

# **Hardware-oriented Numerics**

## **for PDEs**

---

### **Consequences for**

## **Incompressible Flow Solvers**

**S. Turek + FEAST GROUP**

M. Altieri, Chr. Becker, S. Buijssen, S. Kilian, D. Kuzmin, F. Platte  
J. Hron, P. Kotalik, W. Decheng, A. Ouazzi, R. Schmachtel, L. Rivkind...

Institut für Angewandte Mathematik & Numerik  
Universität Dortmund

<http://www.feastflow.de>

## Trends in Numerics for PDEs:

*'A posteriori error control/adaptive meshing'*

*'Iterative (parallel) solution strategies'*

*'Operator-splitting for coupled problems'*

**Reduction of numerical complexity !!!**



## Trends in Processor Technology:

*'Enormous improvements in Processing Data'*

*'Much lower advances in Moving Data'*

**1 PC in 10 years  $\approx$  1 CRAY T3E today !!!**

## Questions:

‘Can we use this enormous computing power?’

‘Particularly: For Numerics for PDEs ???’

‘If not, how to achieve a significant percentage?’



Main components in iterative schemes:

### **Matrix-Vector applications (MV)**

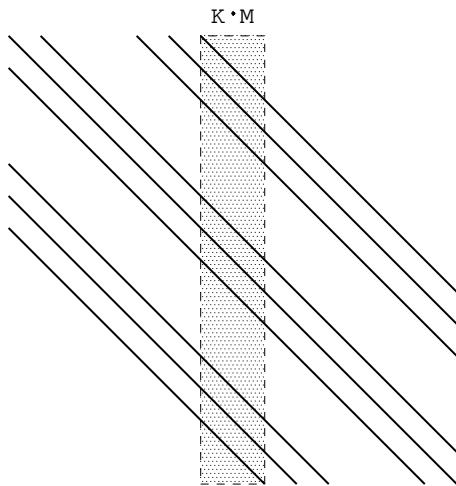
- Krylov-space methods, **Multigrid**, etc.
- **Smoothing**, defect calculations, step-length control, etc.
- Often consuming (at least) **60 - 90%** of CPU time
- **How many MFLOP/s are realistic ?**

# SPARSE Matrix-Vector techniques

```
DO 10 IROW=1,N  
DO 10 ICOL=KLD(IROW),KLD(IROW+1)-1  
10      Y(IROW)=DA(ICOL)*X(KCOL(ICOL))+Y(IROW)
```

- Standard technique in most **FEM** or **FV** codes
- Storage of '**non-zero**' matrix elements only (*CSR*, *CSC*, ...)
- Access via **index vectors**, **linked lists**, **pointers**, etc
- Comparable with 'indexed DAXPY'

## SPARSE BANDED MV techniques



- Typical for **FD** codes on '**tensorproduct meshes**'
- Storage of '**non-zero**' elements in **bands/diagonals**
- MV multiplication 'bandwise' and 'window-oriented'

# Results on locally structured meshes

2D case	N	DAXPY-I	SBB-V	SBB-C	MG-V	MG-C
DEC 21264 (667 MHz) 'ES40'	$65^2$	205 (178)	538	795	370	452
	$257^2$	224 (110)	358	1010	314	487
	$1025^2$	<b>78 (11)</b>	<b>158</b>	<b>813</b>	<b>185</b>	<b>401</b>
HITACHI (375 MHz) 'SR8000'	$65^2$	173 (82)	238	391	191	266
	$257^2$	143 (29)	243	388	198	260
	$1025^2$	<b>144 (7)</b>	<b>226</b>	<b>390</b>	<b>200</b>	<b>267</b>
AMD K7 (850 MHz) 'ATHLON'	$65^2$	203 (195)	101	556	122	355
	$257^2$	29 (27)	78	241	72	166
	$1025^2$	<b>31 (10)</b>	<b>64</b>	<b>236</b>	<b>58</b>	<b>126</b>

## SPARSE MV techniques (DAXPY-I)

- MFLOP/s rates far away from '**Peak Performance**'
- Depending on **problem size + numbering**
- PC partially **faster** than processors in 'supercomputers' !!!

