

# Nonuniformly Weighted Schwarz Smoothers for Spectral Element Multigrid

Jörg Stiller<sup>1</sup> 

Received: 8 December 2015 / Revised: 14 September 2016 / Accepted: 19 December 2016 /  
Published online: 30 December 2016  
© Springer Science+Business Media New York 2016

**Abstract** A hybrid Schwarz/multigrid method for spectral element solvers to the Poisson equation in  $\mathbb{R}^2$  is presented. It extends the additive Schwarz method studied by Lottes and Fischer (J Sci Comput 24:45–78, 2005) by introducing nonuniform weight distributions based on the smoothed sign function. Using a V-cycle with only one pre-smoothing, the new method attains logarithmic convergence rates in the range from 1.2 to 1.9, which corresponds to residual reductions of almost two orders of magnitude. Compared to the original method, it reduces the iteration count by a factor of 1.5–3, leading to runtime savings of about 50%. In numerical experiments the method proved robust with respect to the mesh size and polynomial orders up to 32. Used as a preconditioner for the (inexact) CG method it is also suited for anisotropic meshes and easily extended to diffusion problems with variable coefficients.

**Keywords** Multigrid method · Schwarz methods · Spectral element method ·  $p$ -Version finite element method

## 1 Introduction

High-order finite element methods (FEM) enjoy an increasing interest in computational science and engineering. They include  $hp$ -FEM, spectral element methods (SEM) as well as discontinuous Galerkin methods [7, 16]. The motive that drives the development of high-order methods lies in their potential to deliver accuracy with lower cost in comparison to the first and second order methods used in common simulation tools [29]. However, realizing this advantage in practice is a formidable task. Along with curvilinear mesh generation, the provision of efficient solvers for the resulting algebraic equation systems remains the main challenge.

---

✉ Jörg Stiller  
joerg.stiller@tu-dresden.de

<sup>1</sup> Center of Advancing Electronics Dresden (cfAED), Institute of Fluid Mechanics (ISM), Technische Universität Dresden, 01062 Dresden, Germany

Projection methods for incompressible flow, or implicit discretization of diffusion terms lead to a sequence of linear elliptic problems which are related or equivalent to the Poisson equation or, more generally, the Helmholtz equation [10]. Fast solvers for such equations are therefore a crucial ingredient of competitive high-order methods and, hence, have been in focus of research for almost 30 years [1, 5, 8, 11–13, 15, 17–19, 21, 22, 24–26]. For Helmholtz or Poisson problems discretized on regular meshes, efficient multigrid (MG) techniques have been developed recently. [19] proposed additive Schwarz smoothers based on extended element domains, which attain residual reductions of approximately 0.2 within one sweep. They found that weighting the overlapping Schwarz updates by the inverse of the counting matrix, which corresponds to taking the arithmetic mean, plays a crucial role in obtaining multigrid-like iteration counts. A detailed analysis of the method was given in [18]. [13] presented a similar multigrid approach for the  $p$ -finite element method on locally refined Cartesian meshes. They used a multiplicative Schwarz smoother on element domains which possess only a minimal overlap confined to the element boundaries. [11] developed a  $p$ -multigrid method based on static condensation which, apart from pre- and post-processing, reaches linear complexity. The proposed block smoother can be classified as an additive Schwarz method using a monotonic increasing shape function for blending the overlapping updates. Using this smoother the multigrid method attained convergence rates of about 0.02 combined with a run-time efficiency that comes close to fast direct finite difference solvers. The success of this approach inspired us to extend the idea of nonuniform weighting to the full, “uncondensed” problem and thus led to the present work. The primary goal is to show how nonuniform weighting can be used to boost the performance of high-order spectral-element multigrid techniques. Further, we investigate the influence of the overlap width, smoothing strategies, additive versus multiplicative Schwarz methods and Krylov acceleration on robustness and efficiency. In addition to this, we consider the extension to diffusion problems with variable coefficients.

The remainder of the paper is organized as follows: Sect. 2 provides a brief description of the spectral element discretization. Section 3 presents the solution techniques, namely the weighted additive and multiplicative Schwarz methods, the  $p$ -multigrid method and the inexact multigrid-preconditioned conjugate gradient method. Section 4 proceeds with the discussion of numerical experiments for assessing the solution methods and application to variable diffusion. Finally, Sect. 5 concludes the paper.

## 2 Discretization

As the model problem we consider the Poisson equation

$$-\nabla^2 u = f \quad (1)$$

in the rectangular domain  $\Omega = [0, \ell_x] \times [0, \ell_y]$  with periodic boundaries. For discretization  $\Omega$  is decomposed into  $n_E = n_x \times n_y$  rectangular elements  $\Omega^{mn}$  with dimensions  $\Delta x = \ell_x/n_x$  and  $\Delta y = \ell_y/n_y$ . In each element the solution is approximated as

$$u(x, y)|_{\Omega^{mn}} \simeq \sum_{i,j=0}^p u_{ij}^{mn} \varphi_i(\xi^m(x)) \varphi_j(\eta^n(y)) \quad (2)$$

where  $\varphi_i$  are the Lagrange polynomials to the Gauss-Lobatto-Legendre (GLL) points  $\{\xi_i\}_{i=0}^p$  in the one-dimensional standard region  $\hat{\Omega} = [-1, 1]$  and  $\xi^m(x)$ ,  $\eta^n(y)$  the mapping of

coordinates from  $\Omega^{mn}$  to  $\hat{\Omega}$  [7, 16]. Concatenation of the element coefficients  $\mathbf{u}^{mn} = [u_{ij}^{mn}]$  and enforcing continuity for shared vertices and edges yields the unique global coefficients  $\mathbf{u}$  [see, e.g. [7], pp. 191–194]. Application of the Galerkin spectral element method leads to the discrete equations

$$\mathbf{A}\mathbf{u} = \mathbf{f}. \tag{3}$$

As a consequence of the tensor product ansatz (2) and the Cartesian mesh, the global system matrix in Eq. (3) assumes the tensor product form

$$\mathbf{A} = \mathbf{M}_y \otimes \mathbf{L}_x + \mathbf{L}_y \otimes \mathbf{M}_x, \tag{4}$$

where  $\mathbf{M}_*$  and  $\mathbf{L}_*$  represent the one-dimensional mass and stiffness matrices for directions  $* = x, y$ , respectively. The detailed structure of these operators and underlying spectral element techniques are well described in literature [7, 16] and therefore deliberately skipped here.

