

Efficient Multilevel Solvers and High Performance Computing Techniques for the Finite Element Simulation of Large-Scale Elasticity Problems

Hilmar Wobker¹

¹Institute of Applied Mathematics and Numerics, TU Dortmund, Germany email: hilmar.wobker@math.tu-dortmund.de

January 15, 2010



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...





Efficiency of iterative linear solvers influenced by

- material parameters,
- nonlinear effects,
- the shape of the geometry,
- the size and the quality of the underlying computational mesh,
- algorithmic parameters, and
- the number of processors in a parallel computing system.



Introduction/Motivation



Three aspects of efficiency:

- numerical efficiency: amount of work to achieve a desired goal; convergence rate and robustness
- processor efficiency: ability to exploit the full capacity of modern hardware
- parallel efficiency: communication vs. computation, scalability

Hardware-oriented numerics:

achieve good efficiency in all three aspects

 \rightarrow difficult to realise due to conflicting demands

 \rightarrow our attempt: FEAST





On the one hand:

- hardware-oriented library is tedious to develop and to maintain (multi-core architectures, vector processors, Cache-sizes, co-processors (GPUs, Cell), latency/speed of interconnects, compiler issues, ...)
- mandatory for high efficiency: a priori known data layout, especially of FE matrices

On the other hand:

- different physical applications \Rightarrow different matrix structures (heat transfer, elasticity, Navier-Stokes, Fluid-Solid-Interaction, 2D / 3D, etc.)
- 'application programmers' don't want to be bothered with technical details



Remedy: Realise *multivariate* operations (operators) as a series (set) of *scalar* operations (operators)

Advantages:

- facilitates strict separation of low-level kernel/library functionalities and high-level application code
- 'kernel programmers' can concentrate on the scalar equation case, do not have to heed all the physical applications
- 'application programmers' can concentrate on the application, do not have to heed processor architectures, MPI communication, matrix fill-in patterns, ...
- efficiency of the scalar kernel routines automatically available for multivariate problems
- kernel enhancements (new finite element, GPU solvers, ...) usable without any changes of the application code

FEAST

 \rightarrow prototypically demonstrated with <code>FEASTsolid</code>

Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples
- ML Saddle Point Solvers for Incompressible Elasticity Problems
 Just a Brief Overview...



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

FEAST's Meshing Concept

- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...





Memory wall problem:

'data moving \gg data processing'

- improvement of hardware characteristics per year (1990-2004):
 - processor peak performance: 60 %
 - memory bandwidth: 30 %
 - memory latency: 5.5 %
- 1992: 1 memory access \approx 1 FLOP 2004: 1 memory access \approx 100 FLOPs
- 'memory gap' is still broadening
- unstructured meshes \Rightarrow indirect adressing \Rightarrow expensive memory access
- plus: low arithmetic intensity (sparse matrices)
 ⇒ poor processor efficiency (low MFLOP/s rates)



FEAST's Meshing Concept



Remedy: use structured data with high spatial and temporal locality

• FEAST uses generalised tensor product meshes



rowwise numbering

 \Rightarrow exactly 9 matrix bands for bilinear elements

- direct adressing, caching
- optimised Linear Algebra routines (SparseBandedBLAS)



FEAST's Meshing Concept



- more complex (unstructured) domains by joining several (structured) TP meshes
- local matrices + appropriate border data exchanges
 ⇒ 'virtual' global matrix



Here:

- 64 TP meshes (=64 local matrices), each refined 10x
- distributed over 16 processors
- $\bullet \ \Rightarrow 1.34 \cdot 10^8$ degrees of freedom in total



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples
- ML Saddle Point Solvers for Incompressible Elasticity Problems
 Just a Brief Overview...



FEAST's Adaptivity Concept







Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept

• FEAST's Solver Concept

- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Parallel vs. Numerical Efficiency technische universität

- improvement of hardware characteristics per year (1990–2004):
 - processor peak performance: 60 %
 - network technology (latency, bandwidth): 30 %
- physical constraint: speed of light
- 2004: 1 inter-processor communication \approx 4000 FLOPs 2020: 1 inter-processor communication \approx 670 000 FLOPs
- number of processors in high-end super computers: 2004: ≈ 4000 2010: $> 100\,000$
- similar to the 'memory wall' problem
- ullet \Rightarrow parallel efficiency determined by amount of data exchange

data locality required for parallel efficiency





On the other hand:

- elliptic problems: solution in a point is influenced by all boundary values
- information has to 'travel' at least once through the whole grid
- fast global data exchange necessary for good iterative convergence rates

global information required for numerical efficiency

- conflicting demands: data locality vs. fast global data exchange
- parallel and numerical efficiency hard to combine



Requirements of parallel solution methods:

- low inter-processor communication
- high processor loads
- good scalability
- good convergence behaviour
- robustness with respect to complicated geometries, mesh refinement level, mesh irregularities, and number/size of subdomains

Parallel vs. Numerical Efficiency technische universität

Suitable solver concepts for elliptic problems:

- multigrid methods (MG)
- domain decomposition methods (DD)