## SPARSE BANDED MV techniques (SBB)

- 'Supercomputing' power gets visible (up to **1 GFLOP/s**)
- FEM-Simulation for **complex domains** ???

# FEAST project

**Implementation** techniques and  
**Numerics (!)** adapted to **hardware !!!**



Precise knowledge of **processor** characteristics  
for different tasks in **MG** for **FEM** spaces:

*(Collection of) FEAST INDICES*



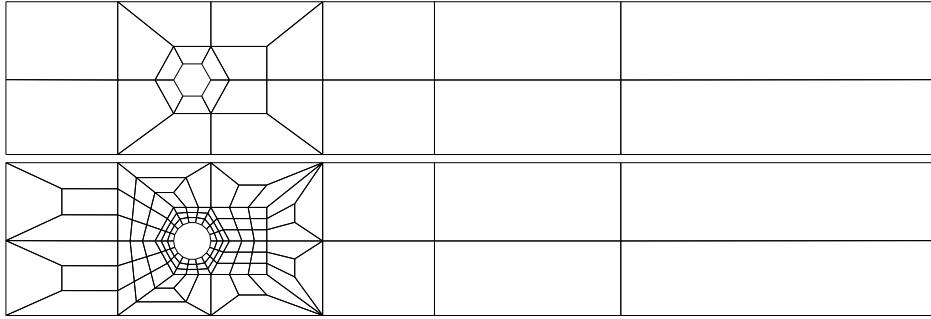
- 1) *Get the optimum (from the hardware)!!!*
- 2) *Prevent from dramatic performance losses !!!*



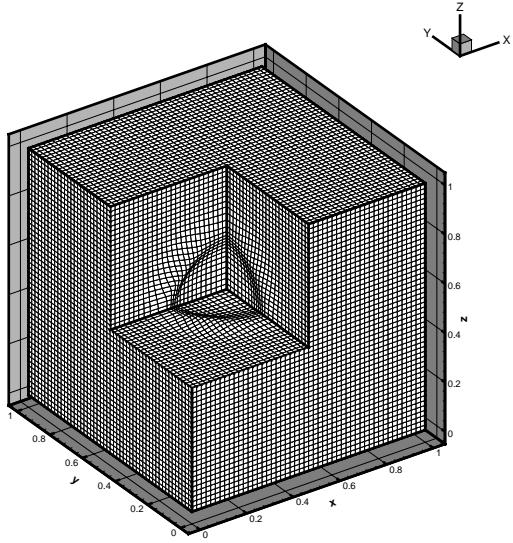
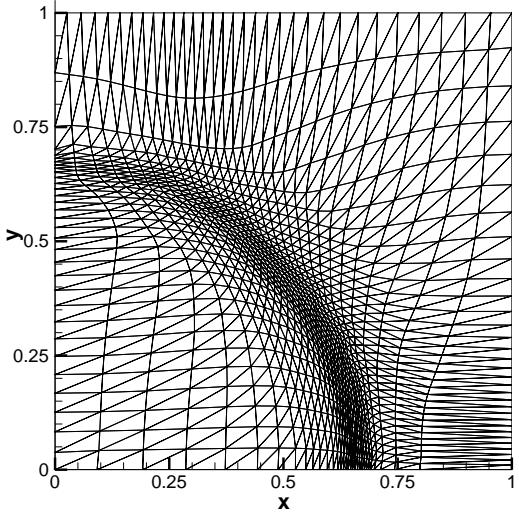
**Hardware-oriented Numerics ???**

# 1) Concepts for adaptive meshing

*1) macro-oriented adaptivity*



*2) patchwise adaptivity*



*3) local adaptivity*

→ hanging nodes, triangles/quads, . . .