### 3 Solution Methods

For solving Eq. (3) we consider polynomial multigrid (MG) and multigrid-preconditioned conjugate gradients (MGCG). Both approaches rely on Schwarz methods for smoothing. We first present the Schwarz methods and then sketch MG and MGCG.

#### 3.1 Schwarz Methods

Schwarz methods are iterative domain decomposition techniques which improve the approximate solution by parallel or sequential subdomain solves, leading to additive or multiplicative methods, respectively. Following [19] we use extended element regions as the subdomains. Figure 1 illustrates how the subdomain  $\Omega_s$  results from the corresponding element domain  $\Omega^{mn}$  by attaching a rectangular strip matching the overlap width  $\delta_o$ . As consequence,  $\Omega_s$  adopts  $n_o$  layers of additional nodes from the neighbor elements. Note, however, that we exclude the outer layer of nodes located on  $\partial\Omega_s$ . For definiteness we define the overlap width in terms of  $n_o$  and the GLL points,

$$\delta_o = \xi_{n_o+1} + 1. \tag{5}$$

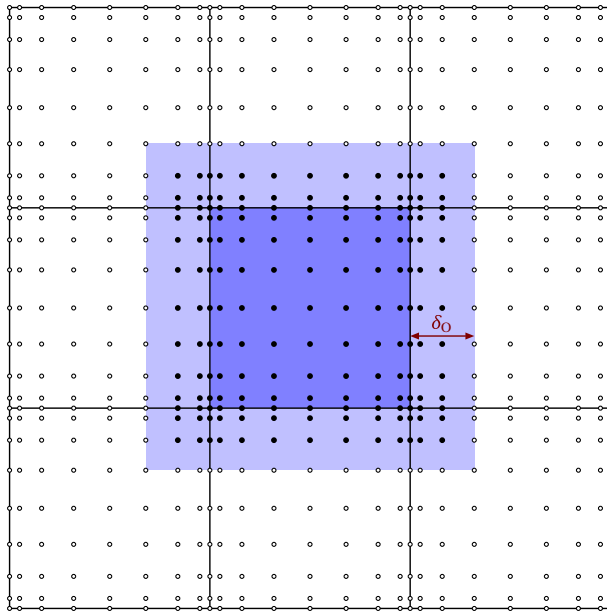
To derive a local correction to some approximate solution  $\tilde{\mathbf{u}}$  we first convert Eq. (3) into the equivalent residual form

$$\mathbf{A}\Delta\mathbf{u} = \mathbf{f} - \mathbf{A}\tilde{\mathbf{u}} = \tilde{\mathbf{r}}, \tag{6}$$

where  $\Delta\mathbf{u} = \mathbf{u} - \tilde{\mathbf{u}}$ . Further we introduce the restriction operator  $\mathbf{R}_s$  such that  $\mathbf{u}_s = \mathbf{R}_s\mathbf{u}$  gives the coefficients associated with  $\Omega_s$ . Conversely, the transposed restriction operator globalizes any local coefficients by adding zeros for exterior nodes. With these prerequisites the correction contributed by  $\Omega_s$  is defined as the solution of the subproblem

$$\mathbf{A}_{ss}\Delta\mathbf{u}_s = \mathbf{r}_s, \tag{7}$$

where  $\mathbf{A}_{ss} = \mathbf{R}_s\mathbf{A}\mathbf{R}_s^t$  represents the restricted system matrix and  $\mathbf{r}_s = \mathbf{R}_s\tilde{\mathbf{r}}$  the restricted residual. Due to the rectangular shape of the subdomain,  $\mathbf{A}_{ss}$  inherits the tensor product structure of  $\mathbf{A}$ . Using the fast diagonalization technique developed by [20] and adopted for SEM e.g. in [6], the inverse subdomain operator can be expressed in the form



**Fig. 1** Example of a subdomain used with the Schwarz method. The shaded area represents  $\Omega_S$  and the dark region in its center the corresponding element. The circles are the GLL nodes for polynomial order  $p = 8$ . Filled circles indicate the nodes that are solved for and updated;  $\delta_0$  is the overlap width

$$\mathbf{A}_{ss}^{-1} = (\mathbf{S}_y \otimes \mathbf{S}_x)(\mathbf{I} \otimes \mathbf{\Lambda}_x + \mathbf{\Lambda}_y \otimes \mathbf{I})^{-1}(\mathbf{S}_y^t \otimes \mathbf{S}_x^t), \tag{8}$$

where  $\mathbf{I}$  is the unity matrix,  $\mathbf{S}_*$  the matrix of eigenvectors to the generalized eigenproblem for the restricted one-dimensional stiffness and mass matrices, and  $\mathbf{\Lambda}_*$  the diagonal matrix of eigenvalues for directions  $* = x, y$ . With equidistant meshes, as in the present case, the operators are identical for all subdomains and, hence, the cost for their pre-computation becomes negligible. Exploiting the tensor-product structure of the inverse, the solution to a single subdomain,  $\Delta \mathbf{u}_s = \mathbf{A}_{ss}^{-1} \mathbf{r}_s$ , can be evaluated with just  $\Theta(2(p + 1 + 2n_0)^3)$  operations.

There exist several options for combining the local solutions. We consider a weighted version of the additive Schwarz method and the multiplicative Schwarz method. The multiplicative Schwarz method solves the subproblems (7) consecutively while continually updating the residual. Note that, in general, one multiplicative Schwarz iteration corresponds to the application of a non-symmetric linear operator, albeit  $\mathbf{A}$  is symmetric. However, for an even number of steps, the method is symmetrized by reversing the order of subdomains in each step, which leads to Algorithm 1.

The weighted additive Schwarz method determines all local corrections independently and computes the global correction as a linear combination of these results, i.e.

$$\Delta \mathbf{u} \simeq \sum_s \mathbf{R}_s^t (\mathbf{W} \Delta \mathbf{u}_s), \tag{9}$$

where  $\mathbf{W}$  is a diagonal local weight matrix. Application of Eq. (9) leads to Algorithm 2. Note that  $\mathbf{W} = \mathbf{I}$  recovers the classical additive Schwarz method. The arithmetic mean employed in [19] is obtained by choosing  $\mathbf{W} = \mathbf{R}_s \mathbf{C}^+ \mathbf{R}_s^t$ , where  $\mathbf{C}^+$  is the pseudoinverse of the counting matrix  $\mathbf{C} = \sum_s \mathbf{R}_s^t \mathbf{R}_s$ . We propose a more flexible approach which elevates the weights gradually from zero at the border to one in the core zone. Due to the regular shape of

