
Preconditioning High–Order Discontinuous Galerkin Discretizations of Elliptic Problems

Paola F. Antonietti¹ and Paul Houston^{2*}

¹ MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano ITALY. paola.antonietti@polimi.it

² School of Mathematical Sciences, University of Nottingham, University Park, Nottingham, NG7 2RD, UK. Paul.Houston@nottingham.ac.uk

1 Introduction

In recent years, attention has been devoted to the development of efficient iterative solvers for the solution of the linear system of equations arising from the discontinuous Galerkin (DG) discretization of a range of model problems. In the framework of two level preconditioners, scalable non-overlapping Schwarz methods have been proposed and analyzed for the h -version of the DG method in the articles [1, 2, 6, 7, 9]. Recently, in [3] it has been proved that the non-overlapping Schwarz preconditioners can also be successfully employed to reduce the condition number of the stiffness matrices arising from a wide class of high–order DG discretizations of elliptic problems. In this article we aim to validate the theoretical results derived in [3] for the multiplicative Schwarz preconditioner and for its symmetrized variant by testing their numerical performance.

2 Model Problem and DG Discretization

In this section we introduce the model problem under consideration and its DG approximation, working, for the sake of simplicity, with the SIPG formulation proposed in [4].

We consider, for simplicity, the weak formulation of the Poisson problem with homogeneous Dirichlet boundary conditions: find $\mathcal{U} \in H_0^1(\Omega)$ such that

$$(\nabla \mathcal{U}, \nabla v)_\Omega = (f, v)_\Omega \quad \forall v \in H_0^1(\Omega), \tag{1}$$

where Ω is a bounded polygonal domain in \mathbb{R}^d , $d = 2, 3$, $f \in L^2(\Omega)$ is a given source term and $(\cdot, \cdot)_\Omega$ is the standard inner product in $[L^2(\Omega)]^d$.

* PH acknowledges the financial support of the EPSRC under the grant EP/H005498.

Let \mathcal{T}_h be a shape-regular, not necessarily matching partition of Ω into disjoint open elements \mathcal{K} (with diameter $h_{\mathcal{K}}$), where each \mathcal{K} is the affine image of a fixed master element $\widehat{\mathcal{K}}$, i.e., $\mathcal{K} = F_{\mathcal{K}}(\widehat{\mathcal{K}})$, where $\widehat{\mathcal{K}}$ is either the open unit d -simplex or the d -hypercube in \mathbb{R}^d , $d = 2, 3$. We define the mesh-size h by $h := \max_{\mathcal{K} \in \mathcal{T}_h} h_{\mathcal{K}}$, and assume that \mathcal{T}_h satisfies a *bounded local variation* property: for any pair of neighboring elements $\mathcal{K}_1, \mathcal{K}_2 \in \mathcal{T}_h$, $h_{\mathcal{K}_1} \approx h_{\mathcal{K}_2}$.

For a given approximation order $p \geq 1$, we define the DG space

$$V_{h,p} := \{v \in L^2(\Omega) : v|_{\mathcal{K}} \circ F_{\mathcal{K}} \in \mathcal{M}^p(\widehat{\mathcal{K}}) \forall \mathcal{K} \in \mathcal{T}_h\},$$

where $\mathcal{M}^p(\widehat{\mathcal{K}})$ is either the space of polynomials of degree at most p on $\widehat{\mathcal{K}}$, if $\widehat{\mathcal{K}}$ is the reference d -simplex, or the space of polynomials of degree at most p in each variable on $\widehat{\mathcal{K}}$, if $\widehat{\mathcal{K}}$ is the reference d -hypercube.

Next, for any internal face $\bar{F} = \partial\mathcal{K}^+ \cap \partial\mathcal{K}^-$ shared by two adjacent elements \mathcal{K}^\pm , with outward unit normal vectors \mathbf{n}^\pm , respectively, we define

$$\begin{aligned} [[\boldsymbol{\tau}]] &:= \boldsymbol{\tau}^+ \cdot \mathbf{n}^+ + \boldsymbol{\tau}^- \cdot \mathbf{n}^-, & [[v]] &:= v^+ \mathbf{n}^+ + v^- \mathbf{n}^-, \\ \{\{\boldsymbol{\tau}\}\} &:= (\boldsymbol{\tau}^+ + \boldsymbol{\tau}^-)/2, & \{\{v\}\} &:= (v^+ + v^-)/2, \end{aligned}$$

