditioners based on non-nested meshes, we study numerically several selected prolongation and restriction operators. Robustness with respect to the mesh size and to jumps in the coefficients is demonstrated.

Key words: multilevel preconditioners, geometric and algebraic multigrid methods, finite elements, non-nested meshes, prolongation

1 Introduction

The topic of this paper is a preconditioning strategy based on non-nested meshes for linear problems arising from finite element discretizations. Using a set of suitable prolongation and restriction operators, we give an explicit construction of a nested space hierarchy with corresponding bases. The analysis of the resulting multilevel preconditioner can be carried out in a natural way by looking at the original non-nested spaces and the connecting operators.

The present approach has the advantage, as compared with (purely) algebraic multigrid methods, that the little geometric information entering the setup leads to a very efficient multilevel hierarchy. Both grid and operator complexity are particularly small. Moreover, in our numerical experiments, we observed robustness of the developed semi-geometric method with respect to jumps in the coefficients; its performance does also not deteriorate for systems of partial differential equations.

Aside from [1, 6, 14], our semi-geometric framework is distinctly motivated by the literature on domain decomposition methods for unstructured meshes, e. g., [4, 5], and the auxiliary space method [15]. We refer to [8] for a more detailed review. To our knowledge, the present paper is the first one to include a numerical comparison of different, to a greater or lesser extent sophisticated candidates for the prolongation between non-nested finite element spaces. We examine operators from or at least motivated by [2, 4, 5, 7, 8, 10, 11].

* Supported by Bonn International Graduate School in Mathematics

2 Multilevel Preconditioners Based on Non-nested Meshes

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

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

For a Galerkin discretization of (1), let $(\mathcal{T}_l)_{l\in\mathbb{N}}$ be a family of *non-nested* shape regular meshes of domains $(\Omega_l)_{l\in\mathbb{N}}$. For a fixed finest level $L \geq 2$, assume for simplicity that $\Omega_L = \Omega$ and, in addition, $\Omega_l \supset \Omega$ for all $l \in \{0,\dots,L-1\}$. Let $h_l : \Omega_l \to \mathbb{R}_+$ be a suitable function, e. g., piecewise constant, reflecting the local mesh size of $\mathcal{T}_l$. We denote the set of nodes of $\mathcal{T}_l$ by $\mathcal{N}_l$ and abbreviate $N_l := |\mathcal{N}_l|$. At each level l, we consider the space X_l of Lagrange conforming finite elements of first order with incorporated homogeneous Dirichlet boundary conditions and denote its nodal basis functions as $(\lambda_p^l)_{p\in\mathcal{N}_l}$ with $\lambda_p^l(q) = \delta_{pq}$, $p,q \in \mathcal{N}_l$. Finally, set $\omega_p := \operatorname{supp}(\lambda_p^l)$ for $p \in \mathcal{N}_l$.

Now, the goal is to develop an efficient method for the iterative solution of the ill-conditioned discrete problem

$$u \in X_L: \quad A_L u = f_L, \tag{2}$$

with the stiffness matrix A_L associated with X_L, namely $(A_L)_{pq} := a(\lambda_p^L, \lambda_q^L)$ for $p,q \in \mathcal{N}_L$, and the right hand side f_L given by $(f_L)_p := f(\lambda_p^L)$. Here and in the following, we do not aspire to distinguish strictly between an operator or function and its representation with respect to a finite element basis.