# ‘Exploit the locally nice structures !’

- Local use of efficient SPARSE BANDED BLAS techniques
- Local superconvergence through orthogonal meshes
- Local storage cost small through the use of matrix stencils

↑↑

## Open problems:

### 1. ‘Optimality’ of the mesh ?

→ *w.r.t. number of unknowns or CPU time ?*

### 2. Error estimators ?

→ *degree of macro-refinement ?*

→ *anisotropic refinement ?*

→ *h/p-refinement ?*

## II) Concepts for iterative (parallel) solvers

### 1) Standard Multigrid

- parallelization of ‘recursive’ smoothers (only blockwise) ?
- complicated geometric structures with local anisotropies ?
- **too few arithmetic work vs. data exchange !**

### 2) Standard Domain Decomposition

- good ratio for communication/aritm. work !
- implementation (overlap, coarse grid problem, **3D**) ?
- **bad convergence behaviour w.r.t. multigrid !**



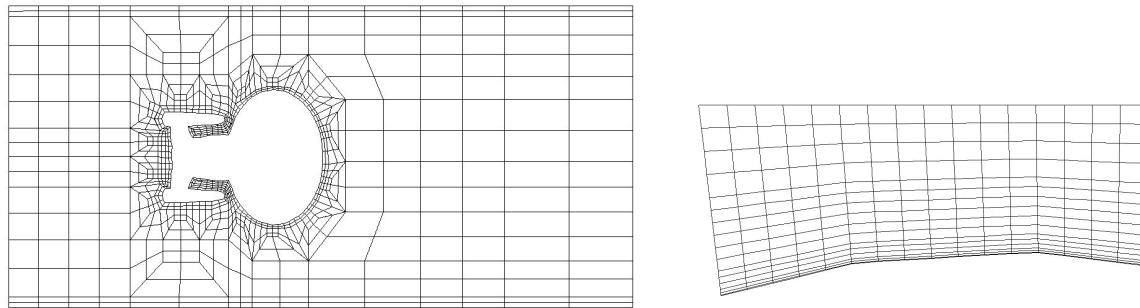
$$1) + 2) = \text{SCARC}$$

*Hide recursively all anisotropies in more  
'local' units! (**robustness**)*

*Perform all Linear Algebra tasks on 'local'  
units only! (**efficiency**)*

# Example: Realization of SCARC in FEAST

2D decomposition and zoomed (macro) element (LEVEL 3) with locally anisotropic refinement towards the wall

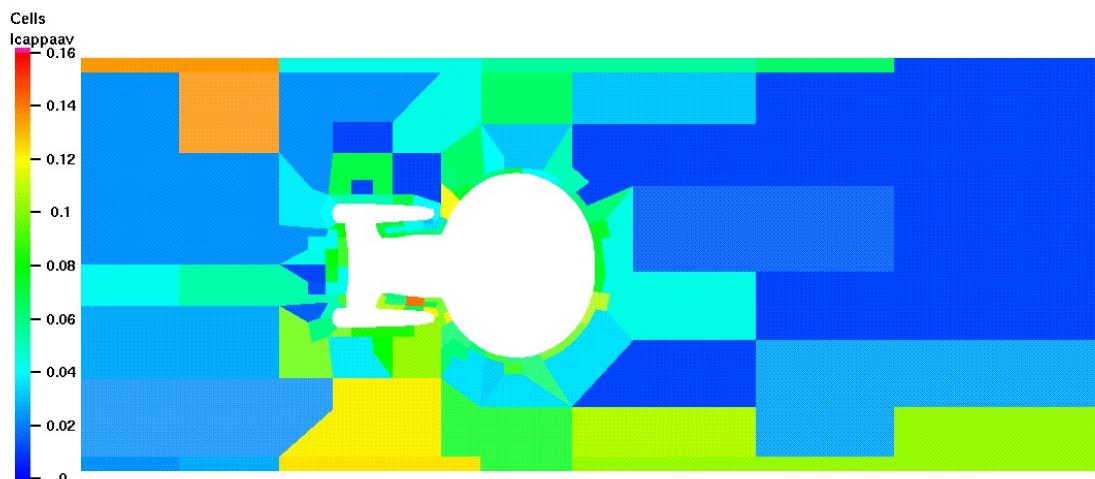


SCARC-CG solver (smoothing steps: 1 global SCARC ; 1 local ‘MG-TriGS’) for locally (an)isotropic refinement

