

Domain Decomposition

Parallelization of Mesh Based Applications

Panagiotis Adamidis
Thomas Bönisch

University of Stuttgart
High-Performance Computing-Center Stuttgart (HLRS)
www.hlrs.de

Höchstleistungsrechenzentrum Stuttgart



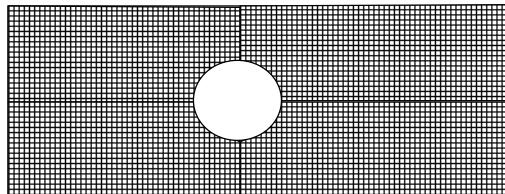
Outline

- Introduction
- Basics
- Boundary Handling
- Example: Finite Volume Flow Simulation on Structured Meshes
- Example: Finite Element Approach on an Unstructured Mesh

Parallelization - Target

- High Application Performance
- Using real big MPP's
 - no loss in efficiency even when using 500 Processors and more
- Using Clusters of SMP's
 - no decrease in Performance due to using external Network connections

A Problem (I)



Flow around a cylinder:
Numerical Simulation using FV, FE or FD

Data Structure: $A(1:n, 1:m)$

Solve: $(A+B+C)x=b$

Parallelization strategies

Work decomposition

Scaling ?

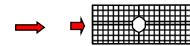
do i=1,100

→ i=1,25

i=26,50

i=51,75

i=76,100



Flow around a cylinder:
Numerical Simulation using FV, FE or FD

→ Data Structure: A(1:n,1:m)

→ Solve: (A+B+C)x=b

Data decomposition

Scales
too much communication ?

A(1:20,1:50)

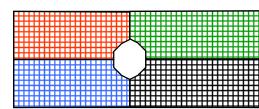
A(1:20,51:100)

A(1:20,101:150)

A(1:20,151:200)

Domain decomposition

Good Chance



H L R I S



Domain Decomposition

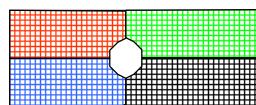
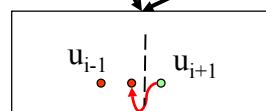
Slide 5

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

Parallelization Problems

- Decomposition (Domain, Data, Work)
- Communication

$$du / dx = (u_{i+1} - u_{i-1}) / dx$$



H L R I S

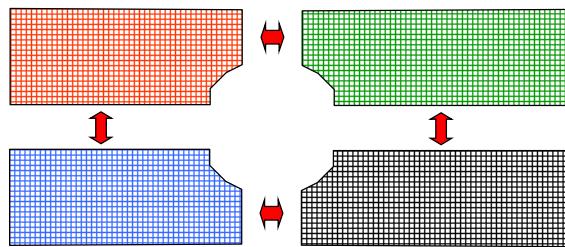


Domain Decomposition

Slide 6

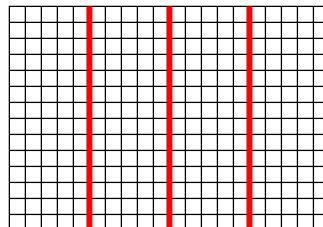
Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

Concepts - Message Passing (II)

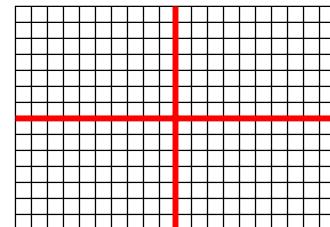


↔ User defined communication

How to split the Domain in the Dimensions (I)



1 - Dimensional



2 (and 3) - Dimensional

How to split the Domain in the Dimensions (II)

- That depends on:
 - computational speed i.e. processor:
vectorprocessor or cache
 - communication speed:
 - latency
 - bandwidth
 - topology
 - number of subdomains needed
 - load distribution (is the effort for every mesh cell equal)

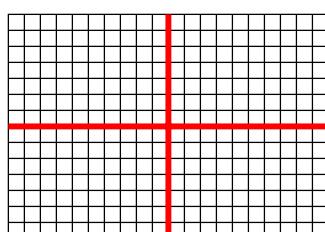
Replication versus Communication (I)

- If we need a value from a neighbour we have basically two opportunities
 - getting the necessary value directly from the neighbour, when needed
 Communication, Additional Synchronisation
 - calculating the value of the neighbour again locally from values known there
 Additional Calculation
- Selection depends on the application