We introduce a rather simple multilevel preconditioner C_L. The delicate point, though, is the construction of an appropriate hierarchy of spaces from the originally unrelated spaces $(X_l)_{l=0,\dots,L}$. For this purpose, let the spaces $(X_l)_{l=0,\dots,L}$ be connected by the prolongation operators $(\mathit{\Pi}_{l-1}^l)_{l=1,\dots,L}$, namely

$$\mathit{\Pi}_{l-1}^l : X_{l-1} \to X_l, \quad \forall\, l \in \{1,\dots,L\}.$$

A closer examination of selected linear operators $(\mathit{\Pi}_{l-1}^l)_{l=1,\dots,L}$ will be the key issue of this paper. We construct a nested sequence of spaces $(V_l)_{l=0,\dots,L}$ via $V_L := X_L$, $V_{L-1} := \mathit{\Pi}_{L-1}^L X_{L-1}$, and further

$$V_l := \mathit{\Pi}_{L-1}^L \cdots \mathit{\Pi}_l^{l+1} X_l, \quad \forall\, l \in \{0,\dots,L-2\}.$$

That way, the images of the operators determine the space hierarchy.

With the nodal bases of the finite element spaces X_{l-1} and X_l a matrix representation of $\mathit{\Pi}_{l-1}^l$ in $\mathbb{R}^{N_l \times N_{l-1}}$ can be computed for $l \in \{1,\dots,L\}$. For convenience, we set $\widetilde{\lambda}_q^L = \lambda_q^L$ for $q \in \mathcal{N}_L$. Then, a basis $(\widetilde{\lambda}_p^l)_{p\in\mathcal{N}_l}$ of V_l for $l \in \{0,\dots,L-1\}$ can recursively be defined by

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

In this manner, basis functions at level $l-1$ are nothing but linear combinations of basis functions at level l induced by the operator Π_{l-1}^{l}; they are piecewise linear with respect to the finest mesh $\mathcal{T}_L$. In particular, one can easily see that the matrix $\Pi_{l-1}^{l} \in \mathbb{R}^{N_l \times N_{l-1}}$ may be regarded as an algebraic representation of the natural embedding of the novel spaces V_{l-1} into V_l.

Then, as customary in a variational approach, the coarse level matrices with respect to the bases $(\widetilde{\lambda}_p^l)_{p\in\mathcal{N}_l}$ are assembled by Galerkin restriction in the following setup phase of the multilevel hierarchy.

Algorithm 1 (Setup semi-geometric multigrid method)
Choose type *of prolongation operator according to Sect. 3.*
setupSGM $(\text{type}, (\mathcal{T}_l)_{l=0,\dots,L})$ {
  for $(l = L, \dots, 1)$ {
    Compute prolongation operator: Π_{l-1}^{l}
    Compute coarse level operator: $A_{l-1} = (\Pi_{l-1}^{l})^T A_l \, \Pi_{l-1}^{l}$
  }
}

If A_L is symmetric positive definite and if $(\Pi_{l-1}^{l})_{l=1,\dots,L}$ have full rank, the respective coarse level operators $(A_l)_{l=0,\dots,L-1}$ are symmetric positive definite, too. In particular, the standard smoothing operators such as ν steps of the (symmetric) Gauß–Seidel or the Jacobi method, denoted by $(S_l^{\nu})_{l=1,\dots,L}$ in the following, are assumed to satisfy a smoothing property in V_l.

Algorithm 2 (Semi-geometric $\mathcal{V}(\nu,\nu)$-cycle)
For (the residual) $r \in V_l$ *compute the value* $C_l r$ *as follows.*
SGM $(l, \Pi_{l-1}^{l}, A_l, S_l^{\nu}, r)$ {
  if $(l = 0)$
    Solve exactly: $C_0 r \leftarrow A_0^{-1} r$
  else {
    Pre-smoothing step: $x \leftarrow S_l^{\nu}(r)$
    Coarse level correction:
      Restriction: $\widetilde{r} \leftarrow (\Pi_{l-1}^{l})^T (r - A_l x)$
      Recursive call: $\widetilde{x} \leftarrow \text{SGM}\ (l-1, \Pi_{l-2}^{l-1}, A_{l-1}, S_{l-1}^{\nu}, \widetilde{r})$
      Prolongation: $x \leftarrow x + \Pi_{l-1}^{l} \widetilde{x}$
    Post-smoothing step: $C_l r \leftarrow x + S_l^{\nu}(r - A_l x)$
  }
  return $C_l r$
}