Multigrid





- use a hierarchy of nested grids to solve Ax = b
- basic components:
 - smoothing (level /): $\mathbf{x}^{(l)} \leftarrow \mathbf{x}^{(l)} + \omega \mathbf{S}^{(l)} (\mathbf{b}^{(l)} \mathbf{A}^{(l)} \mathbf{x}^{(l)})$
 - grid transfer: restriction $(I \rightarrow I 1)$, prolongation $(I 1 \rightarrow I)$ coarse grid solving: $\mathbf{x}^{(1)} = (\mathbf{A}^{(1)})^{-1}\mathbf{b}^{(1)}$
- exploiting smoothing property of elementary methods like Jacobi or Gauß-Seidel
- convergence behaviour independent of the refinement level
- runtime O(#DOF)



Multigrid

technische universität dortmund

Possibilities to traverse the grid hierarchy:

- V-, F- and W-cycle
- arithmetic costs vs. convergence speed / robustness







Block-smoothed Multigrid





1			_			

- smoothing operation highly recursive (Gauß-Seidel, ILU)
- bad parallelisation potential
- remedy: relax the smoothing operation by applying it *block-wise* (additively, block-Jacobi)
- use *minimal overlap* (next slide)
- number of vertices n, number of subdomains M: $M \rightarrow 1$: standard multigrid with selected smoother $M \rightarrow n$: standard multigrid with Jacobi smoother



Minimal Overlap



Switching between global layer and local layer:

- global set of vertices on level I: $\mathcal{V}^{(I)}$
- set of vertices of *i*-th subdomain: $\mathcal{V}_{i}^{(l)}$
- local index of vertex k in subdomain i: $loc_i^{(l)}(k)$
- prolongation matrix $\mathbf{P}_{i}^{(l)}$:

$$(\mathbf{x}^{(l)})_{k} = (\mathbf{P}_{i}^{(l)}\mathbf{x}_{i}^{(l)})_{k} := \begin{cases} (\mathbf{x}_{i}^{(l)})_{\log_{i}^{(l)}(k)} & k \in \mathcal{V}_{i}^{(l)} \\ 0 & k \in \mathcal{V}^{(l)} \setminus \mathcal{V}_{i}^{(l)} \end{cases}$$

- restriction matrix: $\mathbf{R}_{i}^{(l)} := (\mathbf{P}_{i}^{(l)})^{\mathsf{T}}$ • local matrix: $\mathbf{A}_{i}^{(l)} := \mathbf{R}_{i}^{(l)} \mathbf{A}_{i}^{(l)} \mathbf{P}_{i}^{(l)}$
- local matrix: $\mathbf{A}_{i}^{(r)} := \mathbf{R}_{i}^{(r)} \mathbf{A}_{i}^{(r)} \mathbf{P}_{i}^{(r)}$
- interpretation: 'ghost cells', extended Dirichlet boundaries

	1				

\square			_	4	
\square		\vdash		ЦĻ,	
\vdash					
H-				₩.	
		\vdash	++	Hi-	+
\vdash		\vdash		⊢	+
\vdash		l-t-	┿╍┾╼	H.	+++
H	++			++	
H	++	++		++	



Domain Decomposition



interpret DD as preconditioner Ã:

$$\mathbf{x} \leftarrow \mathbf{x} + \omega \mathbf{\tilde{A}} (\mathbf{b} - \mathbf{A} \mathbf{x})$$

- global coupling typically via Krylov space method
- two classes:
 - non-overlapping methods (substructuring / Schur complement), extra interface problem
 - overlapping Schwarz methods







technische universität

dortmund

Overlapping Schwarz Methods



- one-level Schwarz method: dependence on number of subdomains, size of the overlap
- add coarse grid problem: two-level Schwarz method
- generalisation: multilevel Schwarz methods (MLDD)
- MLDD methods can be
 - purely additive (additive within one level / additive between levels)
 - purely multiplicative (multiplicative/multiplicative)
 - hybrid (additive/multiplicative or multiplicative/additive)
- FEAST uses a hybrid MLDD method with minimal overlap that is
 - additive within one level
 - multiplicative between levels







FEAST's Solution Method



• one iteration (multiplicative part between levels):

$$\begin{aligned} \mathbf{x}^{(L)} \leftarrow \mathbf{x}^{(L)} + \tilde{\mathbf{A}}^{(L)} (\mathbf{b}^{(L)} - \mathbf{A}^{(L)} \mathbf{x}^{(L)}) \\ & l = L - 1, \dots, 2: \\ \mathbf{b}^{(l)} \leftarrow \mathbf{R}^{(l)} (\mathbf{b}^{(l+1)} - \mathbf{A}^{(l+1)} \mathbf{x}^{(l+1)}), \quad \mathbf{x}^{(l)} \leftarrow \tilde{\mathbf{A}}^{(l)} \mathbf{b}^{(l)} \\ & l = 1: \\ \mathbf{b}^{(1)} \leftarrow \mathbf{R}^{(2)} (\mathbf{b}^{(2)} - \mathbf{A}^{(2)} \mathbf{x}^{(2)}), \quad \mathbf{x}^{(1)} \leftarrow (\mathbf{A}^{(1)})^{-1} \mathbf{b}^{(1)} \\ & l = 2, \dots, L: \\ \mathbf{x}^{(l)} \leftarrow \mathbf{x}^{(l)} + \mathbf{P}^{(l)} \mathbf{x}^{(l-1)}, \quad \mathbf{x}^{(l)} \leftarrow \mathbf{x}^{(l)} + \tilde{\mathbf{A}}^{(l)} (\mathbf{b}^{(l)} - \mathbf{A}^{(l)} \mathbf{x}^{(l)}) \end{aligned}$$

 connection between global and local layer (additive part within one level):

$$\tilde{\mathbf{A}}^{(l)} := \widetilde{\sum}_{i=1}^{M} \mathbf{P}_i^{(l)} \tilde{\mathbf{A}}_i^{(l)} \mathbf{R}_i^{(l)}.$$

- FEAST
- local preconditioner $\tilde{\mathbf{A}}_{i}^{(l)}$ for the local matrix $\mathbf{A}_{i}^{(l)}$ (slide after next)



Why minimal overlap?