---

**Algorithm 1** Multiplicative Schwarz method

---

```

1: function MSCHWARZ( $u, f, n_1$ )
2:   for  $i = 1, n_1$  do
3:     for  $e = 1, n_E$  do
4:        $s \leftarrow \begin{cases} e & i \text{ odd} \\ n_E + 1 - e & i \text{ even} \end{cases}$ 
5:        $r \leftarrow f - Au$ 
6:        $u \leftarrow u + R_s^t A_{ss}^{-1} R_s r$ 
7:     end for
8:   end for
9:   return  $u$ 
10: end function

```

---



---

**Algorithm 2** Weighted additive Schwarz method

---

```

1: function WSCHWARZ( $u, f, n_1$ )
2:   for  $i = 1, n_1$  do
3:      $r \leftarrow f - Au$ 
4:      $u \leftarrow u + \sum_{s=1}^{n_E} R_s^t W A_{ss}^{-1} R_s r$ 
5:   end for
6:   return  $u$ 
7: end function

```

---

$\Omega_s$  the weights can be cast in the tensor product form  $W = W_y \otimes W_x$ . The one-dimensional weight distributions  $W_*$  are generated from the continuous weighting function

$$w_\kappa(\xi) = \frac{1}{2} \left[ \phi_\kappa \left( \frac{\xi + 1}{\delta_o} \right) - \phi_\kappa \left( \frac{\xi - 1}{\delta_o} \right) \right], \tag{10}$$

where  $\xi$  is the 1D standard coordinate extended beyond  $\hat{\Omega}$  and  $\phi_\kappa$  is a weakly monotonic increasing shape function. In particular we consider the shape functions  $\phi_\kappa$  with  $\kappa \in \{1, 3, 5, \dots\}$  defined as

$$\phi_\kappa(x) = \begin{cases} \hat{\phi}_\kappa(x) & x \in \hat{\Omega} \\ \text{sgn}(x) & \text{else} \end{cases} \tag{11}$$

where  $\hat{\phi}_\kappa$  is a polynomial of degree  $i$  satisfying the conditions


$$\hat{\phi}_\kappa(\pm 1) = \pm 1 \tag{12a}$$

$$\frac{d^k \hat{\phi}_\kappa}{dx^k}(\pm 1) = 0, \quad 0 < k \leq (\kappa - 1)/2. \tag{12b}$$

The  $\hat{\phi}_\kappa$  are strictly monotonic in  $(-1, 1)$  and generate a smooth transition of the weight function in the overlap zone, as exemplified in Fig. 2 for the quintic case. By increasing the polynomial degree the shape function converges toward the sign function, which translates into a top hat weighting function. We remark that omitting the shape function in Eq. (10) yields the arithmetic mean. For reference, Table 1 summarizes all weight functions used in the numerical experiments.

**3.2 Multigrid**

For MG we define a series of polynomial levels  $\{p_l\}$  with  $p_l = 2^l$  increasing from 1 at  $l = 0$  to  $p$  at top level  $L$ . Correspondingly,  $u_l$  denotes the global coefficients and  $A_l$  the system matrix on level  $l$ . On the top level we have  $u_L = u$  and  $A_L = A$ , whereas on lower levels



**Fig. 2** One-dimensional weight distribution for elements of order  $p = 16$  with an overlap of  $n_O = 2$  points using a quintic shape function. The core region and the overlap zone of the subdomain are shaded in dark and light blue, respectively. Filled circles indicate the node positions (Color figure online)

**Table 1** Weight functions (WF) and related shape functions

WF	Shape function	Method
$w_1$	$\hat{\phi}_1 = x$	Linear
$w_3$	$\hat{\phi}_3 = (3x - x^3)/2$	Cubic
$w_5$	$\hat{\phi}_5 = (15x - 10x^3 + 3x^5)/8$	Quintic
$w_7$	$\hat{\phi}_7 = (35x - 35x^3 + 21x^5 - 5x^7)/16$	7th order
$w_l$	$\hat{\phi}_l = \text{sgn}(x)$	Top hat
$w_a$	$\hat{\phi}_a = 0$	Arithmetic mean

$u_l$  is the defect correction and  $A_l$  the counterpart of  $A$  obtained with elements of order  $p_l$ . For transferring the correction from level  $l - 1$  to level  $l$  we use the embedded interpolation operator  $\mathcal{I}_l$ , and for restricting the residual its transpose. These ingredients allow to build a multigrid V-cycle, which is summarized in Algorithm 3. Both, the multiplicative and the weighted additive Schwarz method stated in Algorithm 1 and 2, respectively, can serve as the SMOOTHER. The number of pre- and post-smoothing steps,  $n_{s1,l}$  and  $n_{s2,l}$ , can differ from level to level to allow variable V cycles [4]. Line 8 of Algorithm 3 defines the coarse grid solution by means of the pseudoinverse  $A_0^+$ . In our implementation the coarse problem is solved using the conjugate gradient method. To achieve convergence in spite of singularity, the right side is projected to the null space of  $A_0$ , as proposed by [14].

### 3.3 Preconditioned Conjugate Gradients

For enhancing robustness and efficiency multigrid methods can be accelerated by Krylov subspace methods [27]. In the present case, with symmetric system matrices on all grid levels, one would favor preconditioned conjugate gradients. Unfortunately, weighted additive Schwarz and multiplicative Schwarz with uneven iteration count are both non-symmetric and hence affect the symmetry of MG as well. As a remedy, it is possible to symmetrize the weighting method or to use GMRES instead of CG for acceleration. According to [18], how-

---

**Algorithm 3** Multigrid V-cycle.

---

```

1: function MULTIGRIDCYCLE( $u, f, n_S$ )
2:    $u_L \leftarrow u$ 
3:    $f_L \leftarrow f$ 
4:   for  $l = L, 1$  step  $-1$  do
5:      $u_l \leftarrow \text{SMOOTHER}(u_l, f_l, n_{S1,l})$  ▷ Pre-smoothing
6:      $f_{l-1} \leftarrow \mathcal{I}_l^t(f_l - A_l u_l)$  ▷ Residual restriction
7:   end for
8:    $u_0 \leftarrow A_0^+ f_0$  ▷ Coarse grid solution
9:   for  $l = 1, L$  do
10:     $u_l \leftarrow u_l + \mathcal{I}_l u_{l-1}$  ▷ Correction prolongation
11:     $u_l \leftarrow \text{SMOOTHER}(u_l, f_l, n_{S2,l})$  ▷ Post-smoothing
12:   end for
13:   return  $u \leftarrow u_L$ 
14: end function