The condition number of the preconditioned operator, $\kappa(C_L A_L)$, may be analyzed using the well-known result by Bramble, Pasciak, Wang, and Xu [3, Theorem 1] achieved at the beginning of the nineties. However, we emphasize that the relevant estimates, more precisely, the existence of H^1-stable operators $\mathcal{Q}_l^V : V_L \to V_l$, $l \in \{0, \dots, L-1\}$ satisfying suitable L^2-approximation properties, follow from assumptions on the original spaces $(X_l)_{l=0,\dots,L}$ and the prolongation operators $(\Pi_{l-1}^{l})_{l=1,\dots,L}$ rather than on the spaces $(V_l)_{l=0,\dots,L}$. Note that the possible dependence of the results on the number of levels is not ruled out. For the details we refer to [8].

Theorem 1 (Quasi-optimal preconditioning [8, Theorem 3.5]). *Let $\Pi_{l-1}^l : X_{l-1} \to X_l$, $l \in \{1, \dots, L\}$, be H^1-stable prolongation operators with the L^2-approximation properties*

$$\|h_l^{-1}(v - \Pi_{l-1}^l v)\|_{L^2(\Omega)} \lesssim \|v\|_{H^1(\Omega)}, \quad \forall\, v \in X_{l-1}. \tag{3}$$

Assume there are H^1-stable mappings $\mathcal{Q}_l^X : X_L \to X_l$, $l \in \{0, \dots, L-1\}$, satisfying the analogous L^2-approximation properties. If, additionally, the operators $(S_l^\nu)_{l=1,\dots,L}$ have suitable smoothing properties, then the multilevel method C_L defined by Algorithms 1 and 2 is a quasi-optimal preconditioner, i. e., $\kappa(C_L A_L) \lesssim 1$ uniformly with respect to the mesh size.

3 Looking for Suitable Prolongation Operators

The presented preconditioner based on non-nested meshes is related to both agglomeration multigrid methods [6] and aggregation-based algebraic multigrid methods [1, 12, 14]. The difference is that in our semi-geometric approach the coarsening reflected by the "agglomerates" or "aggregates", respectively, and thus the structures of the coarse level operators are in large part determined by the meshes $(\mathcal{T}_l)_{l=0,\dots,L}$. Still, the second ingredient to the setup in Algorithm 1 is a set of prolongation operators $(\Pi_{l-1}^l)_{l=1,\dots,L}$. It turns out that the little geometric information that enters the method is enough to generate an efficient space hierarchy with relatively smooth coarse level functions. Especially, no additional prolongation smoother [14] is needed.

The paradigm one should keep in mind is that, in the multigrid context, the L^2-projection is a natural way to prolongate a coarse level correction to a finer mesh. As the evaluation of the discrete L^2-projection is computationally inefficient in case of non-nested meshes, one has to seek an alternative.

In this section, some selected (intuitive and more elaborate) candidates for the prolongation operators $(\Pi_{l-1}^l)_{l=1,\dots,L}$ are specified. This is done in preparation for the numerical studies in the last section of the paper; for a more detailed review we refer to [8]. First, we consider the most elementary operator. In the literature on domain decomposition methods for unstructured meshes, the *standard finite element interpolation* $\mathcal{I}_{l-1}^l : \mathcal{C}^0(\Omega) \supset X_{l-1} \to X_l$, $u \mapsto \mathcal{I}_{l-1}^l u := \sum_{p \in \mathcal{N}_l} u(p)\lambda_p^l$, has been proposed to be used with non-nested coarse meshes. Different proofs of the H^1-stability and the L^2-approximation property (3) can be found in [4, 5, 13] in the context of partition lemmas.