- sufficient for robust convergence behaviour of multilevel DD \rightarrow numerical efficiency
- local submeshes preserve tensor product property
 - \rightarrow processor efficiency
- minimal amount of data exchange between subdomains \rightarrow parallel efficiency
- implementation and data structures greatly simplified



The Local Preconditioner



- two strategies to realise the local preconditioner $\tilde{\mathbf{A}}_{i}^{(l)}$:
 - apply one step of an elementary iterative scheme (Jacobi, Gauß-Seidel, ...)
 - apply some iterative or direct solution method to (approximately) solve the local system
- first strategy: MLDD coincides with block-smoothed MG (Ã^(I)_i being the local smoother)
- $\bullet \ \Rightarrow \mathsf{MLDD} \ \mathsf{generalisation} \ \mathsf{of} \ \mathsf{MG}$



The Local Preconditioner

technische universität dortmund

First strategy:

• Notation of the local 9-band matrix (omitting superscript *I*):

$$\mathbf{A}_i = (\mathbf{L}_i^L + \mathbf{L}_i^C + \mathbf{L}_i^U) + (\mathbf{C}_i^L + \mathbf{C}_i^C + \mathbf{C}_i^U) + (\mathbf{U}_i^L + \mathbf{U}_i^C + \mathbf{U}_i^U)$$



• some local smoothers:

$$\begin{split} \tilde{\mathbf{A}}_{i}^{\text{Jacobi}} &:= (\mathbf{C}_{i}^{C})^{-1} \\ \tilde{\mathbf{A}}_{i}^{\text{GS}} &:= (\mathbf{L}_{i}^{L} + \mathbf{L}_{i}^{C} + \mathbf{L}_{i}^{U} + \mathbf{C}_{i}^{L} + \mathbf{C}_{i}^{C})^{-1} \\ \tilde{\mathbf{A}}_{i}^{\text{TriGS}} &:= (\mathbf{L}_{i}^{L} + \mathbf{L}_{i}^{C} + \mathbf{L}_{i}^{U} + \mathbf{C}_{i}^{L} + \mathbf{C}_{i}^{C} + \mathbf{C}_{i}^{U})^{-1} \\ \tilde{\mathbf{A}}_{i}^{\text{MTriGS}} &: \tilde{\mathbf{A}}_{i}^{\text{TriGS}} \text{ applied to column-wise numbered grid} \\ \tilde{\mathbf{A}}_{i}^{\text{ADiTriGS}} &: \text{ alternating application of } \tilde{\mathbf{A}}_{i}^{\text{TriGS}} \text{ and } \tilde{\mathbf{A}}_{i}^{\text{MTriGS}} \end{split}$$





Second strategy:

- solve local systems $\mathbf{A}_i \mathbf{x}_i = \mathbf{b}_i$ (approximately) via
 - *direct solver* if the system is not too large (#DOF< 20000)
 - multigrid method, otherwise
- local MG uses the same smoothers as block-smoothed (global) MG (see previous slide)
- 'automatically toggle' between MG and direct solver via *truncated multigrid method*
- typical local problem size: $\#DOF \approx 10^{6}$ \Rightarrow local multigrid mandatory

Using local multigrid within global MLDD is one of the core ideas of FEAST's solution method!





Comparison of the two strategies (black dot = favoured strategy):

storage requirements	1	$\circ \bullet \circ \circ$	2
flexibility	1	0000	2
'black box'character	1	$\circ \bullet \circ \circ$	2
arithmetic costs	1	• • • •	2
communication vs. computation	1	000•	2
influence on global iteration	1	000•	2
load balancing	1	• • • •	2
using co-processors	1	000•	2



The ScaRC Concept



Basic idea of FEAST's solver concept ScaRC (Scalable Recursive Clustering):

- allow both strategies at the same time
- choose local solver components adaptively according to the 'local situation'

Advantages of ScaRC:

- convergence rates independent of refinement level
- convergence rates independent of number of subdomains (in case of only mild macro anisotropies)
- 'simple' local solvers/smoothers on 'simple' subdomains, 'strong' ones on 'difficult' subdomains
- good balance of computational costs and convergence behaviour
- hide local irregularities from the global solver
- minimise number of global iterations and amount of communication
- high processor loads due to local tensor product grids

The ScaRC Concept



Disadvantages of ScaRC:

- convergence rates dependent on number of subdomains in case of stronger macro anisotropies (block-Jacobi character)
 - can be alleviated by enhancing the global multilevel solver with Krylov space methods
 - further idea: 'Recursive Clustering' (3-layer-ScaRC, merging subdomains)
- bad ratio computation/communication on coarser grid levels (inherent to multilevel approaches!)
- difficult to realise: 'automatic adaptation system' for local solver components
- even more difficult: dynamic load balancing
- local multigrid solvers are nonconforming (next slides)