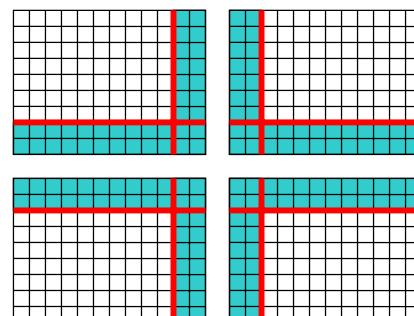
Replication versus Communication (II)

- Normally replicate the values
 - Consider how many calculations you can execute while only sending 1 Bit from one process to another (6 µs, 1.0 Gflop/s → 6000 operations)
 - Sending 16 kByte (20x20x5) doubles (with 300 MB/s bandwidth → 53.3 µs → 53 300 operations)
 - very often blocks have to wait for their neighbours
 - but extra work limits parallel efficiency
- Communication should only be used if one is quite sure that this is the best solution

2- Dimensional DD with two Halo Cells



Mesh Partitioning



Subdomain for each Process



Example:

Parallelization of a 3D Finite Volume Flow Solver



Höchleistungsrechenzentrum Stuttgart

H L R I S



Starting Point: Sequentiell Program

- Written in FORTRAN77
- Using structured meshes
- Parts of the program
 - Preprocessing, reading Data
 - Main Loop
 - **Setup of the equation system**
 - **Preconditioning**
 - **Solving step**
 - Postprocessing, writing Data

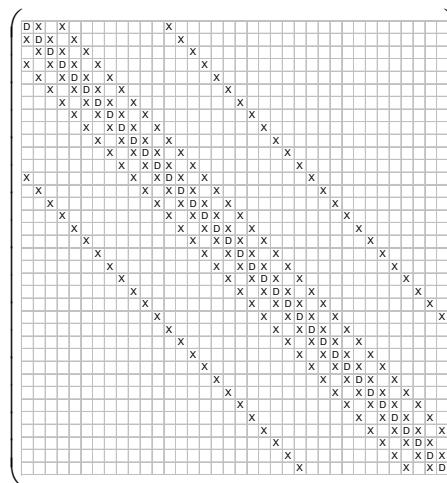
Dynamic Data Structures

- Pure FORTRAN77 is too static
 - number of processors can vary from run to run
 - size of arrays even within the same case can vary
 - dynamic data structures
 - use Fortran90 dynamic arrays
 - use all local memory on a PE for a huge FORTARN77 array and setup your own memory management
 - second method has a problem on SMP's and cc-NUMA's we should only use as much memory as necessary

Main Loop

- Setup of the equation system
 - each cell has 6 neighbours
 - needs data from neighbour cells
 - 2 halo cells at the inner boundaries
- Preconditioning
 - no neighbour information needed at all
 - completely done locally
- Solving step
 - Jacobi Line relaxation with subiterations
 - more complicated (next slides)

Heptadiagonalmatrix



Domain Decomposition Adamidis/Bönisch
Slide 17 Höchstleistungsrechenzentrum Stuttgart

H L R S

Parallelization - Solver (I)

- Sequential:

$$A\vec{u} = \vec{r}$$

$$\vec{y} = A^{-1}\vec{r}$$

- Parallelization problems:

- Matrix A is distributed
 - Matrix inversion is a priori not parallelizable

Domain Decomposition Adamidis/Bönisch
Slide 18 Höchstleistungsrechenzentrum Stuttgart

H L R T S

Parallelization - Solver (II)

$$A\vec{u} = \vec{r}$$

$$A = L + M$$

$$(L + M)\vec{u} = \vec{r}$$

$$L\vec{u} = \vec{r} - M\vec{u}$$

$$L\vec{u}^{(j)} = \vec{r} - M\vec{u}^{(j-1)}$$

$$\vec{u}^{(j)} = L^{-1} \left(\vec{r} - M\vec{u}^{(j-1)} \right)$$



Domain Decomposition
Slide 19

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S

Parallelization - Solver (III)

$$A = L + M$$

$$A = \begin{pmatrix} m_1 & u_1 \\ l_2 & m_2 & u_2 \\ l_3 & m_3 & u_3 \\ l_4 & m_4 \\ \hline m_5 & u_5 \\ l_6 & m_6 & u_6 \\ l_7 & m_7 & u_7 \\ l_8 & m_8 \\ \hline m_9 & u_9 \\ l_{10} & m_{10} & u_{10} \\ l_{11} & m_{11} & u_{11} \\ l_{12} & m_{12} \end{pmatrix} + \begin{pmatrix} & & \\ & u_4 & \\ l_5 & & \\ & & \\ & u_8 & \\ l_9 & & \end{pmatrix}$$