```

---



---

**Algorithm 4** Inexact multigrid preconditioned conjugate gradients.

---

```

1: function MGCG( $u, f, n_S, i_{\max}, r_{\max}$ )
2:    $r_{\text{old}} \leftarrow \mathbf{0}$ 
3:    $r \leftarrow f - Au$ 
4:    $p \leftarrow \text{MULTIGRIDCYCLE}(\mathbf{0}, r, n_S)$ 
5:    $\delta \leftarrow p^t r$ 
6:   for  $i = 1, i_{\max}$  do
7:      $q \leftarrow Ap$ 
8:      $\alpha \leftarrow \delta / (p^t q)$ 
9:      $u \leftarrow u + \alpha p$ 
10:     $r \leftarrow r - \alpha q$ 
11:    if  $\|r\| \leq r_{\max}$  exit
12:     $z \leftarrow \text{MULTIGRIDCYCLE}(\mathbf{0}, r, n_S)$ 
13:     $\beta \leftarrow q^t (r - r_{\text{old}}) / \delta$ 
14:     $p \leftarrow z + \beta p$ 
15:     $\delta \leftarrow z^t r$ 
16:     $r_{\text{old}} \leftarrow r$ 
17:   end for
18:   return  $u$ 
19: end function

```

---

ever, symmetrization can deteriorate the efficiency of the method. This detrimental behavior was confirmed in own tests and, hence, the symmetrized method is not considered here.

Recently, generalizations of the conjugate gradient method have been developed that allow for relaxing some restrictions of standard CG and, thus, promise a cheaper alternative to GMRES. The use of inaccurately solved and non-symmetric preconditioners in CG-like methods has been justified, e.g., in [2,9,23]. Moreover, [3] demonstrated the suitability of the so-called flexible PCG in conjunction with non-symmetric multigrid preconditioners. Following this approach, we use the MGCG method summarized in Algorithm 4. This method is equivalent to the flexible PCG of [23], but can be regarded also as a variant of the inexact PCG proposed by [9]. The main difference to standard PCG consists in the application of the Polak-Ribière formula for  $\beta$ , instead of the Fletcher-Reeves formula, on line 13 of the algorithm. We also note that, as before with the coarse problem, the right side  $f$  must be in the null space of  $A$  if the system is singular.

## 4 Results

Numerical tests were performed to assess the influence of weighting, overlap and cycling strategy on the computational efficiency and robustness of MG and MGCG. The methods were implemented in Fortran and compiled using the GNU compiler collection 6.0 with `-O3`.

All results are based on test cases with the source  $f$  evaluated analytically from the exact solution and starting from a random initial guess confined to the interval  $[0,1]$ . The convergence speed is evaluated using the number of cycles  $n_{10}$  needed to reduce the norm of the residual by a factor of  $10^{10}$  and the average logarithmic convergence rate according to [28]

$$\bar{r} = \frac{1}{n} \log_{10} \frac{\|\mathbf{r}^{(0)}\|}{\|\mathbf{r}^{(n)}\|}, \quad (13)$$

where  $\mathbf{r}^{(n)}$  is the Euclidean norm of the residual vector after the  $n$ th cycle. Note that  $n_{10}$  is nearest integer greater than or equal to  $10/\bar{r}$ .

As an efficiency measure we define the approximate number of operator applications required for reducing the residual by a factor of  $10^k$ ,

$$\bar{\omega}_k = \frac{k}{\bar{r}} \frac{W_{\text{cyc}}}{W_{\text{op}}}, \quad (14)$$

where  $W_{\text{cyc}}$  is the cost for one cycle and  $W_{\text{op}}$  the cost for one application of the system matrix  $\mathbf{A}$ . Exploiting sum factorization [7, 16],  $W_{\text{op}}$  can be estimated as  $2n_p^3 n_E$ , where  $n_p = p + 1$  and  $n_E$  is the number of elements. According to Sect. 3.1, the cost of one Schwarz iteration is approximately  $2(n_p + 2n_o)^3 n_E$ . Assuming a maximum relative overlap of  $n_o/n_p$  this yields the estimate

$$W_{\text{cyc}} = \left[ 4 \left( 1 + 2 \frac{n_o}{n_p} \right)^3 c_s n_s + 2c_s + c_{\text{CG}} \right] n_p^3 n_E, \quad (15)$$

where  $n_s$  is the number of pre- and post-smoothing steps on the finest level,  $c_s = 4/3$  for the classical V-cycle and  $c_s = 2$  for a variable V-cycle doubling the number of smoothing steps when changing to the next lower level, and  $c_{\text{CG}} = 2$  is the extra cost for conjugate gradients with MGCG. Since the bracketed term is constant, the overall cost of one multigrid cycle scales approximately with  $pN$ , where  $N = p^2 n_E$  denotes the number of unknowns. Among the multigrid components, the smoother is by far the most expensive part, accounting for about 80–90% of the total cost in typical applications.

### 4.1 Weighting and Overlap

We consider the Poisson problem (1) in the domain  $\Omega = [0, 2]^2$ , which is uniformly subdivided in  $8 \times 8$  square elements with order  $p$  ranging from 4 to 32. The right hand side is chosen to match the exact solution  $u = \sin(\pi x) \sin(\pi y)$ . In the first test series we set the overlap to  $n_o = 1$  on all levels  $l > 0$ . Table 2 shows the measured convergence rates for MG with one pre-smoothing. Column “ $w_a$ ” corresponds to the weighted additive Schwarz method using the arithmetic mean in overlap areas. Compared to [19] our results agree well for  $p = 4$ , but show a faster convergence with higher polynomial orders. This could be attributed to using periodic instead of Dirichlet boundary conditions.

The remaining columns in Table 2 display the convergence rates for additive Schwarz smoothing with the gradual weighting introduced in Sect. 3.1 and multiplicative Schwarz. In comparison to the arithmetic mean, weighting using a smooth—cubic, quintic or 7th order—shape function roughly doubled the convergence rate for orders 4, 8 and 16, while linear



**Table 2** Convergence rates  $\bar{r}$  for MG with additive Schwarz smoothing using one pre- and no post-smoothing steps,  $n_E = 8 \times 8$  elements and a fixed overlap of  $n_O = 1$

$p$	$w_a$	$w_1$	$w_3$	$w_5$	$w_7$	$w_l$	Mult
4	0.66	0.86	1.01	1.17	1.25	0.72	1.01
8	0.40	0.83	1.17	1.29	1.23	0.52	1.29
16	0.34	0.80	0.84	0.84	0.84	0.42	1.26
32	0.32	0.43	0.43	0.43	0.43	0.38	0.76

The weighting methods are referred to as defined in Table 1. Results for the multiplicative smoother (mult) are included for comparison

**Table 3** Convergence rates for a level-dependent overlap of  $n_{O,l} = \lfloor p_l/8 \rfloor$ . For caption see Table 2

$p$	$w_a$	$w_1$	$w_3$	$w_5$	$w_7$	$w_l$	Mult
4	0.63	0.91	0.98	0.96	0.79	0.31	1.03
8	0.40	0.75	1.06	1.28	1.28	0.64	1.30
16	0.51	1.07	1.36	1.28	1.12	0.53	1.40
32	0.71	1.39	1.48	1.50	1.51	0.19	1.56

and top hat weighting yielded a lower, but still remarkable improvement. As expected, the multiplicative Schwarz smoother attained the fastest convergence. At  $p = 32$  all methods suffer a serious performance degradation, except for arithmetically weighted Schwarz, which nonetheless remains the slowest.