Notations



- first strategy: 1-layer-ScaRC (multigrid scheme only on global layer)
- second strategy: 2-layer-ScaRC (multigrid schemes on global and local layer)
- short notation
 - 1-layer-ScaRC: MG__JAC__D MG(1e-6,V44,0.7)__JAC__D
 - 2-layer-ScaRC:

 $\texttt{MG__MG-ADI-D__D}$

MG(1e-6,F22,0.7)__MG(1e-1,V22)-ADI-D__D

'__': layer change
'D': direct coarse grid solver
'1e-6': relative stopping criterion ('gain 6 digits')
'V22': cycle type and pre-/postsmoothing steps
'0.7': damping parameter



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples
- ML Saddle Point Solvers for Incompressible Elasticity Problems
 Just a Brief Overview...



Nonconforming Local MG





			_			





- minimal overlap \Rightarrow extended Dirichlet boundaries
- local domain size increases with coarser grid levels
- $\bullet\,$ nonnested local grids $\Rightarrow\,$ nonconforming local multigrid method
- coarse grid correction not optimal
- (massive) convergence problems of standard multigrid schemes
- loss of level independency



Adaptive Coarse Grid Correction technische universität

Remedy:

• adaptively damp the suboptimal coarse grid correction:

$$\mathbf{x}^{(l)} \leftarrow \mathbf{x}^{(l)} + \alpha \mathbf{P}^{(l)} \mathbf{x}^{(l-1)}$$

(subscripts for local subdomains omitted)

• minimise error in energy norm:

$$\left(\boldsymbol{x}^{(\textit{l})}-\boldsymbol{x}^{*}\right)^{\mathsf{T}}\boldsymbol{\mathsf{A}}^{(\textit{l})}(\boldsymbol{x}^{(\textit{l})}-\boldsymbol{x}^{*})$$

 $(\mathbf{x}^* \text{ exact solution})$

• for symmetric **A**^(*I*):

$$\alpha = \frac{\mathbf{c}^{(l)^{\mathsf{T}}} \mathbf{d}^{(l)}}{\mathbf{c}^{(l)^{\mathsf{T}}} \mathbf{A}^{(l)} \mathbf{d}^{(l)}}$$

 $\begin{aligned} & (\mathbf{c}^{(l)} := \mathbf{P}^{(l)} \mathbf{x}^{(l-1)} \text{ prolongated coarse grid correction,} \\ & \mathbf{d}^{(l)} := \mathbf{b}^{(l)} - \mathbf{A}^{(l)} \mathbf{x}^{(l)} \text{ current defect)} \end{aligned}$


Adaptive Coarse Grid Correction technische universität

Advantages:

- established technique
- easy to implement
- relatively low computational costs (one matrix vector multiplication, two scalar products)
- damping parameter computed per subdomain:
 - $\alpha \ll 1$ for subdomains that suffer strongly from ext. Dirichl. bound.
 - $\alpha \approx 1$ otherwise
 - \Rightarrow overall solution process not unnecessarily impaired in general

Heuristic used in FEAST before ACGC:

- prolongated values corresponding to subdomain boundary nodes were not fully added (divided by number of incident subdomains)
- equally affects all local solves
- fails in some cases







- (outer) geometric boundary conditions: Neumann ('N'), Dirichlet ('D')
- (inner) extended Dirichlet boundaries ('E')
- only consider local solve on top right subdomain (center figure)
- reference computation with standard Dirichlet instead of extended Dirichlet boundary conditions (right figure)
- solve standard Poisson problem $-\Delta u = 1$
- o local solver 1e-6,MG([V|F]22,0.7)-JAC-D





isotropic subdomain, Jacobi smoother



FEAST











- anisotropic example
- consider top right subdomain again
- o local solver MG(1e-6, [V|F]22,0.6)-JAC-D





V-cycle (left) and F-cycle (right)

Summary for W-cycle:

- W-cycle with heuristic: converges in most, but not in all cases \Rightarrow not reliable
- standard W-cycle: always converges, but strange oscillations depending on number of grid levels
- W-cycle + ACGC: always converges, no oscillations



technische universität

dortmund





- now: use ADiTriGS instead of Jacobi as local smoother
- increase anisotropies
- consider top right subdomain again
- local solver MG(1e-6, [V|F]22,1.0)-ADI-D
- solve standard Poisson problem $-\Delta u = 1$







anisotropic subdomain, ADiTriGS smoother



V-cycle (left) and F-cycle (right)

- hence, standard/heuristic F-cycle reliable?
- no! (next slide)





anisotropic subdomain, ADiTriGS smoother, F-cycle



isotr. operator $-\Delta u$ (left), anisotr. operator $-10\partial_{xx}u - \partial_{yy}u$ (right)

Summary for W-cycle:

- standard W-cycle and W-cycle with heuristic: always converge, but strange oscillations depending on number of grid levels
- W-cycle + ACGC: always converges, no oscillations





Now: consider arithmetic costs of the global solver scheme

• use total arithmetic efficiency:

$$\mathsf{TAE} = -\frac{\#\mathsf{FLOPs}}{\#\mathsf{DOF} \times \#\mathsf{iter} \times \mathsf{log}_{10}(c)}$$

 \Rightarrow How many FLOPs are needed per DOF to gain one digit?