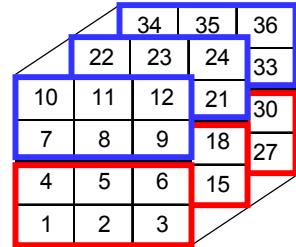
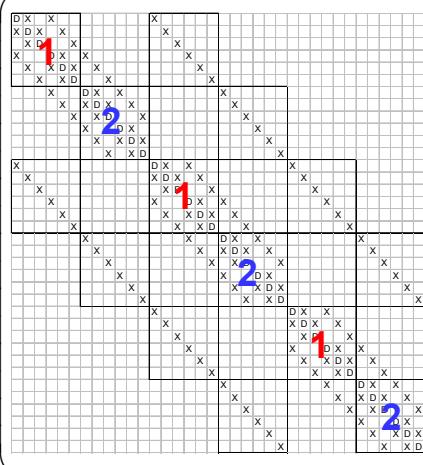


Domain Decomposition
Slide 20

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S

A Real Heptadiagonalmatrix



Domain Decomposition Adamidis/Bönisch
Slide 21 Höchstleistungsrechenzentrum Stuttgart

H L R I S

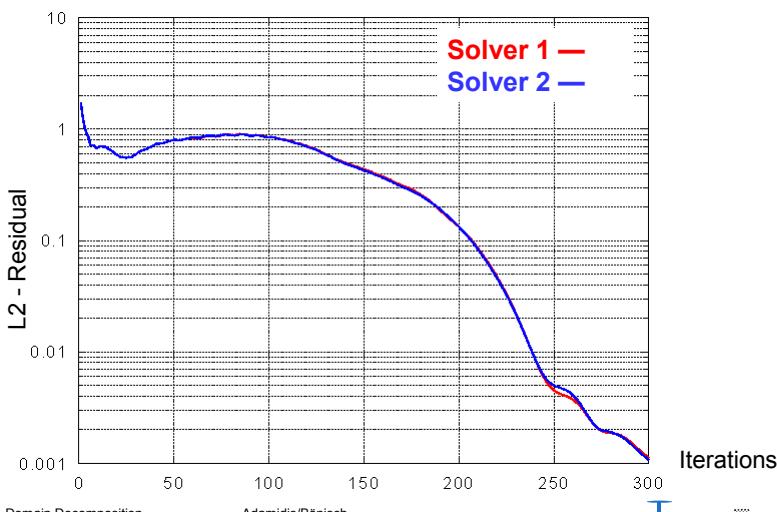
Difference between strong / weak coupling

- Solver with weak coupling
 - Extra computational effort due to additional solving step (but no factorization)
 - Additional update of right hand side
 - Two times communication
 - one after each solving step
- Solver with stronger coupling
 - Jacobi line relaxation method with subiterations
 - collapse iterations from line relaxation and parallelization
 - no additional iterations
 - much more communication

Domain Decomposition Adamidis/Bönisch
Slide 22 Höchstleistungsrechenzentrum Stuttgart

H L R I S

Comparison of Solvers - Convergence



Domain Decomposition

Slide 23

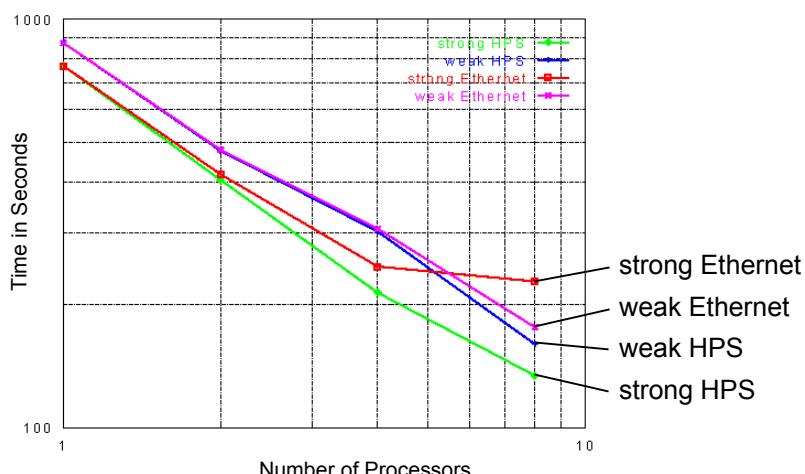
Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S