where $\boldsymbol{\tau}^\pm$ and v^\pm denote the traces on $\partial\mathcal{K}^\pm$ taken from the interior of \mathcal{K}^\pm of the (sufficiently regular) functions $\boldsymbol{\tau}$ and v , respectively (cf. [5]). On a boundary face $\bar{F} = \partial\mathcal{K} \cap \partial\Omega$, we set $[[\boldsymbol{\tau}]] := \boldsymbol{\tau} \cdot \mathbf{n}$, $[[v]] := v \mathbf{n}$, $\{\{\boldsymbol{\tau}\}\} := \boldsymbol{\tau}$, and $\{\{v\}\} := v$.

We collect all interior (respectively, boundary) faces in the set \mathcal{F}_h^I (respectively, \mathcal{F}_h^B), define $\mathcal{F}_h := \mathcal{F}_h^I \cup \mathcal{F}_h^B$, and introduce on $V_{h,p} \times V_{h,p}$ the following bilinear form

$$\begin{aligned} \mathcal{A}(u, v) &:= \sum_{\mathcal{K} \in \mathcal{T}_h} \int_{\mathcal{K}} \nabla u \cdot \nabla v \, dx + \sum_{\mathcal{K} \in \mathcal{T}_h} \int_{\mathcal{K}} \nabla u \cdot \mathcal{R}([[v]]) \, dx \\ &\quad + \sum_{\mathcal{K} \in \mathcal{F}_h} \int_{\mathcal{K}} \mathcal{R}([[u]]) \cdot \nabla v \, dx + \sum_{F \in \mathcal{F}_h} \int_F \alpha \frac{p^2}{|F|} [[u]] \cdot [[v]] \, ds, \end{aligned}$$

where $\alpha > 0$ is a parameter at our disposal. The lifting operator $\mathcal{R}(\cdot)$ is defined as: $\mathcal{R}(\boldsymbol{\tau}) := \sum_{F \in \mathcal{F}_h} r_F(\boldsymbol{\tau})$, where $r_F : [L^2(F)]^d \rightarrow [V_{h,p}]^d$ is given by

$$\int_{\Omega} r_F(\boldsymbol{\tau}) \cdot \boldsymbol{\eta} \, dx := - \int_F \boldsymbol{\tau} \cdot \{\{\boldsymbol{\eta}\}\} \, ds \quad \forall \boldsymbol{\eta} \in [V_{h,p}]^d \quad \forall F \in \mathcal{F}_h.$$

The DG discretization of problem (1) reads:

$$\text{Find } u \in V_{h,p} \text{ such that } \quad \mathcal{A}(u, v) = \int_{\Omega} f v \, dx \quad \forall v \in V_{h,p}. \quad (2)$$

Let $\boldsymbol{\varphi}_j$, $j = 1, \dots, N_h^p := \dim(V_{h,p})$, be a set of basis functions that span $V_{h,p}$, then (2) can be written in the following equivalent form: Find $\mathbf{u} \in \mathbb{R}^{N_h^p}$ such that $\mathbf{A}\mathbf{u} = \mathbf{f}$, where here (and in the following) we use the bold notation to denote the spaces of

degrees of freedom (vectors) and discrete linear operators (matrices). The following result provides an estimate for the spectral condition number of \mathbf{A} ; we refer to [3] for the proof.

Proposition 1 ([3]). *For a set of basis functions which are orthonormal on the reference element $\widehat{\mathcal{K}} \subset \mathbb{R}^d$, $d = 2, 3$, the condition number $\kappa(\mathbf{A})$ of the stiffness matrix \mathbf{A} can be bounded by*

$$\kappa(\mathbf{A}) \lesssim \alpha \frac{p^4}{h^2}.$$

Remark 1. We are working, for the sake of simplicity, with the SIPG formulation proposed in [4], but the results shown in Proposition 1 and in Theorem 1 below also hold for a wide class of DG methods; we refer to [3] for details.

3 Two Level Non-overlapping Schwarz Preconditioners

In this section we introduce the non-overlapping Schwarz preconditioners.