- 2-layer-ScaRC solver with Jacobi smoother: MG(1e-6,V11)__MG(1e-1,V22,0.7)-JAC-D__D
- 2-layer-ScaRC solver with ADiTriGS smoother: MG(1e-6,F22)__MG(1e-1,V22)-ADI-D__D
- only vary maximum MG level, fix coarse grid level to 1



2-layer-ScaRC with local Jacobi smoother, isotropic operator



isotropic configuration (left), anisotropic configuration (right)



technische universität

dortmund



2-layer-ScaRC with local ADiTriGS smoother, anisotropic configuration



isotropic operator (left), anisotropic operator (right)





Adaptive coarse grid correction:

- method to alleviate the negative effects of the extended Dirichlet boundary conditions ('increasing subdomain size')
- easy to implement
- Iow arithmetic costs
- facilitates the use of V- and F-cycle multigrid
- smoothes oscillatory convergence behaviour of W-cycle
- often less expensive than standard W-cycle (GPUs!)
- and: can also improve standard (conforming) multigrid methods



Overview



Introduction/Motivation

Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



1-layer-ScaRC vs. 2-layer-ScaRC to technische universität

- Is 2-layer-ScaRC really superior to 1-layer-ScaRC?
 - use slightly anisotropic operator $-\partial_{xx}u 4\partial_{yy}u = 1$
 - use multigrid-Krylov solvers:
 - 1-layer-ScaRC: MG-FGMRES4__JAC__D
 - 2-layer-ScaRC: MG-FGMRES4__MG(T7)-BICG-JAC-D__D
 - use more complex grids (see next slides):
 - CROSSOVER-ISO: 64 subdomains, 16.8M DOF (level 9), aspect ratio 2.91
 - CROSSOVER-ANISO: aspect ratio 20.4
 - ASMO: 70 subdomains, 18.4M DOF (level 9), aspect ratio 18.2





CROSSOVER







ASMO







1-layer-ScaRC vs. 2-layer-ScaRC to technische universität dortmund

CROSSOVER-ISO (left), CROSSOVER-ANISO (right), ASMO (bottom)





Number of global smoothing steps $(\Rightarrow \text{ amount of communication})$

	M	G-FGMRES	34	MG-FGMRES4MG				
lev	CR-I	CR-A1	AS-I	CR-I	CR-A1	AS-I		
5	88	352	776	32	32	48		
6	96	464	1016	32	32	48		
7	104	552	1296	32	32	48		
8	112	624	1504	32	32	48		
9	120	680	1688	32	32	56		

- 2-layer-ScaRC successfully hides the local irregularities (anisotropies) from the global solver
- ullet \Rightarrow strongly favoured in a massively parallel computation





But:

- for ADiTriGS as local smoother, examples to show superiority of 2-layer-ScaRC not found yet
- current favourite: 1-layer-ScaRC solver BiCG-MG__ADI__D
- reason: for ADiTriGS, mesh anisotropies are not really irregularities
 ⇒ there is 'nothing to hide' within a 2-layer-ScaRC solver
- future attempts: massively parallel computations, more complicated physical equations, 3D



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

3 Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...





Elasticity:

- continuum mechanical discipline about deformation of solid elastic bodies
- solid and deformable: forces change the body's shape, but not its continuous coherence
- elastic: deformation process is reversible
- assumptions:
 - material is homogeneous and isotropic
 - only small deformations (linearised elasticity)





- solid body $\bar{\Omega} \subset \mathbb{R}^3$
- boundary $\Gamma:=\partial\Omega,\ \Gamma=\Gamma_D\cup\Gamma_N$
- deformation mapping ${f \Phi}: ar \Omega o \mathbb{R}^3$ with det $(
 abla {f \Phi}) > 0$
- displacements $\mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), u_2(\mathbf{x}), u_3(\mathbf{x}))^{\mathsf{T}}$ of material point $\mathbf{x} \in \overline{\Omega}$, $\mathbf{\Phi} = \mathbf{id} + \mathbf{u}$
- kinematic relation between displacements and strains: linearised strain tensor $\boldsymbol{\varepsilon} = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^{\mathsf{T}})$





Stress principle of Euler and Cauchy:

existence of a vector field $t:\bar\Omega\times S_1\to\mathbb{R}^3$ (Cauchy stress vector) with:

• for arbitrary $V \subset \overline{\Omega}$ (g applied surface force):

$$\mathbf{t}(\mathbf{x}, \mathbf{n}) = \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Gamma_{N} \cap \partial V$$

• axiom of force balance (f applied body force):

$$\int_V \mathbf{f}(\mathbf{x}) \, \mathrm{d} x + \int_{\partial V} \mathbf{t}(\mathbf{x}, \mathbf{n}) \, \mathrm{d} \mathbf{a} = \mathbf{0}$$

• axiom of balance of angular momenta:

$$\int_{V} \mathbf{x} \times \mathbf{f}(\mathbf{x}) \, \mathrm{d}x + \int_{\partial V} \mathbf{x} \times \mathbf{t}(\mathbf{x}, \mathbf{n}) \, \mathrm{d}a = \mathbf{0}$$