Inspired by these observations, several tests were run with overlaps depending on the polynomial degree on each mesh level. Table 3 shows the convergence rates for the case  $n_{O,l} = \lfloor p_l/8 \rfloor$ . Note that this choice implies  $n_O = 0$  for degrees less than 8, while reaching  $n_O = 4$  with  $p = 32$ . As a consequence, the convergence rates for  $p = 4$  are slightly lower than with  $n_O = 1$ , except for multiplicative Schwarz. For  $p \geq 16$  the increased overlap yields a considerable speedup. This improvement is most pronounced for cubic and quintic weighting, which come remarkably close to multiplicative Schwarz.

As a résumé of the first study we conclude that 1) gradual weighting with a smooth shape function yields a decisive improvement over arithmetic weighting, and 2) increasing the overlap with growing  $p$  is crucial for robustness.

### 4.2 Robustness and Efficiency

Next we investigate robustness with respect to the mesh size and aspect ratio. First, MG with one pre-smoothing is applied on uniform meshes consisting of  $4^2$  to  $1024^2$  elements with  $p$  ranging from 4 to 32 and up to four million unknowns. Table 4 compiles the results for quintically weighted and multiplicative Schwarz smoothers with overlap  $n_{O,l} = \lceil p_l/8 \rceil$ . Except in coarse quadrangulations, where periodicity can induce interference effects, the convergence characteristics are virtually independent of the number of elements  $n_E$ . The convergence rate  $\bar{r}$  shows a moderate growth for increasing order  $p$  and is similar for both smoothers, with a slight advantage for the weighted additive Schwarz method. As a consequence, the equivalent number of operator applications required for reducing the residual by an order of magnitude drops almost to one third when increasing  $p$  from 4 to 32 and, thus, mitigates the higher operator cost per DOF.

**Table 4** Robustness with respect to problem size: MG using additive and multiplicative Schwarz smoothers with overlap  $n_{o,l} = \lceil pl/8 \rceil$

$p$	$\sqrt{n_E}$	MG(1,0), add, $w_5$			MG(1,0), mult		
		$\bar{r}$	$n_{10}$	$\bar{\omega}_1$	$\bar{r}$	$n_{10}$	$\bar{\omega}_1$
4	32	1.17	9	9.2	0.87	12	12.4
	64	1.17	9	9.3	0.86	12	12.6
	128	1.17	9	9.3	0.85	12	12.7
	256	1.17	9	9.3	0.85	12	12.7
8	16	1.30	8	5.4	1.28	8	5.5
	32	1.29	8	5.4	1.26	8	5.5
	64	1.29	8	5.4	1.26	8	5.5
	128	1.28	8	5.4	1.26	8	5.5
16	8	1.33	8	5.1	1.44	7	4.7
	16	1.37	8	4.9	1.42	8	4.8
	32	1.36	8	5.0	1.46	7	4.6
	64	1.36	8	5.0	1.46	7	4.6
32	4	1.90	6	3.5	1.65	7	4.0
	8	1.58	7	4.2	1.59	7	4.2
	16	1.87	6	3.6	1.63	7	4.1
	32	1.93	6	3.4	1.64	7	4.0
	64	1.93	6	3.4	1.65	7	4.0

In the second test we fixed the mesh to  $16 \times 16$  elements of order  $p = 16$ , but increased the aspect ratio  $AR = \Delta x/\Delta y$  by enlarging the domain into the  $x$  direction, i.e.,  $\Omega = [0, 2AR] \times [0, 2]$ . Table 5 reports the results for MG and MGCG using additive weighted Schwarz with  $w_5$ ,  $n_{o,l} = \lceil pl/8 \rceil$  and one pre-smoothing step. As expected, the stand-alone MG performs well for small aspect ratios, but degrades for  $AR > 2$ . MGCG is slightly less efficient than MG for  $AR \leq 2$ , but proves more robust at higher aspect ratios. At  $AR = 8$  it converges approximately twice as fast as MG.

While these observations hold almost uniformly for all orders  $p$  considered, it remains to investigate the impact of solver parameters such as smoothing steps and overlap. Figure 3 presents selected results of the corresponding study for  $p = 16$  and aspect ratios  $AR = 1$  to 16. In particular we considered several variants of MGCG(1,1), each applying one pre- and one post-smoothing. In one case, indicated by “var”, we employed a variable V-cycle in which the number of smoothing steps doubles with each coarser level, i.e.  $n_{s1,l} = n_{s2,l} = 2^{L-l}$ . The study included quintically weighted additive (“add,  $w_5$ ”) as well as multiplicative (“mult”) Schwarz smoothers with a level-dependent overlap of  $n_{o,l} = \lceil pl/8 \rceil$ . Additionally we tested multiplicative Schwarz with  $n_o = 0$  and  $n_{s1} = n_{s2} = 2$ , which corresponds to the method of [13], and the arithmetically averaged additive Schwarz smoother using a constant overlap of  $n_o = 1$ . Figure 3a depicts the achieved convergence rates. Compared to the case of only one smoothing, the additional post-smoothing raises  $\bar{r}$  by a factor between 1.5 and 2, which is well in the expected range. Switching to multiplicative Schwarz yields an even higher gain for increasing aspect ratios. A similar effect is achieved using additive Schwarz with the variable V-cycle. MGCG(2,2) with zero overlap attains a convergence rate similar to MGCG(1,0) with level-dependent overlap. The arithmetically averaged Schwarz method with two smoothing steps falls about two thirds behind the quintically weighted method with only one smoothing for  $AR = 1$ , but gains a slight advantage over the latter for higher aspect ratios.