Subdomain partition. We decompose the domain Ω into N non-overlapping subdomains Ω_i , i.e., $\overline{\Omega} = \cup_{i=1}^N \overline{\Omega}_i$. Next, we consider two levels of *nested* partitions of the domain Ω : (i) a coarse partition \mathcal{T}_H (with mesh-size H); (ii) a fine partition \mathcal{T}_h (with mesh-size h). We will suppose that the subdomain partition does not cut any element of \mathcal{T}_H (and therefore of \mathcal{T}_h).

Local solvers. For $i = 1, \dots, N$, we define the local DG spaces as

$$V_{h,p}^i := \{v \in L^2(\Omega_i) : v|_{\mathcal{K}} \circ F_{\mathcal{K}} \in \mathcal{M}^p(\widehat{\mathcal{K}}) \quad \forall \mathcal{K} \in \mathcal{T}_h, \mathcal{K} \subset \Omega_i\}.$$

Denoting by $R_i^T : V_{h,p}^i \longrightarrow V_{h,p}$ the classical injection operator from $V_{h,p}^i$ to $V_{h,p}$, the local solvers $\mathcal{A}_i : V_{h,p}^i \times V_{h,p}^i \longrightarrow \mathbb{R}$ are defined as

$$\mathcal{A}_i(u_i, v_i) := \mathcal{A}(R_i^T u_i, R_i^T v_i) \quad \forall u_i, v_i \in V_{h,p}^i, \quad i = 1, \dots, N. \tag{3}$$

Coarse solver. For an integer $0 \leq q \leq p$, we define the coarse space $V_{H,q}^0$ as

$$V_{H,q}^0 := \{v \in L^2(\Omega) : v|_{\mathcal{D}} \circ F_{\mathcal{D}} \in \mathbb{M}^{q,q}(\widehat{\mathcal{K}}) \quad \forall \mathcal{D} \in \mathcal{T}_H\},$$

and the *coarse solver* $\mathcal{A}_0 : V_{H,q}^0 \times V_{H,q}^0 \longrightarrow \mathbb{R}$ as

$$\mathcal{A}_0(u_0, v_0) := \mathcal{A}(R_0^T u_0, R_0^T v_0) \quad \forall u_0, v_0 \in V_{H,q}^0, \tag{4}$$

where $R_0^T : V_{H,q}^0 \longrightarrow V_{h,p}$ is the classical injection operator from $V_{H,q}^0$ to $V_{h,p}$.

Let the *local* projection operators be defined as

$$\begin{aligned} \tilde{P}_i : V_{h,p} &\rightarrow V_{h,p}^i : \mathcal{A}_i(\tilde{P}_i u, R_i^T v_i) := \mathcal{A}(u, R_i^T v_i) \quad \forall v_i \in V_{h,p}^i, \quad i = 1, \dots, N, \\ \tilde{P}_0 : V_{h,p} &\rightarrow V_{H,q}^0 : \mathcal{A}_0(\tilde{P}_0 u, R_0^T v_0) := \mathcal{A}(u, R_0^T v_0) \quad \forall v_0 \in V_{H,q}^0, \end{aligned} \tag{5}$$

and define the projection operators as $P_i := R_i^T \tilde{P}_i : V_{h,p} \rightarrow V_{h,p}$, $i = 0, 1, \dots, N$. The multiplicative Schwarz operator and its symmetrized variant are then defined as

$$P_{\text{mu}} := I - (I - P_N)(I - P_{N-1}) \cdots (I - P_0), \tag{6}$$

$$P_{\text{mu}}^S := I - (I - P_0)^T \cdots (I - P_N)^T (I - P_N) \cdots (I - P_0), \tag{7}$$

respectively (cf. [10]). The Schwarz method consists in solving either $P_{\text{mu}} u = g_{\text{mu}}$ or $P_{\text{mu}}^S u = g_{\text{mu}}^S$, for suitable right hand sides g_{mu} and g_{mu}^S , respectively. It can be shown that the operator defined in (7) is symmetric and positive definite; we therefore consider the conjugate gradient (CG) algorithm for the solution of $P_{\text{mu}}^S u = g_{\text{mu}}^S$. An estimate of the condition number of P_{mu}^S is