Cauchy's theorem:

existence of a symmetric tensor field $\sigma: \bar{\Omega} \to \mathbb{M}^3$ (Cauchy stress tensor) that satisfies

- $\mathbf{t}(\mathbf{x},\mathbf{n}) = \boldsymbol{\sigma}(\mathbf{x})\mathbf{n}, \quad \mathbf{x} \in \bar{\Omega}, \mathbf{n} \in S_1$,
- the PDE

$$-\operatorname{div}(\sigma(\mathbf{x})) = \mathbf{f}(\mathbf{x}), \qquad \mathbf{x} \in \Omega,$$

and the boundary conditions

$$\boldsymbol{\sigma}(\mathbf{x})\mathbf{n} = \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Gamma_{\mathsf{N}}.$$



Basic Relations of Elasticity



- constitutive law: relation between strains and stresses
- Hooke's law for isotropic material:

$$\boldsymbol{\sigma} = 2\mu\boldsymbol{\varepsilon} + \lambda\operatorname{tr}(\boldsymbol{\varepsilon})\mathbf{I}$$

- Lamé constants μ and λ
- relation to Young's modulus E and Poisson ratio ν :

$$\mu = \frac{E}{2(1+\nu)}, \qquad \lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}$$

• $\nu \to 0.5 \ (\lambda \to \infty)$: compressibility of material decreases

Material	$E [N/m^2]$	ν[-]	$\mu [N/m^2]$	$\lambda [N/m^2]$
Steel/Iron	$2.1 \cdot 10^{11}$	0.29	$8.14 \cdot 10^{10}$	$1.1\cdot10^{11}$
Lead	$1.6 \cdot 10^{10}$	0.44	$5.6 \cdot 10^{9}$	$4.1 \cdot 10^{10}$
Rubber	$2.5 \cdot 10^{7}$	0.499–0.5	$8.2 \cdot 10^{6}$	$4.1 \cdot 10^{9}$





• resulting boundary value problem - the Lamé Equation:

$$\begin{split} -2\mu \; \operatorname{div}(\boldsymbol{\varepsilon}) &-\lambda \, \boldsymbol{\nabla} \operatorname{div}(\mathbf{u}) = \mathbf{f} & \quad \text{in } \boldsymbol{\Omega} \\ \mathbf{u} &= \bar{\mathbf{u}} & \quad \text{on } \boldsymbol{\Gamma}_{\mathrm{D}} \\ \boldsymbol{\sigma} \mathbf{n} &= \mathbf{g} & \quad \text{on } \boldsymbol{\Gamma}_{\mathrm{N}} \end{split}$$

bilinear form:

$$k(\mathbf{u},\mathbf{v}) := \int_{\Omega} 2\mu \varepsilon(\mathbf{u}) : \varepsilon(\mathbf{v}) + \lambda \operatorname{div}(\mathbf{u}) \operatorname{div}(\mathbf{v}) \operatorname{dx}$$

• weak formulation: Find $\mathbf{u} - \bar{\mathbf{u}} \in \mathbf{V} := \left\{ \mathbf{v} \in H^1(\Omega)^3 \mid \mathbf{v} = \mathbf{0} \text{ on } \Gamma_D \right\}$ such that

$$k(\mathbf{u},\mathbf{v}) = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, \mathrm{d}x + \int_{\Gamma_{\mathsf{N}}} \mathbf{g} \cdot \mathbf{v} \, \mathrm{d}a, \qquad \mathbf{v} \in \mathbf{V}$$



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

3 Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Operator Splitting



- motivation: reduction to scalar components
- rewrite left hand side of BVP (now 2D):

$$-2\mu \operatorname{div}(\varepsilon) - \lambda \nabla \operatorname{div}(\mathbf{u})$$
$$= \begin{pmatrix} (2\mu + \lambda)\partial_{11} + \mu\partial_{22} & (\mu + \lambda)\partial_{12} \\ (\mu + \lambda)\partial_{21} & \mu\partial_{11} + (2\mu + \lambda)\partial_{22} \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}$$

• rewrite left hand side of weak form:

$$k(\mathbf{u}, \mathbf{v}) = \int_{\Omega} (2\mu + \lambda)\partial_1 u_1 \partial_1 v_1 + \mu \partial_2 u_1 \partial_2 v_1 \, dx$$

+ ... (mixed terms)
+ $\int_{\Omega} \mu \partial_1 u_2 \partial_1 v_2 + (2\mu + \lambda)\partial_2 u_2 \partial_2 v_2 \, dx$
= ... (next slide)



Operator Splitting

technische universität dortmund

• rewrite left hand side of weak form:

$$k(\mathbf{u},\mathbf{v}) = \underbrace{\int_{\Omega} \left[\begin{pmatrix} 2\mu + \lambda & 0 \\ 0 & \mu \end{pmatrix} \nabla u_1 \right] \cdot \nabla v_1 \, \mathrm{d}x}_{=:k_{11}(u_1,v_1)} + k_{12}(u_2,v_1)$$
$$+ k_{21}(u_1,v_2) + \underbrace{\int_{\Omega} \left[\begin{pmatrix} \mu & 0 \\ 0 & 2\mu + \lambda \end{pmatrix} \nabla u_2 \right] \cdot \nabla v_2 \, \mathrm{d}x}_{=:k_{22}(u_2,v_2)}$$