## Global (parallel) convergence rates

#NEQ	Dirichlet 'Velocity'		Neumann 'Pressure'	
	$AR \approx 10$	$AR \approx 10^6$	$AR \approx 10$	$AR \approx 10^6$
210,944	0.17 (8)	0.18 (8)	0.21 (9)	0.15 (8)
843,776	0.17 (8)	0.17 (8)	0.20 (9)	0.17 (8)
3,375,104	0.18 (9)	0.19 (9)	0.22 (10)	0.22 (10)
13,500,416	0.19 (9)	0.18 (9)	0.23 (10)	0.23 (10)

## Local convergence rates (for $AR \approx 10^6$ )



### III) MPSC: Incompressible Flow Solvers

#### *Results with FEATFLOW*

Computer	total time	'memory intensive'	'floating point'
IBM SP2 (160 MHz)	<b>2748</b>	1587 ( <b>58%</b> )	1161 ( <b>42%</b> )
PC PII (400 MHz)	<b>4927</b>	2401 ( <b>49%</b> )	2526 ( <b>51%</b> )
IBM SP2 (66 MHz)	<b>5970</b>	3824 ( <b>64%</b> )	2146 ( <b>36%</b> )

- ‘Discrete Projection’: Burgers + Pressure Poisson
- Streamline-Diffusion FEM discretization in 3D
- Implicit adaptive time stepping
- *Faster solvers useless!*



How to increase the  
‘floating point’ intensive parts ?

*Larger time steps !*

**How ???**

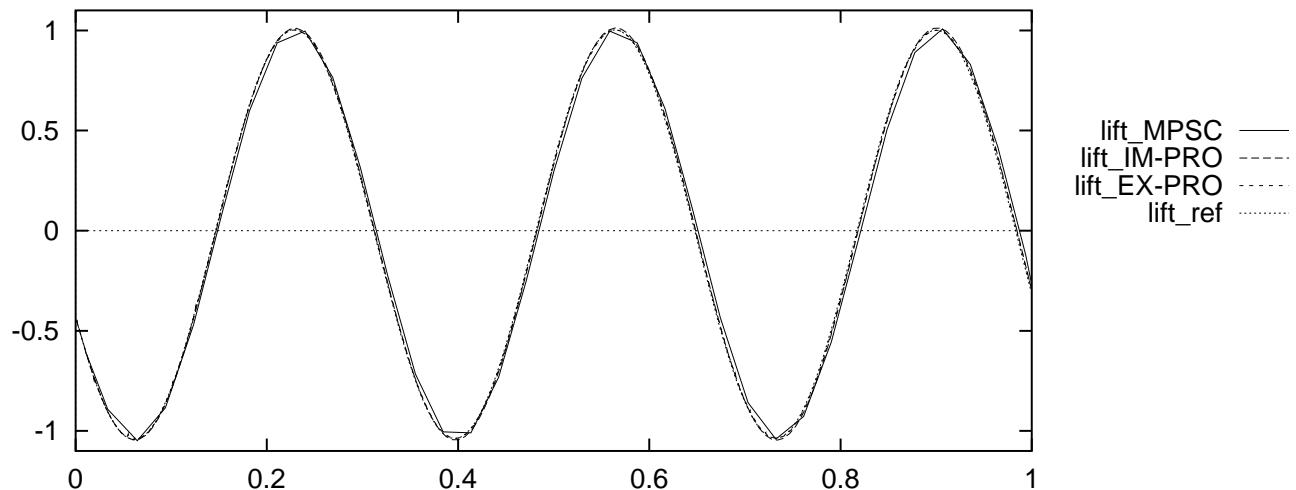
# Stronger velocity-pressure coupling:

*'Perform more arithmetic work with the once assembled matrices and vectors'!*



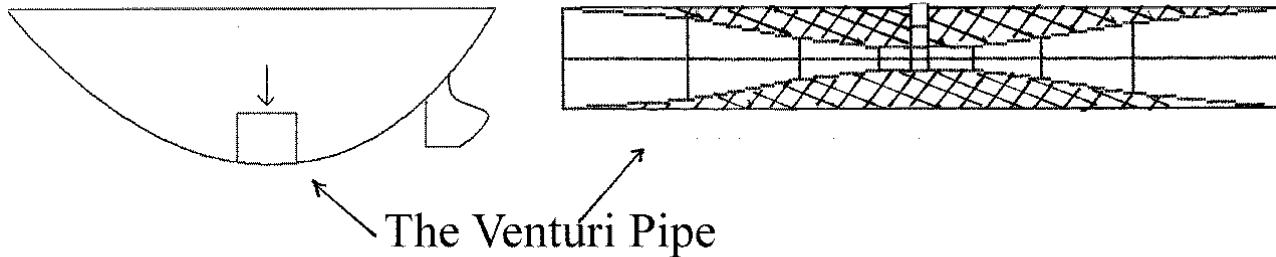
## Multilevel Pressure Schur Complement

CPU (solvers)	Method	#NT	Lift		Drag	
			mean	peak	mean	peak
14,358 ( <b>81%</b> )	fully impl. MPSC	39	1%	1%	0%	2%
42,679 ( <b>51%</b> )	semi-impl. DPM	165	0%	0%	0%	0%
64,485 ( <b>54%</b> )	semi-expl. DPM	889	0%	8%	0%	0%



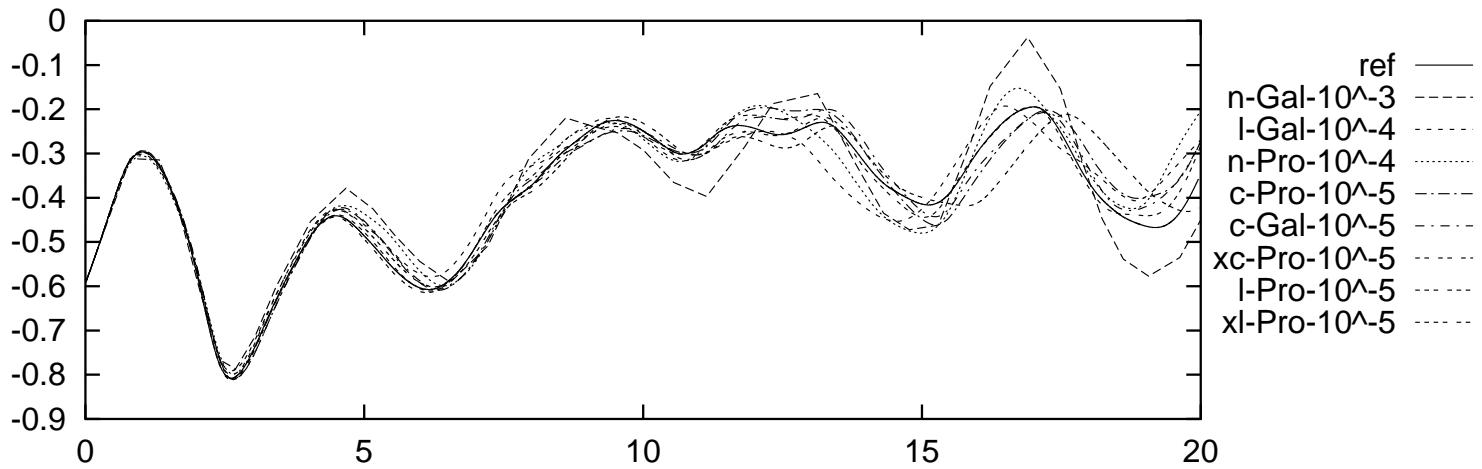
*'Ready' for SPARSE BANDED BLAS  
'Ready' for Galerkin-type error control*