$$\kappa(P_{\text{mu}}^S) := \frac{\lambda_{\max}(P_{\text{mu}}^S)}{\lambda_{\min}(P_{\text{mu}}^S)},$$

where $\lambda_{\max}(P_{\text{mu}}^S)$ and $\lambda_{\min}(P_{\text{mu}}^S)$ are the extremal eigenvalues of the operator P_{mu}^S . On the other hand, the multiplicative operator P_{mu} is non-symmetric; we therefore consider a Richardson iteration applied to $P_{\text{mu}} u = g_{\text{mu}}$, and show that the norm of the error propagation operator $E_{\text{mu}} := (I - P_N)(I - P_{N-1}) \cdots (I - P_0)$ is strictly less than one, i.e.,

$$\|E_{\text{mu}}\|_{\mathcal{A}}^2 := \sup_{\substack{v \in V_{h,p} \\ v \neq 0}} \frac{\mathcal{A}(E_{\text{mu}} v, E_{\text{mu}} v)}{\mathcal{A}(v, v)} < 1,$$

and therefore a Richardson iteration applied to the preconditioned system converges. The following result provides a bound for the norm of the error propagation operator of the multiplicative Schwarz operator, and for the condition number of the symmetrized Schwarz operator (we refer to [3] for the proof).

Theorem 1 ([3]). *There exists constants $C_1, C_2 \geq 1$, independent of the mesh-size and the polynomial degree, such that*

$$\|E_{\text{mu}}\|_{\mathcal{A}}^2 \leq 1 - \frac{h}{C_1 \alpha p^2 H}, \quad \kappa(P_{\text{mu}}^S) \leq C_2 \alpha p^2 \frac{H}{h}.$$

Theorem 1 also guarantees that the multiplicative Schwarz method can be accelerated with the GMRES iterative solver. Indeed, according to [8], the GMRES method applied to the preconditioned system $P_{\text{mu}} u = g_{\text{mu}}$ does not stagnate (i.e., the iterative method makes some progress in reducing the residual at each iteration step) provided that: (i) $\|P_{\text{mu}}\|_{\mathcal{A}}$ is bounded; (ii) the symmetric part of P_{mu} is positive definite, i.e., there exists $c_p > 0$ such that $\mathcal{A}(v, P_{\text{mu}} v) \geq c_p \mathcal{A}(v, v)$ for all $v \in V_{h,p}$. Condition (i) follows directly from the definition of P_{mu} and Theorem 1: $\|P_{\text{mu}}\|_{\mathcal{A}} = \|I - E_{\text{mu}}\|_{\mathcal{A}} \leq 1 + \|E_{\text{mu}}\|_{\mathcal{A}} < 2$. To prove condition (ii), it can be shown that




Fig. 1. Initial Cartesian and triangular coarse and fine grids on a 16 subdomain partition. (a) Initial coarse grids (mesh-size H_0) and (b) initial fine grids (mesh-size h_0)

$$\mathcal{A}(P_{\text{mu}}v, v) = \mathcal{A}(v, v) - \mathcal{A}(E_{\text{mu}}v, v) \geq (1 - \|E_{\text{mu}}\|_{\mathcal{A}}) \mathcal{A}(v, v).$$

Therefore, condition (ii) holds true with $c_p = 1 - \|E_{\text{mu}}\|_{\mathcal{A}}$ which is positive due to Theorem 1.

4 Numerical Results

In this section we present some numerical experiments to highlight the practical performance of the multiplicative and symmetrized non-overlapping Schwarz preconditioners. From the algebraic point of view, the Schwarz operators (6) and (7) can be written as the product of a suitable preconditioner, namely \mathbf{B}_{mu} , $\mathbf{B}_{\text{mu}}^{\text{S}}$, respectively, and \mathbf{A} . Indeed, the local components can be constructed as $\mathbf{A}_i = \mathbf{R}_i \mathbf{A} \mathbf{R}_i^T$, see (3) for $i = 1 \dots, N$, and (4) for $i = 0$. From the definition (5) of the local projection $\tilde{\mathbf{P}}_i = \mathbf{A}_i^{-1} \mathbf{R}_i \mathbf{A}$, and therefore $\mathbf{P}_i = \mathbf{R}_i^T \tilde{\mathbf{P}}_i = \mathbf{R}_i^T \mathbf{A}_i^{-1} \mathbf{R}_i \mathbf{A}$. In practice, only the action of the preconditioner on a vector is needed. Algorithm 2 shows how to compute the action of \mathbf{B}_{mu} on a vector $\mathbf{x} \in \mathbb{R}^{N_p}$. Throughout this section we have set the