• FE discretisation (Q_1) \Rightarrow block structured linear equation system:

$$\begin{pmatrix} \textbf{K}_{11} & \textbf{K}_{12} \\ \textbf{K}_{21} & \textbf{K}_{22} \end{pmatrix} \begin{pmatrix} \textbf{u}_1 \\ \textbf{u}_2 \end{pmatrix} = \begin{pmatrix} \textbf{f}_1 \\ \textbf{f}_2 \end{pmatrix}$$

('separate displacement ordering'; not used, e.g., in FEAP)

• K₁₁ and K₂₂ correspond to scalar elliptic operators ('anisotropic Laplace operator')



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

3 Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting

Solution Concept

• Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Solution Concept



• basic iteration - block preconditioned Richardson method:

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \mathbf{\tilde{K}}^{-1}(\mathbf{f} - \mathbf{K}\mathbf{u}^k)$$

o defect computation:

$$\mathbf{f} - \mathbf{K}\mathbf{u}^k = \begin{pmatrix} \mathbf{f}_1 - \mathbf{K}_{11}\mathbf{u}_1^k - \mathbf{K}_{12}\mathbf{u}_2^k \\ \mathbf{f}_2 - \mathbf{K}_{21}\mathbf{u}_1^k - \mathbf{K}_{22}\mathbf{u}_2^k \end{pmatrix}$$

• e.g., block Jacobi preconditioner:

$$\boldsymbol{\tilde{K}}_{\mathsf{BJac}}^{-1} = \begin{pmatrix} \boldsymbol{\mathsf{K}}_{11}^{-1} & \boldsymbol{\mathsf{0}} \\ \boldsymbol{\mathsf{0}} & \boldsymbol{\mathsf{K}}_{22}^{-1} \end{pmatrix}$$

- reduction to scalar components ⇒ exploit FEAST concepts!
- note: each global matrix K_{ij} is represented by local 9-band matrices corresponding to TP subdomains



Solution Concept



• condition number of the preconditioned system:

$$\kappa := \kappa \big(\tilde{\mathbf{K}}_{\mathsf{BJac}}^{-1} \mathbf{K} \big) \leqslant \frac{1}{c_{\mathsf{K}}} \Big(1 + \frac{1}{1 - 2\nu} \Big)$$

three observations:

κ does not depend on mesh size parameter h
 ⇒ single-grid outer solver (i. e., Krylov method) can be sufficient

 κ depends on Korn's constant c_K = c_K(Ω, Γ_D)
 ⇒ convergence depends on boundary conditions / geometry (e. g., cantilever beam)
 (Korn's inequality: ∃ c_K = c_K(Ω, Γ_D) > 0 : ||ε(v)||₀² ≥ c_K ||v||₁², v ∈ V)
 κ depends on Poisson ratio ν
 ⇒ condition of the system increases with incompressibility of the material (ν → 0.5)

block Gau
ß-Seidel or block SOR instead of block Jacobi
 ⇒ no such estimate, but more efficient in most cases



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples
- ML Saddle Point Solvers for Incompressible Elasticity Problems
 Just a Brief Overview...



Selected Numerical Examples



- dependence on mesh anisotropies
- dependence on geometry
- 1-layer-ScaRC vs. 2-layer-ScaRC
- parallel efficiency



Dependence on Mesh Anisotr.



4×2 subdomains (left), 8×4 subdomains (right)

·			_	_		
		_				

							`	C	,	'		
_	_	_	_	_	_	_		_	_			



				_		


Dependence on Mesh Anisotr.



BiCG-BJac (left), BiCG-BSor (right)

- convergence independent of mesh refinement level (outer single-grid Krylov solver!)
- BSor roughly halves number of iterations compared to BJac
- anisotropies are fully 'absorbed' by the scalar subsolves



technische universität

dortmund

Dependence on Geometry





- typical cantilever beam configuration: left side fixed, other sides free, vertical body force applied
- increasing global anisotropy: L/H = 4, 16, 64
- two meshings:

4/16/64 isotropic subdomains vs. 1 anisotropic subdomain

- two solvers:
 - BiCG-BSor: outer single-grid Krylov solver
 - BiCG-MG-BSor: multigrid (applied to block system!) as preconditioner for Krylov solver



Dependence on Geometry







- single-grid Krylov solver not sufficient: convergence (more or less) independent of refinement level, but strongly dependent on global anisotropy
- additional multigrid greatly weakens this dependency (at least in case of isotropic subdomains)



applying multigrid scheme only to scalar subsystems is not sufficient!

1-layer-ScaRC vs. 2-layer-ScaRC to technische universität dortmund





- two variants: ISO (max. aspect ratio 2.2) and ANISO (max. aspect ratio 63.6) (only ANISO shown)
- o compare three solvers:
 - using 1-layer-ScaRC: BiCG-MG(1,V11)-BSor[MG(1e-1)-FGMRES4_JAC_D]
 - using 2-layer-ScaRC: BiCG-MG(1,V11)-BSor[MG(1e-1)-FGMRES4_MG-BICG-JAC-D_D]
 - 'standard solver': MG(F11)-BiCG(2)-Jac (using point Jacobi smoother, disregarding the block structure)



1-layer-ScaRC vs. 2-layer-ScaRC to technische universität



- convergence behaviour of 1-layer-ScaRC and 2-layer-ScaRC variants (nearly) independent of the configuration
- the standard solver suffers signifcantly
- 2-layer-ScaRC always needs 1 iteration per call, 1-layer-ScaRC between 6 (iso) and 70 (aniso) iterations
 ⇒ communication amount of 2-layer-ScaRC variant much smaller!