# Test: Flow through venturi Pipe



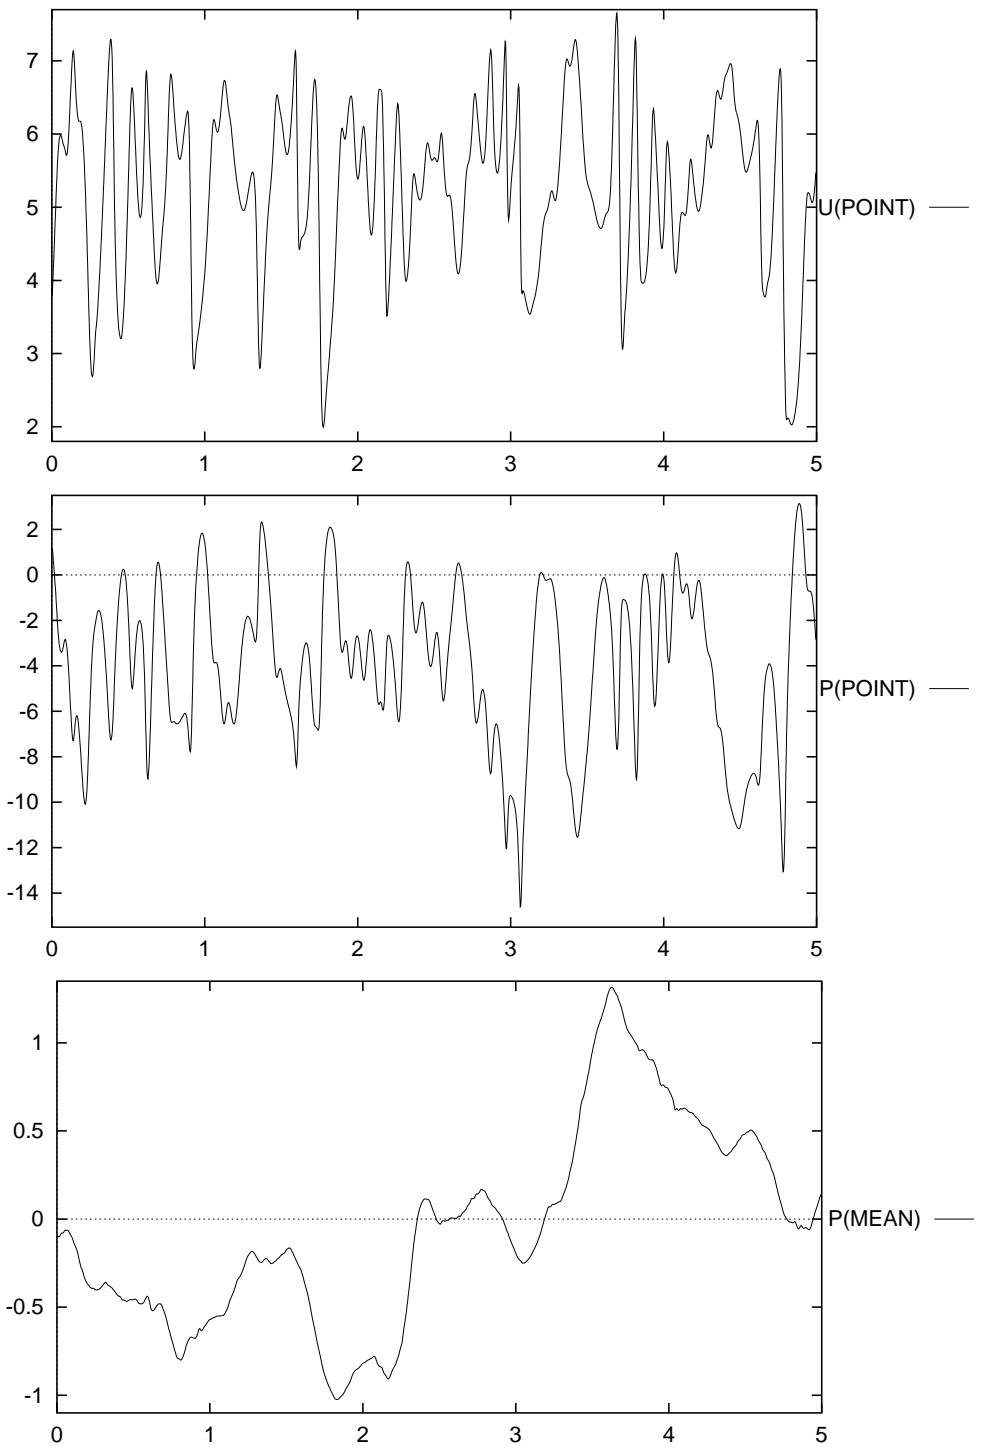
**Adaptive** time step control for **velocity** (global,  $l_2$ ) and **flux** (through upper device !!!)