**Table 5** Robustness with respect to aspect ratio: MG versus MGCG using additive Schwarz with  $w_5$  and overlap  $n_{o,l} = \lceil pl/8 \rceil$

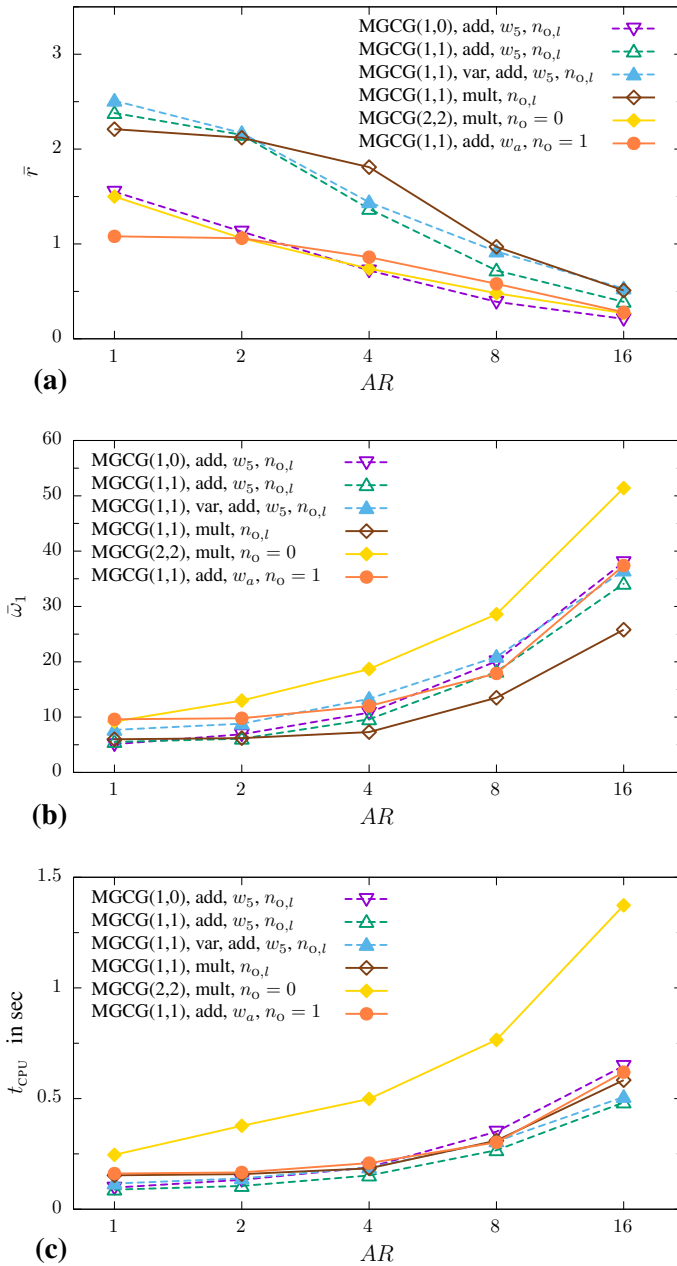
$p$	$AR$	MG(1,0),add, $w_5$			MGCG(1,0),add, $w_5$		
		$\bar{r}$	$n_{10}$	$\bar{\omega}_1$	$\bar{r}$	$n_{10}$	$\bar{\omega}_1$
4	1	1.17	9	9.2	1.30	8	9.3
	2	0.99	11	11.0	1.10	10	11.0
	4	0.39	26	27.8	0.59	17	20.5
	8	0.12	85	91.2	0.28	36	42.5
8	1	1.30	8	5.4	1.33	8	6.1
	2	0.86	12	8.2	1.03	8	7.9
	4	0.43	24	16.3	0.65	16	12.5
	8	0.16	63	43.9	0.34	30	23.8
16	1	1.37	8	4.9	1.55	7	5.1
	2	0.95	11	7.1	1.14	9	6.9
	4	0.50	20	13.5	0.72	14	10.8
	8	0.17	59	39.9	0.39	26	20.1
32	1	1.87	6	3.6	2.01	5	3.8
	2	1.23	9	5.4	1.42	8	5.4
	4	0.65	16	10.3	0.83	12	9.2
	8	0.22	46	30.4	0.44	23	17.5

As the convergence rate does not account for the cost, it is of limited value when comparing methods of different computational complexity. A better measure is the equivalent number of operator applications required for reducing the residual by one order of magnitude,  $\bar{\omega}_1$ , which is shown in Fig. 3b. In this metric, the multiplicative MGCG(1,1) with level-dependent overlap performs best, especially for higher aspect ratios. It is followed by its additive counterpart with quintic weighting, which is at level for  $AR \leq 2$ , but needs ca 34 instead of 26 operator applications for  $AR = 16$ . The comparison also reveals that the benefit of the variable V-cycle is lost due to the higher computational complexity. Generally, the influence of smoothing and overlap parameters lessens with increasing aspect ratio (exempting the case of  $n_o = 0$ ), which indicates that the role of the conjugate gradient method gets more important.

Figure 3c depicts the runtimes measured on a 3.1 GHz Intel Core i7-5557U CPU. Note that MGCG(1,1) with quintic weighting attained the best performance despite its higher operation count in comparison to MGCG(1,1) with multiplicative Schwarz. This is because the additive Schwarz method evaluates the residual for all elements at once, yielding a single, highly efficient BLAS3 operation. In contrast, multiplicative Schwarz requires a series of local residual updates, which is harder to optimize. Consistently, the multiplicative MGCG(2,2) with  $n_o = 0$  remains the least efficient method for all aspect ratios. Compared to MGCG(1,1) with  $n_o = 1$  and arithmetic weighting, the method with  $n_{o,l} = \lceil pl/8 \rceil$  and quintic weighting succeeds twice as fast for  $AR = 1$  and still gains 23% at  $AR = 16$ . Though other choices may yield even better performance, the study documents that the method is not too sensitive to parameter variations, such that only minor improvements can be expected.

### 4.3 Variable Diffusion

Since many problems in physics involve variable coefficients, it is interesting to explore if the multigrid method is capable to retain its efficiency in such applications. As an example we consider the diffusion equation



**Fig. 3** Performance of MGCG for different aspect ratios. All cases use  $16 \times 16$  elements of order  $p = 16$  and  $n_{o,l} = \lceil pl/8 \rceil$ , if not specified otherwise. **a** Average logarithmic convergence rate. **b** Operator applications required for  $10^1$  residual reduction. **c** Solver runtime for  $10^{10}$  residual reduction

$$-\nabla \cdot (\nu \nabla u) = f \tag{16}$$

with variable diffusivity  $\nu$  in the periodic domain  $\Omega = [0, AR] \times [0, 1]$ . From a physical perspective it seems reasonable that the solution and the diffusivity vary on similar scales. Following this idea we set  $u = \sin(2\pi x) \sin(2\pi y)$  and

$$\nu = 1 + \hat{\nu} \sin(2\pi(x - s)) \sin(2\pi(y - s)), \tag{17}$$

where  $\hat{\nu}$  is the amplitude and  $s$  the shift of the diffusivity fluctuation. According to preliminary studies, a non-zero shift poses an additional difficulty to the solver. Taking this into account,  $s = 0.2$  is chosen in all tests reported below. The source is analytically computed from  $f = -(\nu \nabla^2 u + \nabla \nu \cdot \nabla u)$ .