1-layer-ScaRC vs. 2-layer-ScaRC tu technische universität



- arithmetic costs of 2-layer-ScaRC very high on ISO configuration
- arithmetic costs of 1-layer-ScaRC and the standard solver increase drastically on the ANISO configuration \Rightarrow 2-layer-ScaRC superior
- just 'proof of concept'! (with ADiTriGS: TAE pprox 1000)



Parallel Efficiency





- weak scalability: increase problem size and resources by the same factor
- 'perfect weak scalability': runtime remains constant
- \bullet smallest problem: 4 \times 1 subdomains, 4.2 $\rm M$ vertices, 4 processors
- \bullet largest problem: 16 \times 8 subdomains, 134.2 $\rm M$ vertices, 128 proc.



Parallel Efficiency





- good weak scalability
- comparable to scalar FEAST solvers
- parallel efficiency of the scalar FEAST library fully transfers to elasticity solvers



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Overview



Introduction/Motivation

2 Multilevel Solvers for Scalar Elliptic Equations

- FEAST's Meshing Concept
- FEAST's Adaptivity Concept
- FEAST's Solver Concept
- Local Multigrid Solvers
- 1-layer-ScaRC vs. 2-layer-ScaRC

Multilevel Solvers for Compressible Elasticity Problems

- Basic Relations of Elasticity
- Operator Splitting
- Solution Concept
- Selected Numerical Examples

ML Saddle Point Solvers for Incompressible Elasticity Problems Just a Brief Overview...



Incompressible Elast. Problems

- rubber-like materials are (nearly) incompressible, important for many industrial applications
- basic problem: $\nu \rightarrow 0.5 \Rightarrow \lambda \rightarrow \infty$
- pure displacement formulation: 'volume locking'
- significant deterioration of FE approximation and solver behaviour
- \bullet remedy: mixed displacement-pressure \mathbf{u}/p formulation
- $\bullet\,$ leads to Stokes-like saddle point problem $\mathcal{A} {\bf x} = {\bf b}$ with

$$\mathcal{A} := \begin{pmatrix} \textbf{A} & \textbf{B} \\ \textbf{B}^\mathsf{T} & \textbf{C} \end{pmatrix}, \qquad \textbf{x} := \begin{pmatrix} \textbf{u} \\ \textbf{p} \end{pmatrix}, \qquad \textbf{b} := \begin{pmatrix} \textbf{f} \\ \textbf{0} \end{pmatrix}$$

- C contains compressibility and, if necessary, stabilisation terms (e. g., for Q_1/Q_1)
- distinguish segregated (uncoupled) and coupled solution methods



technische universität

Saddle Point Solvers



Segregated methods:

- solve subsystems for **u** and *p*
- Uzawa / pressure Schur complement methods:

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \mathbf{\tilde{S}}^{-1} (\mathbf{B}^{\mathsf{T}} \mathbf{A}^{-1} \mathbf{f} - (\mathbf{B}^{\mathsf{T}} \mathbf{A}^{-1} \mathbf{B} - \mathbf{C}) \mathbf{p}_k),$$

 \tilde{S}^{-1} preconditioner for Schur complement $S := B^T A^{-1} B - C$ • block-triangular preconditioners:

$$\begin{pmatrix} \mathbf{u}_{k+1} \\ \mathbf{p}_{k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{u}_k \\ \mathbf{p}_k \end{pmatrix} + \begin{pmatrix} \mathbf{\tilde{A}} & \mathbf{0} \\ \mathbf{B}^{\mathsf{T}} & -\mathbf{\tilde{S}} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{f} - \mathbf{A}\mathbf{u}_k - \mathbf{B}\mathbf{p}_k \\ \mathbf{0} - \mathbf{B}^{\mathsf{T}}\mathbf{u}_k - \mathbf{C}\mathbf{p}_k \end{pmatrix},$$

essential for both variants:
 a good Schur complement preconditioner



Saddle Point Solvers



Segregated methods (continued):

- Schur complement preconditioning often involves solution of scalar systems
 - \Rightarrow scalar ScaRC solvers!
- solve Ã-systems with techniques from compressible elasticity
 ⇒ scalar ScaRC solvers!
- hence: also the solution of saddle point systems can essentially be brought down to the solution of scalar systems
- multigrid can be applied to the whole saddle point system
 ⇒ overall solver potentially contains four (!) nested MG schemes
- BEAM configurations: three nested MG schemes can be beneficial in terms of numerical efficiency (with corresponding bad parallel efficiency)



Saddle Point Solvers



Coupled methods:

- multigrid with Vanka smoothers
- solve whole saddle point system at once, but reduction to small subdomains (one element, patch of few elements)
- combine local solutions in multiplicative fashion (block Gauß-Seidel)
- mesh anisotropies need special treatment
- no Schur complement preconditioner needed
- no reduction to scalar subsystems, ScaRC solvers not applicable

In the context of my work: segregated methods clearly superior to coupled Vanka methods (up to 30x - 40x faster)





Thank you for your attention!