Algorithm 2 $\mathbf{z} = \mathbf{B}_{\text{mu}} \mathbf{x}$

```

 $\mathbf{z} = \mathbf{R}_0^T \mathbf{A}_0^{-1} \mathbf{R}_0 \mathbf{x}$ 
for  $i = 1 \rightarrow N$  do
     $\mathbf{z} = \mathbf{z} + \mathbf{R}_i^T \mathbf{A}_i^{-1} \mathbf{R}_i (\mathbf{x} - \mathbf{A} \mathbf{z})$ 
end for
    
```

penalty parameter $\alpha := 10$ (see (2)). We consider a subdomain partition consisting of $N = 16$ squares, and consider the initial Cartesian and unstructured triangular partitions shown in Fig. 1, and denote by H_0 and h_0 the corresponding initial coarse and fine mesh-sizes, respectively. We consider n successive global uniform refinements of these initial grids so that the resulting mesh-sizes are $H_n = H_0/2^n$ and $h_n = h_0/2^n$, with $n = 0, 1, 2, 3$, respectively. The (relative) tolerance is set equal to 10^{-9} (respectively, 10^{-6}) for the CG (respectively, GMRES) iterative solver. We first address the performance of the multiplicative Schwarz preconditioner by keeping the mesh fixed, and varying the polynomial approximation degree p . In Table 1 we compare the GMRES iteration counts for both the preconditioned and non-preconditioned (in

Table 1. GMRES iteration counts. Multiplicative Schwarz preconditioner with a piecewise constant coarse solver ($q = 0$). Unstructured triangular grids.

	$h = h_0/2$ $H = H_0$	$h = h_0/4$ $H = H_0$	$h = h_0/4$ $H = H_0/2$
$p = 1$	23 (94)	33 (199)	25 (199)
$p = 2$	45 (259)	64 (540)	49 (540)
$p = 3$	66 (470)	93 (996)	74 (996)
$p = 4$	85 (713)	124 (1546)	97 (1546)
$p = 5$	105 (1004)	153 (2187)	123 (2187)
$p = 6$	124 (1342)	183 (2924)	144 (2924)
$p = 7$	143 (1727)	209 (3742)	167 (3742)
$p = 8$	162 (2148)	235 (4673)	189 (4673)
$p - rate$	0.93 (1.63)	0.88 (1.66)	0.93 (1.66)

parenthesis) systems, for different polynomial approximation degrees and different mesh configurations. These results have been obtained on unstructured triangular grids (cf. Fig. 1). Comparing the iteration counts of the preconditioned systems with the unpreconditioned ones for a fixed p , it is clear that the proposed preconditioner is very efficient. Indeed, we observe a reduction in the number of iterations needed to achieve convergence of around one order of magnitude when the proposed preconditioner is employed. The last row of Table 1 shows the computed growth rate in the number of iterations: we observe that the number of iterations needed to obtain convergence increases linearly as a function of p for the preconditioned system of equations, whereas this quantity grows almost quadratically for the non-preconditioned problem. In Fig. 2 we report the condition number estimates of the symmetrized Schwarz operator and the corresponding iteration counts versus the polynomial degree p . The solid lines refer to the mesh configuration $h = h_0/2, H = H_0$, whereas the dashed lines refer to the mesh configuration $h = h_0/4, H = H_0/2$. This set of numerical experiments has been obtained on Cartesian meshes, employing a piecewise linear coarse solver. As predicted by the theoretical estimates, the condition number of the preconditioned system grows quadratically as a function of p . Moreover, we clearly observe that, for fixed p , by refining both the fine and the coarse grid, but keeping the ratio of the fine and coarse mesh-sizes constant, the condition number (and therefore the number of iterations needed to obtain convergence) remains constant.