Comparison of Solvers - Performance



Domain Decomposition

Slide 24

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

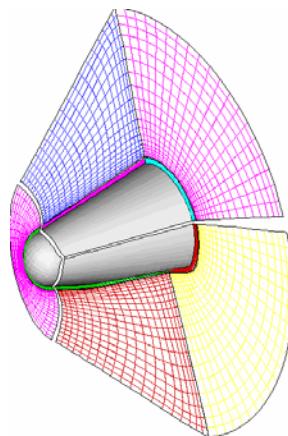
H L R I S



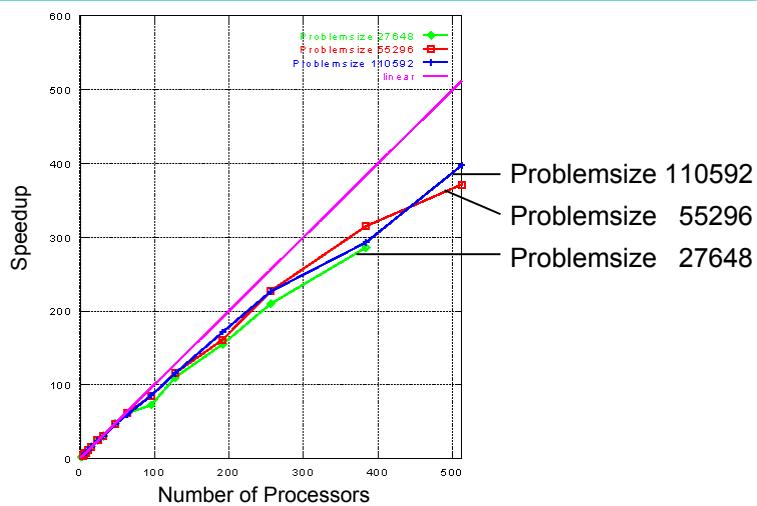
Results - Solving Method

- the presented solver with weak coupling works fine for this CFD problems
- Solutions differ in the scale of one percent
- convergence rate is nearly equal to the sequential program

Domain Decomposition



Speedup on Cray T3E



Domain Decomposition

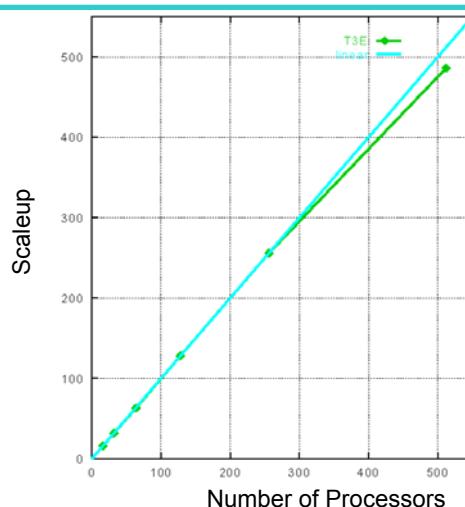
Slide 27

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart



Scaleup on Cray T3E



Domain Decomposition

Slide 28

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

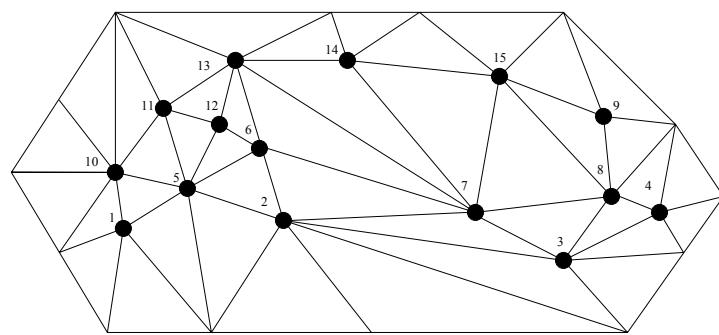


Domain Decomposition of Unstructured Grids

Höchstleistungsrechenzentrum Stuttgart

H L R I S

Unstructured FEM Grid with Global Numbering



Shape of Corresponding System Matrix

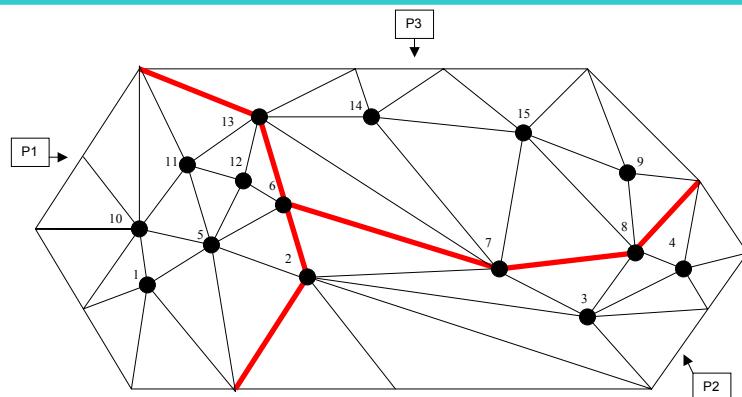
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	x				x				x						
2		x	x		x	x	x								
3		x	x	x				x	x						
4			x	x					x	x					
5	x	x			x	x				x	x	x			
6		x			x	x	x				x	x			
7		x	x			x	x	x				x	x	x	
8			x	x			x	x	x				x		
9			x			x	x						x		
10	x				x			x	x						
11		x					x	x	x	x					
12		x	x					x	x	x					
13			x	x				x	x	x	x				
14				x					x	x	x				
15					x	x	x			x	x				