Discretization using rectangular spectral elements yields the linear system

$$B u = f, \tag{18}$$

where  $B(\nu)$  represents the discrete diffusion operator or, equivalently,

$$B \Delta u = f - B \tilde{u} = \tilde{r}, \tag{19}$$

for the correction  $\Delta u$  to a given approximation  $\tilde{u}$ . Application of the Schwarz method described in Sect. 3.1 leads to the local correction equation

$$B_{s,s} \Delta u_s = r_s, \tag{20}$$


where  $B_{s,s}$  and  $r_s$  are the diffusion operator and, respectively, the residual restricted to the subdomain  $\Omega_s$ . In comparison to the subdomain problem for Poisson case (7), Eq. (20) is more expensive to solve, because the fast diagonalization technique is no longer applicable. Yet, the smoothing property is more important for multigrid than accurate solution of the subproblems. This motivates the reintroduction of the discrete Laplacian by approximating the restricted diffusion operator on the left side of (20) by  $B_{s,s} \approx \bar{\nu}_s A_{s,s}$ , where the diffusivity  $\bar{\nu}_s$  is assumed to be constant in  $\Omega_s$ . For simplicity,  $\bar{\nu}_s$  is set to the average of  $\nu$  over the embedded element. The correction can then be approximated as

$$\Delta u_s \approx \frac{1}{\bar{\nu}_s} A_{s,s}^{-1} r_s, \tag{21}$$


where, again,  $A_{s,s}^{-1}$  stands for the application of the factored inverse obtained from fast diagonalization. As a result, the solution techniques developed in Sect. 3 can be utilized with no change except for the residual evaluation.

The performance of the scheme was studied in two test series. In the first series, the aspect ratio was fixed to  $AR = 1$  and the domain  $\Omega = [0, 1]^2$  decomposed into  $8^2$  square elements of order  $p = 16$ . The diffusivity fluctuation amplitude  $\hat{\nu}$  was gradually increased from 0 to 0.9, where the latter corresponds to variations of the magnitude up to 90%. Figure 4 shows the measured convergence rates for MG und MGCG using one pre- and post-smoothing based on additive Schwarz with a level-dependent overlap of  $n_{o,l} = \lceil p_l/8 \rceil$  and quintic weighting. The results indicate that MG retains its efficiency up to fluctuation amplitudes of about 30%, but then degrades with rising  $\hat{\nu}$ . As expected, Krylov acceleration improves the robustness, such that MGCG achieves  $\bar{r} = 0.91$  for  $\hat{\nu} = 0.9$ , which is nearly twice the convergence rate obtained with MG. Compared to  $\hat{\nu} = 0$ , this corresponds to an increase of the cycle count and, hence, in runtime, by a factor of just 2.2. A similar behavior was observed for polynomial orders  $p = 4, 8$  and 32.

In the second test series, we increased the domain extension in the  $x$ -direction, while keeping the diffusivity fluctuation amplitude  $\hat{\nu}$  at a constant level. The number of elements



**Fig. 4** MG and MGCG convergence rates for different diffusivity fluctuation amplitudes. Discretization is based on an isotropic mesh comprising  $8^2$  elements of order  $p = 16$ . One pre- and post-smoothing with an overlap of  $n_{o,l} = \lceil p_l/8 \rceil$  and quintic weighting were applied in both cases



**Fig. 5** MGCG convergence rate for variable diffusivity on anisotropic meshes with increasing aspect ratio using a variable V-cycle with one pre- and post-smoothing, overlap  $n_{o,l} = \lceil p_l/2 \rceil$  and quintic weighting

is fixed and identical in both directions, such that the element aspect ratio equals  $AR$ . For achieving a robustness similar to the Poisson case it proved necessary to increase the subdomain overlap with growing  $\hat{\nu}$ . Figure 5 shows the convergence rates  $\bar{r}$  for MGCG using a variable V-cycle and additive Schwarz smoothing for an amplitude of 90%, which represents the most challenging test in the series. Comparing the results for  $p = 8$  and  $p = 16$  one observes that  $\bar{r}$  strongly depends on the quadrangulation, but only marginally on the polynomial order. Using a finer mesh yields considerably higher convergence rates and better robustness. Orders 4 and 32 fit nicely into this picture, but are not shown for clarity. The congruence of different orders using the same mesh suggests, that the performance depends on how well the diffusivity fluctuation is resolved by the element mean values adopted for  $\bar{v}_s$ . This presents a possible limitation of the approach, which needs further consideration in subsequent work. Nevertheless, the study demonstrates the suitability of the proposed method for problems involving variable diffusivity, as long as the latter is sufficiently resolved.

## 5 Conclusions

We have developed a nonuniformly weighted additive Schwarz method acting as the smoother in multigrid solvers for the spectral element discretization of the Poisson equation in  $\mathbb{R}^2$ . The method generalizes the Schwarz/multigrid method proposed in [19] and was inspired from weighting techniques devised in [11]. In each step, it determines the solution for a subdomain corresponding to an extended element region. These local solutions are blended according to a polynomial shape function which features a smooth transition from zero at the border toward one in the core of the subdomain. As an alternative we considered a multiplicative Schwarz method with no weighting required. Both Schwarz methods were integrated in a polynomial multigrid method which, in turn, was embedded in a preconditioned CG method.

The performance of these methods was assessed in a series of numerical experiments with ansatz orders  $p$  ranging from 4 to 32 and up to  $n_E = 256^2$  elements of aspect ratios  $AR$  from 1 to 16. For unit-aspect ratio elements the proposed weighting improved the logarithmic MG convergence rate and reduced the cost by a factor of 1.5–3 in comparison to the original method. The study indicates that for robustness the subdomain overlap has to be bounded, i.e., the number  $n_O$  of node layers adopted from neighbor elements must grow with increasing order. Thus, with MG, the number of layers varies from level to level. A reasonable choice is to use an overlap of  $\lceil p_l/8 \rceil$  layers, where  $p_l$  denotes the polynomial order on level  $l$ . The resulting multigrid method is robust with respect to the mesh size, i.e.  $p$  and  $n_E$ , but degrades with increasing aspect ratio. This behavior can be mitigated by Krylov subspace acceleration: Using MG as a preconditioner for the inexact conjugate gradient method [9] improves the convergence rate for higher aspect ratios considerably.