Next, we consider the performance of the symmetrized Schwarz preconditioner when varying the coarse and fine mesh-size, and keeping the polynomial approximation degree p fixed. In Table 2 (left) we report the condition number estimates for the symmetrized Schwarz operator employing piecewise biquadratic elements ($p = 2$) and a piecewise constant coarse solver ($q = 0$); whereas, in Table 2 (right) the analogous results obtained with piecewise bicubic elements ($p = 3$) and a piecewise linear coarse solver ($q = 1$) are shown. We clearly observe that the condition number grows




Fig. 2. Condition number estimates of the symmetrized Schwarz operator and corresponding iteration counts versus the polynomial degree p on Cartesian grids for different discretization steps (solid line: $h = h_0/2, H = H_0$; dashed line $h = h_0/4, H = H_0/2$). Piecewise linear coarse solver

Table 2. Condition number estimates for the symmetrized Schwarz operator with $p = 2, q = 0$ (left) and $p = 3, q = 1$ (right). Cartesian grids.

$h \downarrow H \rightarrow$	H_0	$H_0/2$	$H_0/4$	$H_0/8$	H_0	$H_0/2$	$H_0/4$	$H_0/8$
h_0	5.32e2	1.12e3	4.01e3	7.08e3	4.81e1	9.5925e1	1.92e2	3.91e2
$h_0/2$	2.74e2	4.71e2	2.80e3	5.59e3	2.14e1	4.35e1	8.70e1	1.75e2
$h_0/4$	–	2.60e2	1.18e3	3.42e3	–	2.09e1	4.24e1	8.44e1
$h_0/8$	–	–	3.45e2	1.75e3	–	–	2.05e1	4.26e1
$\kappa(\mathbf{A})$	2.88e5	1.18e6	4.89e6	1.99e7	7.44e5	2.81e6	1.11e7	4.55e7

as $O(Hh^{-1})$, as predicted by Theorem 1. Moreover, we clearly observe that employing a piecewise linear coarse solver ($q = 1$) rather than a piecewise constant coarse solver ($q = 0$) significantly improves the performance of the preconditioner. Indeed, comparing the condition number estimates of the preconditioned system with the analogous ones obtained for the non-preconditioned problem (last row of Table 2) we clearly observe that the condition number of the non-preconditioned system is reduced with respect to the condition number of the preconditioned system by approximately 5 orders of magnitude for $q = 1$ and 4 orders of magnitude for $q = 0$.

Bibliography

[1] P. F. Antonietti and B. Ayuso. Schwarz domain decomposition preconditioners for discontinuous Galerkin approximations of elliptic problems: non-overlapping case. *M2AN Math. Model. Numer. Anal.*, 41(1):21–54, 2007.

- [2] P. F. Antonietti and B. Ayuso. Multiplicative Schwarz methods for discontinuous Galerkin approximations of elliptic problems. *M2AN Math. Model. Numer. Anal.*, 42(3):443–469, 2008.
- [3] P. F. Antonietti and P. Houston. A class of domain decomposition preconditioners for hp -discontinuous Galerkin finite element methods. *J. Sci. Comp.*, 46(1):124–149, 2011.
- [4] D. N. Arnold. An interior penalty finite element method with discontinuous elements. *SIAM J. Numer. Anal.*, 19(4):742–760, 1982.
- [5] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. *SIAM J. Numer. Anal.*, 39(5):1749–1779 (electronic), 2001/02.
- [6] A. T. Barker, S. C. Brenner, P. Eun-Hee, and L.-Y. Sung. Two-level additive Schwarz preconditioners for a weakly over-penalized symmetric interior penalty method. *J. Sci. Comp.*, 47:27–49, 2011.
- [7] S. C. Brenner and K. Wang. Two-level additive Schwarz preconditioners for C^0 interior penalty methods. *Numer. Math.*, 102(2):231–255, 2005.
- [8] S. C. Eisenstat, H. C. Elman, and M. H. Schultz. Variational iterative methods for nonsymmetric systems of linear equations. *SIAM J. Numer. Anal.*, 20(2):345–357, 1983.
- [9] X. Feng and O. A. Karakashian. Two-level additive Schwarz methods for a discontinuous Galerkin approximation of second order elliptic problems. *SIAM J. Numer. Anal.*, 39(4):1343–1365 (electronic), 2001.
- [10] A. Toselli and O. Widlund. *Domain decomposition methods—algorithms and theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 2005.