x = edge of FEM grid

Domain Decomposition

1. Nonoverlapping Domain Decomposition
2. Grid Points Separated into Inner and Boundary Points
3. Renumbering of the Inner and Boundary Points

1. Nonoverlapping Domain Decomposition

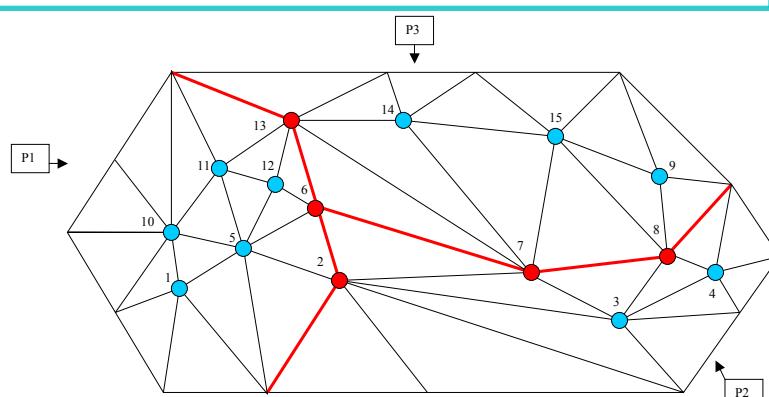


Domain Decomposition
Slide 33

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S

2. Grid Points Separated into Inner and Boundary Points

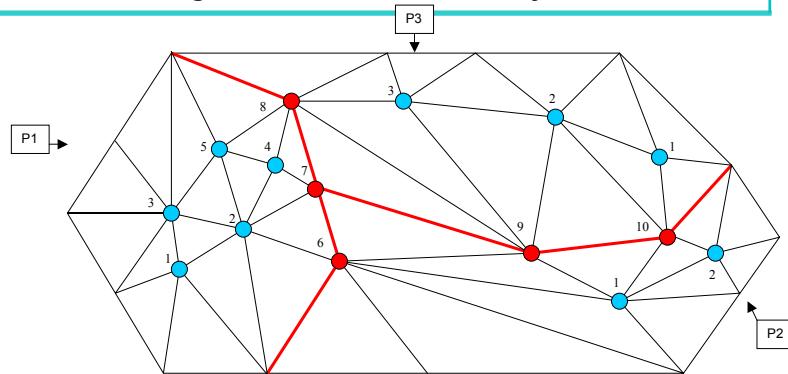


Domain Decomposition
Slide 34

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S

3. Renumbering of the Inner and Boundary Points

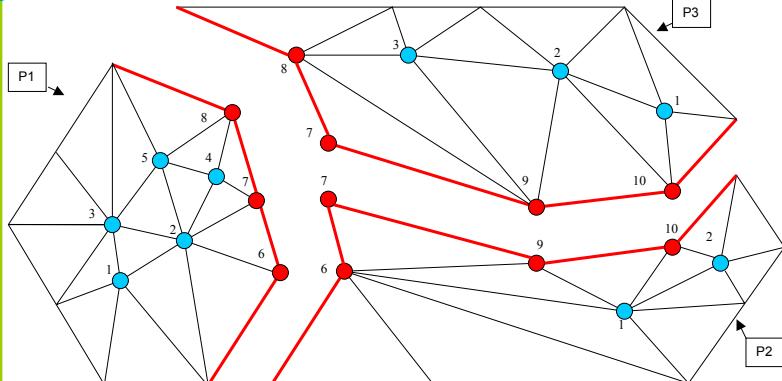


- Inner points are numbered from 1 up to the number of inner points
 - Boundary points are numbered globally starting from the maximum of all inner points of all partitions plus 1