Whereas the above mapping operates on continuous functions only, the rest of the operators comprise a weighting and are thus well-defined on appropriate Lebesgue spaces. The *Clément quasi-interpolation operator* first introduced in the influential paper [7] is defined by

$$\mathcal{R}_{l-1}^l : L^2(\Omega) \supset X_{l-1} \to X_l, \quad u \mapsto \mathcal{R}_{l-1}^l u := \sum_{p \in \mathcal{N}_l} (Q_p u)(p)\lambda_p^l, \tag{4}$$

with the L^2-projections Q_p onto the local polynomial spaces $\mathbb{P}_r(\omega_p)$ of degree $r \in \mathbb{N}$. It is probably most famous for its frequent use in proofs of the reliability of a posteriori error estimators. In Sect. 4, we employ $\mathcal{R}_{l-1}^l$ with $r = 0$.

The following *L^2-quasi-projection operator* has been analyzed in [2] to construct approximation operators replacing the L^2-projection from the space $H^1(\Omega)$ to the discrete spaces X_l. It is defined by

$$\widehat{\mathcal{Q}}_{l-1}^l : L^2(\Omega) \supset X_{l-1} \to X_l, \quad u \mapsto \widehat{\mathcal{Q}}_{l-1}^l u := \sum_{p \in \mathcal{N}_l} \frac{(\lambda_p^l, u)_{L^2(\Omega)}}{\int_{\omega_p} \lambda_p^l} \lambda_p^l. \tag{5}$$

Note that this is the operator obtained from the discrete L^2-projection by lumping the mass matrix associated with X_l.

In [8], we have investigated a new operator, primarily motivated by [10, 11]. For its definition, choose a set of functions $(\psi_p^l)_{p \in \mathcal{N}_l}$ with $\psi_p^l \in \mathcal{C}^0(\omega_p)$ for all $p \in \mathcal{N}_l$ such that $(\psi_p^l, \lambda_q^l)_{L^2(\Omega)} = \delta_{pq} \int_{\omega_p} \lambda_p^l$, $\forall\, p, q \in \mathcal{N}_l$. Then, the *pseudo-$L^2$-projection operator* $\mathcal{P}_{l-1}^l : L^2(\Omega) \supset X_{l-1} \to X_l$ is defined by a Petrov–Galerkin variational formulation with trial space X_l and test space $Y_l := \operatorname{span}\{\psi_p^l |\ p \in \mathcal{N}_l\} \not\subset \mathcal{C}^0(\Omega)$ yielding the representation formula

$$\mathcal{P}_{l-1}^l u = \sum_{p \in \mathcal{N}_l} \frac{(\psi_p, u)_{L^2(\Omega)}}{\int_{\omega_p} \lambda_p^l} \lambda_p^l. \tag{6}$$

For this last operator, the authors have proved the H^1-stability and the L^2-approximation property in case of shape regular meshes. Therefore, the multilevel preconditioner C_L defined by Algorithms 1 and 2 is quasi-optimal with the choice $\Pi_{l-1}^l = \mathcal{P}_{l-1}^l$ for $l \in \{1, \ldots, L\}$; see [8, Theorem 5.7]. Note that the present considerations do not yield estimates which are independent of the number of levels L, though.

Let us remark that, if the meshes $\mathcal{T}_{l-1}$ and $\mathcal{T}_l$ are nested, the operators $\mathcal{R}_{l-1}^l$ and $\widehat{\mathcal{Q}}_{l-1}^l$ do not coincide with the L^2-projection, which is the same as the finite element interpolation in this case; see [8]. In contrast, the operator $\mathcal{P}_{l-1}^l$ is always a projection, especially the L^2-projection in the nested case.

To evaluate (4), (5) or (6) in the setup phase (Algorithm 1) exactly, one has to compute the intersections of elements in consecutive meshes. In practice, good results may be obtained by an approximate numerical integration via a quadrature rule solely based on the finer mesh.