$$|f(\mathbf{u}_h^k(t_n), p_h^k(t_n)) - f(\mathbf{u}(t_n), p(t_n))| \leq TOL$$



	Meth.	#NT	Flux		Pressure	
			$l_2$	mean	$l_2$	mean
	n-MPSC	<b>38</b>	15%	1%	19%	1%
O	l-MPSC	<b>222</b>	11%	2%	13%	0%
K	n-DPM	<b>263</b>	11%	2%	14%	3%
	l-DPM	<b>649</b>	4%	0%	5%	0%
	xl-DPM	<b>6941</b>	11%	0%	14%	0%

# Test: Perpendicular inflow into pipe



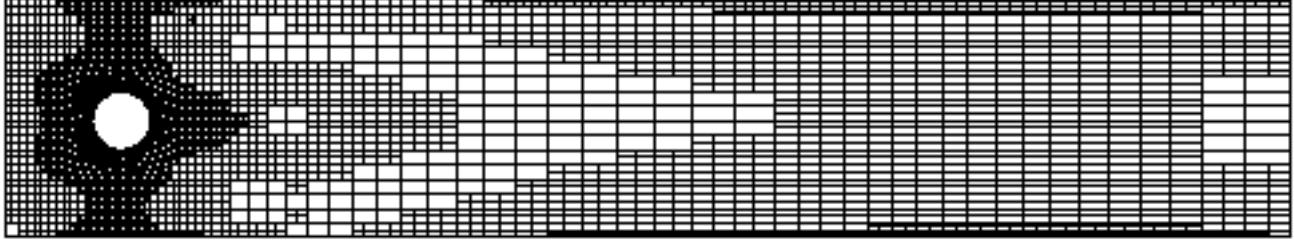
Rigorous adaptive error control of  
**mean values** ('wall pressure') in  
space/time ?!?

# Hardware-Oriented Numerics for PDEs:

## 1) *Patch-oriented adaptivity*

‘Many’ (local) **tensorproduct grids** (S-B BLAS)

‘Few’ (local) **unstructured grids** (SPARSE)



## 2) *Generalized MG-DD solver: SCARC*

Exploit locally ‘regular’ structures (**efficiency**)

Recursive ‘clustering’ of anisotropies (**robustness**)

‘Strong local solvers improve global convergence !’

## 3) *Navier-Stokes solvers: Full Galerkin MPSC*

Develop Navier-Stokes solvers with more ‘arithmetic’ than ‘memory intensive’ operations

‘Exploit locally regular structures !!!’

## ‘My’ conclusions:

*Everybody can/must test the absolute (!)  
computational performance!*

*Compromise between ‘flexible/reusable/abstract’  
and ‘machine-dependent/hardware-oriented’  
implementation styles is required!*

*Modern Numerics has to consider recent and  
future hardware trends!*



**Future: Huge Potential**

*‘Hardware-Oriented Numerics’*