Finally, it has been shown that the proposed multigrid method is easily adapted and well suited for solving diffusion problems with varying coefficients, provided the mesh is fine enough to approximate diffusivity fluctuations by element mean values. Improving the treatment of variable coefficients and extending the approach to three space dimensions are topics of ongoing and future work.

**Acknowledgements** Funding by German Research Foundation (DFG) in frame of the Project STI 157/4-1 is gratefully acknowledged.

## References

1. Bastian, P., Blatt, M., Scheichl, R.: Algebraic multigrid for discontinuous Galerkin discretizations of heterogeneous elliptic problems. *Numer. Linear Algebra Appl.* **19**(2), 367–388 (2012)
2. Blaheta, R.: GPCG-generalized preconditioned CG method and its use with non-linear and non-symmetric displacement decomposition preconditioners. *Numer. Linear Algebra Appl.* **9**(6–7), 527–550 (2002)
3. Bouwmeester, H., Dougherty, A., Knyazev, A.V.: Nonsymmetric preconditioning for conjugate gradient and steepest descent methods. *Procedia Comput. Sci.* **51**, 276–285 (2015). ISSN 1877-0509
4. Bramble, J.: *Multigrid Methods*. Pitman Research Notes in Mathematics Series, vol. 294. Longman Scientific & Technical, Harlow (1995)
5. Brenner, S.C., Zhao, J.: Convergence of multigrid algorithms for interior penalty methods. *Appl. Num. Anal. Comp. Math.* **2**(1), 3–18 (2005)
6. Couzy, W., Deville, M.O.: A fast Schur complement method for the spectral element discretization of the incompressible Navier–Stokes equations. *J. Comput. Phys.* **116**, 135–142 (1995)
7. Deville, M.O., Fischer, P.F., Mund, E.H.: *High-Order Methods for Incompressible Fluid Flow*, vol. 1. Cambridge University Press, Cambridge (2002)
8. Fischer, P.F., Lottes, J.W.: Hybrid Schwarz-multigrid methods for the spectral element method: Extensions to Navier–Stokes. In: Kornhuber, R., Périaux, J., Widlund, O., Hoppe, R., Pironneau, O., Xu, J. (eds.) *Domain Decomposition Methods in Science and Engineering Series*, pp. 35–49. Springer, Berlin (2004)

9. Golub, G.H., Ye, Q.: Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM J. Sci. Comput.* **21**(4), 1305–1320 (1999). ISSN 10648275
10. Guermund, J.L., Mineev, P., Shen, J.: An overview of projection methods for incompressible flows. *Comput. Methods Appl. Mech. Eng.* **195**, 6011–6045 (2006)
11. Haupt, L., Stiller, J., Nagel, W.: A fast spectral element solver combining static condensation and multigrid techniques. *J. Comput. Phys.* **255**, 384–395 (2013)
12. Heinrichs, W.: Line relaxation for spectral multigrid methods. *J. Comput. Phys.* **77**, 166–182 (1988)
13. Janssen, B., Kanschat, G.: Adaptive multilevel methods with local smoothing for  $H^1$ - and  $H^{\text{curl}}$ -conforming high order finite element methods. *SIAM J. Sci. Comput.* **33**(4), 2095–2114 (2011)
14. Kaasschieter, E.F.: Preconditioned conjugate gradients for solving singular systems. *J. Comput. Appl. Math.* **24**, 265–275 (1988)
15. Kanschat, G.: Multilevel methods for discontinuous Galerkin FEM on locally refined meshes. *Comput. Struct.* **82**, 2437–2445 (2004)
16. Karniadakis, G.E., Sherwin, S.J.: *Spectral/hp Element Methods for Computational Fluid Dynamics*, 2nd edn. Oxford University Press, Oxford (2005)
17. Kraus, J.K., Tomar, S.K.: A multilevel method for discontinuous Galerkin approximation of three-dimensional anisotropic elliptic problems. *Numer. Linear Algebra Appl.* **15**(5), 417–438 (2008)
18. Loisel, S., Nabben, R., Szyld, D.: On hybrid multigrid-Schwarz algorithms. *J. Sci. Comput.* **36**(2), 165–175 (2008)
19. Lottes, J.W., Fischer, P.F.: Hybrid multigrid/Schwarz algorithms for the spectral element method. *J. Sci. Comput.* **24**, 45–78 (2005)
20. Lynch, R.E., Rice, J.R., Thomas, D.H.: Direct solution of partial difference equations by tensor product methods. *Numer. Math.* **6**, 185–199 (1964)
21. Maday, Y., Munoz, R.: Spectral element multigrid. II. Theoretical justification. *J. Sci. Comput.* **3**, 323–353 (1988)
22. Mitchell, W.F.: The hp-multigrid method applied to hp-adaptive refinement of triangular grids. *Numer. Linear Algebr.* **17**, 211–228 (2010)
23. Notay, Y.: Flexible conjugate gradients. *SIAM J. Sci. Comput.* **22**(4), 1444–1460 (2000)
24. Olson, L.: Algebraic multigrid preconditioning of high-order spectral elements for elliptic problems on a simplicial mesh. *SIAM J. Sci. Comput.* **29**(5), 2189–2209 (2007)
25. Pasquetti, R., Rapetti, F.: p-Multigrid method for Fekete-Gauss spectral element approximations of elliptic problems. *Commun. Comput. Phys.* **5**(5), 667–682 (2009)
26. Rønquist, E., Patera, A.: Spectral element multigrid. I. Formulation and numerical results. *J. Sci. Comput.* **2**, 389–406 (1987)
27. Trottenberg, U., Oosterlee, C.W., Schüller, A.: *Multigrid*. Academic Press, Oxford (2000)
28. Varga, R.S.: *Matrix Iterative Analysis*. Springer Series in Computational Mathematics, vol. 27, 2nd ed. Springer, Berlin (2000)
29. Wang, Z.J., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, R., Hillewaert, K., Huynh, H., Kroll, N., May, G., Persson, P.-O., van Leer, B., Visbal, M.: High-order CFD methods: current status and perspective. *Int. J. Numer. Methods Fluids* **72**(8), 811–845 (2013). ISSN 1097-0363