Without loss of generality, we may assume that the prolongation operators do not contain any zero columns; otherwise the respective coarse degrees of freedom are not coupled to the original problem (2) and should be removed in Algorithm 1. As a measure for the efficiency of the multilevel hierarchy itself, in addition to iteration counts or convergence rates, we put forward the notions of grid complexity $\mathcal{C}_{\text{gr}}$ and operator complexity $\mathcal{C}_{\text{op}}$ defined by

$$\mathcal{C}_{\text{gr}} := \sum_{l=0}^{L} N_l / N_L, \quad \mathcal{C}_{\text{op}} := \sum_{l=0}^{L} n_l / n_L, \tag{7}$$

that are common in the AMG literature. Here, n_l is the number of non-zero entries in A_l, $l \in \{0, \dots, L\}$. A prevalent technique to keep $\mathcal{C}_{\text{gr}}$ and $\mathcal{C}_{\text{op}}$ small (and the application of Algorithm 2 efficient) is truncation of the prolongation operators by deleting the entries that are smaller than $\varepsilon_{\text{tr}} = 0.2$ times the maximal entry in the respective row and rescaling afterwards; see [12].

4 Numerical Results

Because of the geometric nature of the coarsening procedure, it is important to analyze its dependence on the meshes. This is done, in each single study, by choosing an independently generated fine mesh $\mathcal{T}_L$ (of varying mesh size) approximating the unit ball. In the fashion of an auxiliary space method [15], we use nested coarse meshes $(\mathcal{T}_l)_{l<L}$ associated with the unit cube (structured; 189, 1,241 and 9,009 nodes, respectively) and standard interpolation between the levels $l < L$. Note that the different fine meshes yield different coarse spaces $(V_l)_{l<L}$ although the respective coarse meshes are unchanged.

We report on the convergence of the conjugate gradient method (until the residual norm is below 10^{-16}) preconditioned by the semi-geometric $\mathcal{V}(3,3)$-cycle (Algorithm 2) with symmetric Gauß–Seidel smoothing. In view of the affinity of our method to aggregation-based AMG, it is reasonable to examine whether an over-relaxation of the coarse level correction can improve the convergence; see [1, 12]. For the model problem (1) and $d = 3$, we identify the scaling factors 1.0, 1.1, 1.1, 1.2 for the four operators, respectively, in Fig. 1. But note that over-correction is not really necessary. The results of our experiments can be found in Table 1 for a constant coefficient $\alpha = 1$ and in Table 2 for a coefficient function α constant on each element in $\mathcal{T}_L$ with a randomly chosen value 1 or 10,000. For each operator, we give the number of iteration steps and an approximation of the asymptotic convergence rate. The last column (with caption "none") always contains the values for the one-level pcg with the symmetric Gauß–Seidel method as preconditioner.

A semi-geometric approach has the best chance of generating a very efficient multilevel hierarchy. This can be verified by noting that the complexities (7) are quite small in all numerical studies, namely down to $\mathcal{C}_{\text{gr}} = 1.035$ and $\mathcal{C}_{\text{op}} = 1.070$, at the same time with convergence rates of 0.0713 and 0.0747, respectively, for the scalar problem with 204,675 nodes; see Table 3.

Finally, let us remark that the convergence behavior does not deteriorate for systems of PDEs; see Table 4. This is in contrast to most algebraic multigrid methods; we refer to [9] for a discussion and an exemplary comparison of different algorithms. Certainly, one reason for this robustness is the fact that we treat the different physical unknowns separately, e. g., the (scalar) displacement in direction of a chosen basis of $\mathbb{R}^d$ in case of (linear) elasticity problems. Thus, the coarse level hierarchy is the same in each component. In the present linear elastic example, we observe a superior performance of the projections $\mathcal{I}$ and $\mathcal{P}$ over the operators $\mathcal{R}$ and $\widehat{\mathcal{Q}}$.

Table 1. Convergence of the pcg for a constant coefficient α. In this and the other convergence studies (see Tables 2 and 4), we give the number of needed pcg iterations and an approximate asymptotic convergence rate for $\Pi = \mathcal{I}, \mathcal{R}, \widehat{\mathcal{Q}}, \mathcal{P}$. Both quantities appear to be reasonably bounded.