Arrow Shaped System Matrix after Renumbering

	P1			P2			P3			logical boundary					
	1	2	3	4	5	1	2	1	2	3	6	7	8	9	10
1	x	x	x								x	x			
2	x	x	x	x	x										
3	x	x	x												
4		x	x	x							x	x			
5		x	x	x	x							x			
1				x	x						x		x	x	
2				x	x									x	
1						x	x								x
2						x	x	x	x			x	x		
3						x	x	x	x			x	x		
6	x			x							x	x	x		
7	x	x									x	x	x	x	
8		x	x					x			x	x	x	x	
9			x				x	x	x	x	x	x	x	x	
10				x	x	x	x	x	x	x			x		

Data Distribution



Domain Decomposition

Slide 37

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S



Linear System

$$\begin{pmatrix} A_H^1 & & A_{IB}^1 \\ A_H^2 & \ddots & A_{IB}^2 \\ & \ddots & \ddots & \ddots \\ & & A_H^n & A_{IB}^n \\ A_{BI}^1 & A_{BI}^2 & \cdots & A_{BI}^n & A_{BB} \end{pmatrix} \begin{pmatrix} x_I^1 \\ x_I^2 \\ \vdots \\ x_I^n \\ x_B \end{pmatrix} = \begin{pmatrix} b_I^1 \\ b_I^2 \\ \vdots \\ b_I^n \\ b_B \end{pmatrix}$$

with $A_{BB} = \sum_{i=1}^n A_{BB}^i$

- Available on local memories

$$\begin{pmatrix} A_H^i & A_{IB}^i \\ A_{BI}^i & A_{BB}^i \end{pmatrix} \quad i = 1, \dots, n$$



Domain Decomposition

Slide 38

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S



Direct Substructuring

1. Direct factorization
2. Assembling of the Schur complement system
3. Solving of the Schur complement system
4. Solving of the interior unknowns



Domain Decomposition
Slide 39

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S



Transformations of the Original System

$$A = \begin{pmatrix} I & & & & \\ & I & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & I \end{pmatrix} \begin{pmatrix} A_{II}^1 & & & & \\ & A_{II}^2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & S^1 + \dots + S^n \end{pmatrix} \begin{pmatrix} I & & & A_{II}^{(1)-1} & A_{IB}^1 \\ & I & & A_{II}^{(2)-1} & A_{IB}^2 \\ & & \ddots & & \ddots \\ & & & \ddots & I \end{pmatrix}$$

$$S^i = A_{BB}^i - A_{BI}^i A_{II}^{(i)-1} A_{IB}^i$$

$$\begin{pmatrix} A_{II}^1 & & & & \\ & A_{II}^2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & A_{BI}^1 & A_{BI}^2 & & I \end{pmatrix} \begin{pmatrix} I & & & A_{II}^{(1)-1} & A_{IB}^1 \\ & I & & A_{II}^{(2)-1} & A_{IB}^2 \\ & & \ddots & & \ddots \\ & & & \ddots & S^1 + \dots + S^n \end{pmatrix} \begin{pmatrix} x_I^1 \\ x_I^2 \\ \vdots \\ x_b \end{pmatrix} = \begin{pmatrix} b_I^1 \\ b_I^2 \\ \vdots \\ b_b \end{pmatrix}$$



Domain Decomposition
Slide 40

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S



Schurcomplement System

$$A_{II}^i x_I^i = b_I^i - A_{IB}^i x_B$$

$$A_{SH} x_B = b_{SH}$$

$$A_{SH} = \sum_{i=1}^n \left(A_{BB}^i - A_{BI}^i (A_{II}^i)^{-1} A_{IB}^i \right)$$