#dof	$\mathcal{I}$		$\mathcal{R}$		$\widehat{\mathcal{Q}}$		$\mathcal{P}$		None	
47,348	11	0.0287	11	0.0243	11	0.0230	12	0.0394	91	0.5832
53,460	12	0.0398	11	0.0255	12	0.0345	12	0.0330	97	0.6418
64,833	12	0.0381	12	0.0362	12	0.0288	12	0.0279	102	0.6434
72,525	13	0.0486	12	0.0333	12	0.0311	12	0.0288	106	0.6576
87,244	13	0.0503	13	0.0437	13	0.0452	12	0.0313	114	0.6954
127,787	14	0.0572	15	0.0469	13	0.0421	14	0.0416	125	0.6792
204,675	16	0.0761	16	0.0684	16	0.0737	16	0.0713	146	0.7358

Table 2. Convergence of the pcg for α randomly jumping between 1 and 10,000.

#dof	$\mathcal{I}$		$\mathcal{R}$		$\widehat{\mathcal{Q}}$		$\mathcal{P}$		None	
47,348	15	0.0405	14	0.0235	14	0.0270	14	0.0296	98	0.6222
53,460	15	0.0387	14	0.0286	14	0.0277	14	0.0305	103	0.6767
64,833	15	0.0393	15	0.0357	15	0.0388	15	0.0408	109	0.6810
72,525	17	0.0653	15	0.0399	15	0.0342	15	0.0374	110	0.6470
87,244	16	0.0454	16	0.0439	16	0.0408	16	0.0412	121	0.7225
127,787	18	0.0586	17	0.0524	17	0.0553	17	0.0504	130	0.6880
204,675	20	0.0831	20	0.0790	20	0.0802	20	0.0747	152	0.7206

Table 3. Grid and operator complexity depend on the prolongation type and the problem size. As we keep the coarse meshes fixed (but not the coarse spaces) throughout the presented studies, both $\mathcal{C}_{\text{gr}}$ and $\mathcal{C}_{\text{op}}$ decrease with increasing problem size.

#elem.	#dof	$\mathcal{C}_{\text{gr}}(\mathcal{I})$	$\mathcal{C}_{\text{op}}(\mathcal{I})$	$\mathcal{C}_{\text{gr}}(\mathcal{R})$	$\mathcal{C}_{\text{op}}(\mathcal{R})$	$\mathcal{C}_{\text{gr}}(\widehat{\mathcal{Q}})$	$\mathcal{C}_{\text{op}}(\widehat{\mathcal{Q}})$	$\mathcal{C}_{\text{gr}}(\mathcal{P})$	$\mathcal{C}_{\text{op}}(\mathcal{P})$
262,365	47,348	1.144	1.364	1.144	1.490	1.144	1.430	1.143	1.338
297,620	53,460	1.128	1.320	1.128	1.418	1.128	1.373	1.127	1.296
361,907	64,833	1.106	1.263	1.106	1.330	1.106	1.298	1.105	1.242
405,256	72,525	1.095	1.233	1.095	1.291	1.095	1.263	1.095	1.214
490,617	87,244	1.079	1.190	1.079	1.230	1.079	1.210	1.079	1.173
719,951	127,787	1.055	1.128	1.055	1.149	1.055	1.138	1.055	1.117
1,161,926	204,675	1.035	1.076	1.033	1.085	1.033	1.080	1.035	1.070

Table 4. Linear elastic problem on the unit ball with Poisson ratio 0.3. Some differences in the performance of the prolongation operators may be observed.

#dof	$\mathcal{I}$		$\mathcal{R}$		$\widehat{\mathcal{Q}}$		$\mathcal{P}$		None	
142,044	17	0.0915	24	0.1969	23	0.1754	19	0.1194	137	0.7398
160,380	17	0.0974	25	0.1936	23	0.1667	19	0.1319	146	0.7936
194,499	18	0.1068	28	0.2215	26	0.1767	20	0.1473	152	0.7501
217,575	18	0.1101	29	0.2297	26	0.2107	21	0.1447	160	0.8159
261,732	19	0.1285	31	0.2382	27	0.2095	21	0.1600	175	0.8334
383,361	20	0.1439	32	0.2547	28	0.2279	22	0.1692	186	0.8395
614,025	24	0.1895	34	0.2411	31	0.2267	24	0.1969	217	0.8760

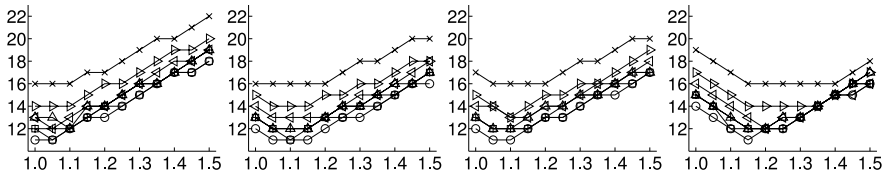

Fig. 1. Over-relaxation of the coarse level correction $\Pi_{l-1}^{l}\widetilde{x}$ in Algorithm 2. Study of the number of pcg iterations depending on the scaling factor for the choices $\Pi = \mathcal{I}$, $\mathcal{R}$, $\widehat{\mathcal{Q}}$, $\mathcal{P}$ (from *left to right*). Each *line* represents a different problem size.

References

1. R. Blaheta. Algebraic multilevel methods with aggregations: an overview. In I. Lirkov, S. Margenov, and J. Waśniewski, editors, *Large-Scale Scientific Computing*, volume 3743 of *Lecture Notes in Computer Science*, pp. 3–14. Springer, 2006.
2. J.H. Bramble, J.E. Pasciak, and P.S. Vassilevski. Computational scales of Sobolev norms with applications to preconditioning. *Math. Comput.*, 69(230):463–480, 2000.
3. 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.
4. X. Cai. The use of pointwise interpolation in domain decomposition methods with non-nested meshes. *SIAM J. Sci. Comput.*, 16(1):250–256, 1995.
5. 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.
6. T. Chan, J. Xu, and L. Zikatanov. An agglomeration multigrid method for unstructured grids. In J. Mandel, C. Farhat, and X.-C. Cai, editors, *Domain Decomposition Methods 10*, volume 218 of *Contemporary Mathematics*, pp. 67–81. AMS, Providence, RI, 1998.
7. P. Clément. Approximation by finite element functions using local regularization. *RAIRO Anal. Numér.*, 9(R-2):77–84, 1975.
8. T. Dickopf and R. Krause. A pseudo-L^2-projection for multilevel methods based on non-nested meshes. INS Preprint No. 0908 (University of Bonn) and ICS Preprint No. 2009-04 (University of Lugano), 2009.
9. M. Griebel, D. Oeltz, and M.A. Schweitzer. An algebraic multigrid method for linear elasticity. *SIAM J. Sci. Comput.*, 25(2):385–407, 2003.
10. L.R. Scott and S. Zhang. Finite element interpolation of nonsmooth functions satisfying boundary conditions. *Math. Comput.*, 54(190):483–493, 1990.
11. O. Steinbach. On a generalized L^2-projection and some related stability estimates in Sobolev spaces. *Numer. Math.*, 90(4):775–786, 2002.
12. K. Stüben. An introduction to algebraic multigrid. In U. Trottenberg et al., editors, *Multigrid*, pp. 413–532. Academic Press, London, 2001.
13. A. Toselli and O. Widlund. *Domain Decomposition Methods – Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer, Berlin 2005.
14. P. Vaněk, M. Brezina, and J. Mandel. Convergence of algebraic multigrid based on smoothed aggregation. *Numer. Math.*, 88(3):559–579, 2001.
15. J. Xu. The auxiliary space method and optimal multigrid preconditioning techniques for unstructured grids. *Computing*, 56(3):215–235, 1996.