$$b_{SH} = b_B - \sum_{i=1}^n A_{BI}^i (A_{II}^i)^{-1} b_I^i$$

or

$$A_{SH} = \sum_{i=1}^n \left(A_{BB}^i - A_{BI}^i Z_i \right)$$

$$Z_i = (A_{II}^i)^{-1} A_{IB}^i \quad A_{II}^i Z_i = A_{IB}^i$$

$$b_{SH} = b_B - \sum_{i=1}^n A_{BI}^i z_i$$

where

$$z_i = (A_{II}^i)^{-1} b_I^i \quad A_{II}^i z_i = b_I^i$$

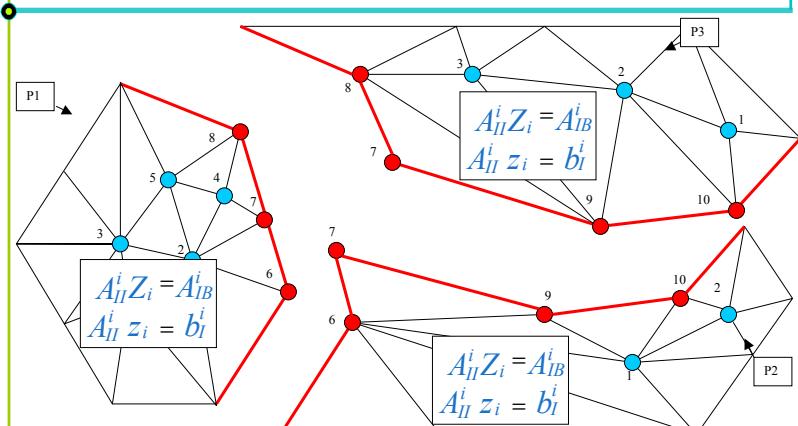
H L R I S



Domain Decomposition
Slide 41

Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

1. Direct factorization

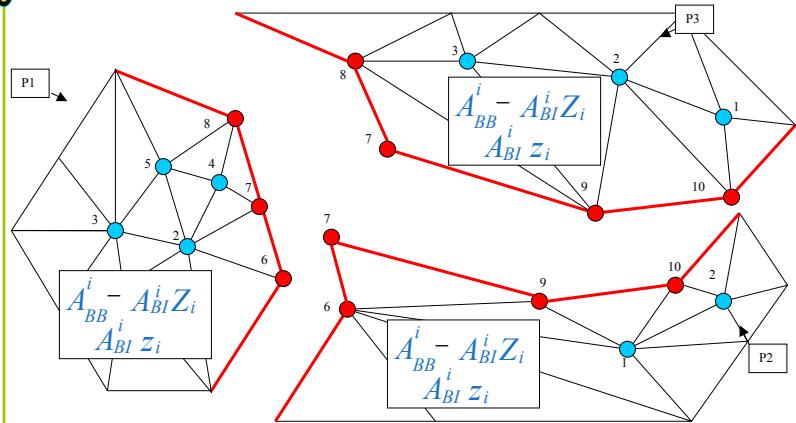


Domain Decomposition
Slide 42

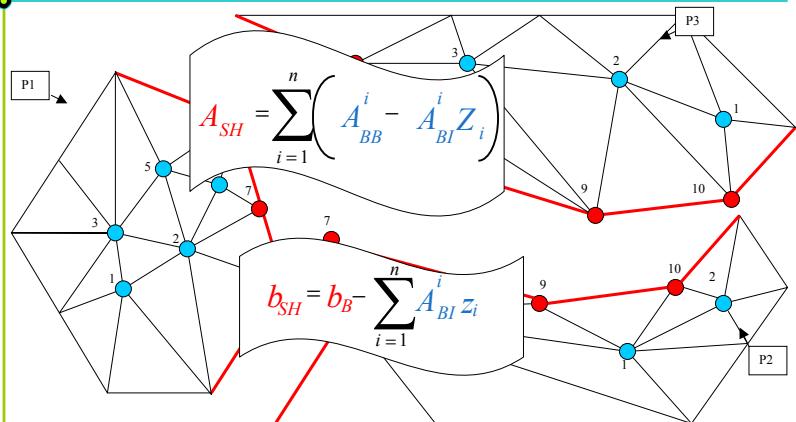
Adamidis/Bönisch
Höchstleistungsrechenzentrum Stuttgart

H L R I S

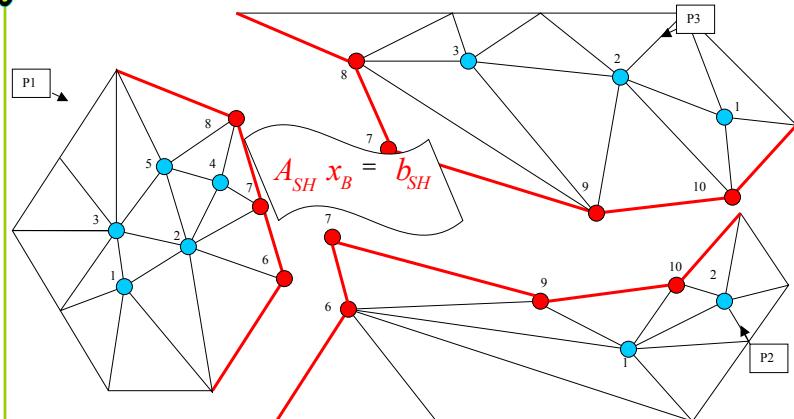
2. Assembling of the Schur complement system (I)



2. Assembling of the Schur complement system (II)



3. Solving of the Schur complement system



Domain Decomposition

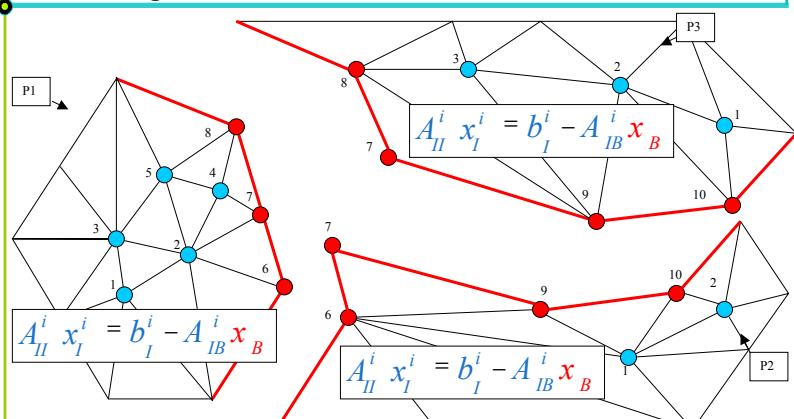
Slide 45

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S

4. Solving of the interior unknowns



Domain Decomposition

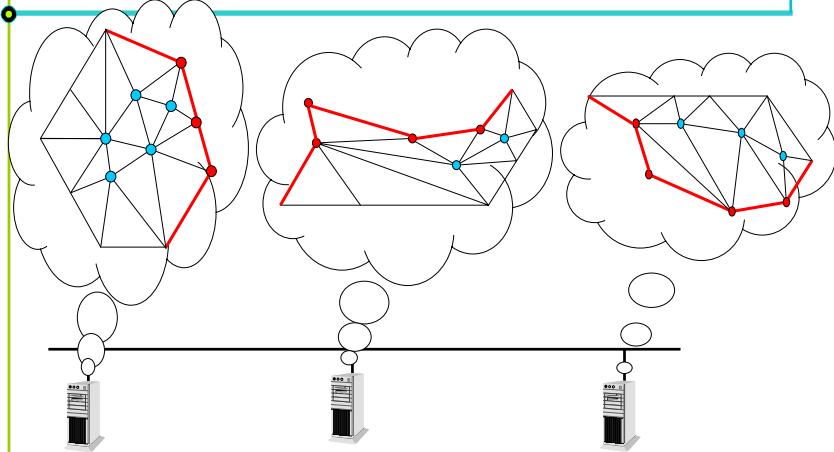
Slide 46

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S

Parallel Computations on the Subdomains



Domain Decomposition

Slide 47

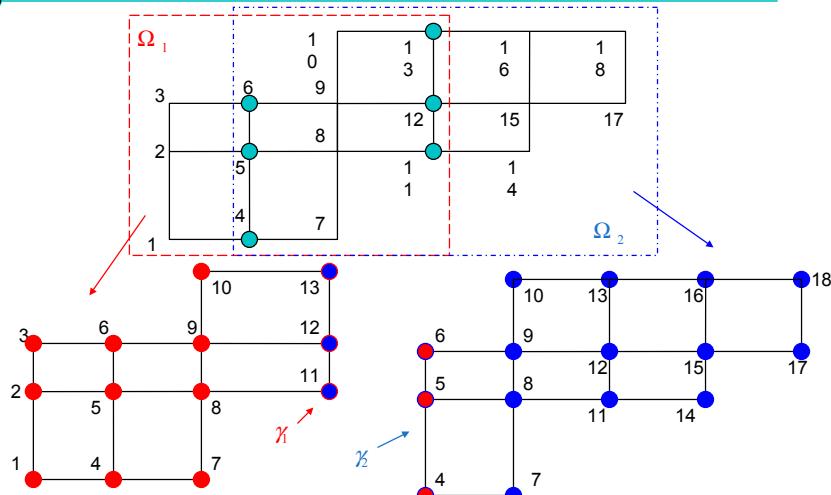
Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S



Overlapping Domain Decomposition



Domain Decomposition

Slide 48

Adamidis/Bönisch

Höchstleistungsrechenzentrum Stuttgart

H L R I S



Literature

- Barry F. Smith, Petter E. Bjørstad, William D. Gropp:
Domain Decomposition, Parallel Multilevel Methods for Elliptic Partial Differential Equations, Cambridge University Press, 1996,
ISBN 0-521-49589-X
- David E. Keyes, Youcef Saad, Donald G. Truhlar:
Domain-Based Parallelism and Problem Decomposition Methods in Computational Science and Engineering, SIAM, 1995,
ISBN 0-89871-348-X