

HPC for Environmental Simulations

PD Dr. rer. nat. habil. Ralf-Peter Mundani

Computation in Engineering / BGU

Scientific Computing in Computer Science / INF

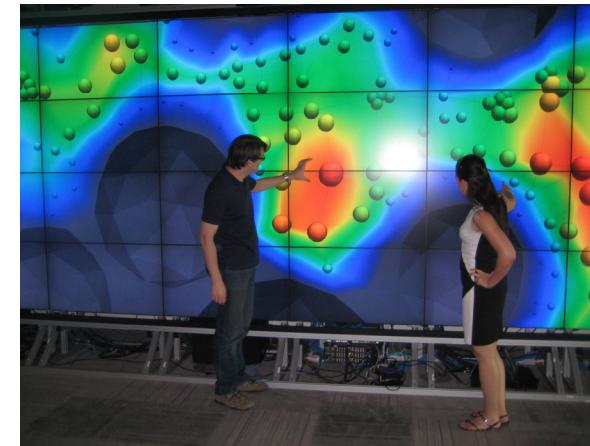
18th International Symposium on Symbolic and Numerical Algorithms for Scientific Computing

September 24–27, 2016

Timisoara, Romania

Motivation

"High-performance computing must now assume a broader meaning, encompassing not only flops, but also the ability, for example, to efficiently manipulate vast and rapidly increasing quantities of both numerical and non-numerical data."



The White House.
High-performance-computing



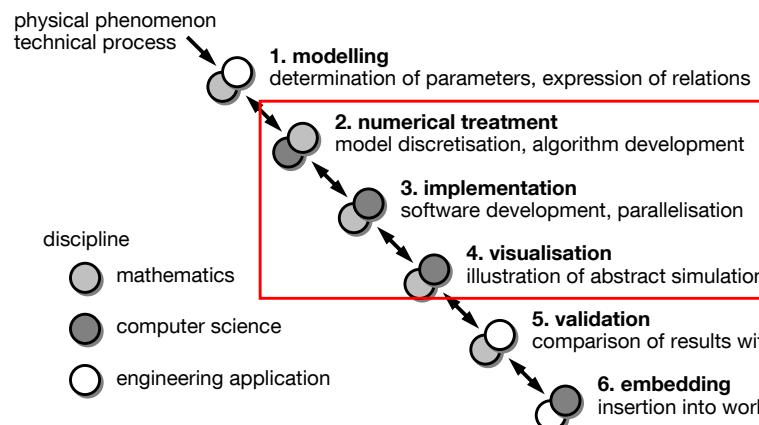
AESOP: 40×46" NEC panels
with total res. of 13,600 ×
3,072 pixels (~42 MPixel)

PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

2

Motivation

- simulation – from phenomena to prediction



Motivation

- why parallel programming and HPC?
 - complex problems (especially the so called 'grand challenges') demand for more computing power
 - climate or geophysics simulation (tsunami, e.g.)
 - structure or flow simulation (crash test, e.g.)
 - development systems (CAD, e.g.)
 - large data analysis (Large Hadron Collider at CERN, e.g.)
 - military applications (crypto analysis, e.g.)
 - ...
 - performance increase due to
 - faster hardware, more memory ('work harder')
 - more efficient algorithms, optimisation ('work smarter')
 - parallel computing ('get some help')

PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

4

Motivation

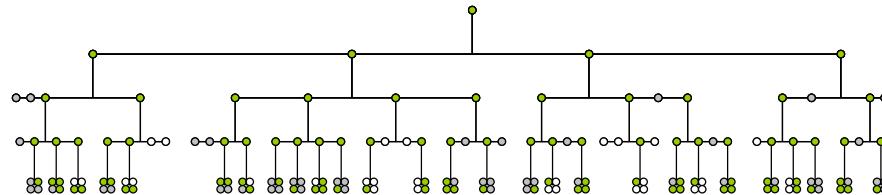
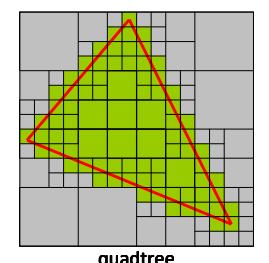
- **objectives** (in case all resources would be available N -times)
 - **throughput**: compute N problems simultaneously
 - running N instances of a sequential program with different data sets ('embarrassing parallelism'); SETI@home, e.g.
 - **response time**: compute one problem at a fraction ($1/N$) of time
 - running one instance (i.e. N processes) of a parallel program for jointly solving a problem; finding prime numbers, e.g.
 - **problem size**: compute one problem with N -times larger data
 - running one instance (i.e. N processes) of a parallel program, using the sum of all local memories for computing larger problem sizes; iterative solution of SLE, e.g.

▪ overview

- geometric and physical modelling
- foundations / parallel architectures
- multigrid methods
- towards massive parallel HPC...
- interactive visual data exploration

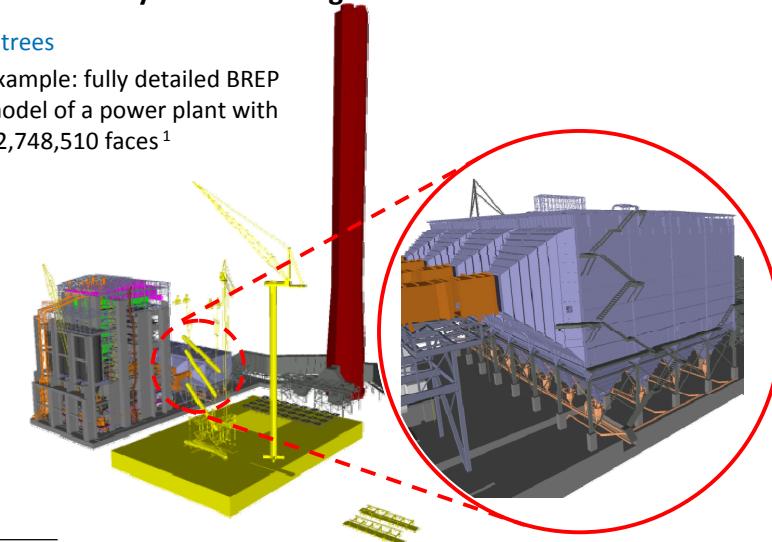
Geometric and Physical Modelling

- **spacetrees**
 - hierarchical data structure (cf. quadtrees in 2D and octrees in 3D)
 - built via recursive bi-section in every dimension
→ 2^D children / node
 - reduced complexity, i.e. amount of voxels compared to equidistant discretisation $O(N^3)$
→ $O(N)$ in 2D and $O(N^2)$ in 3D on average



Geometric and Physical Modelling

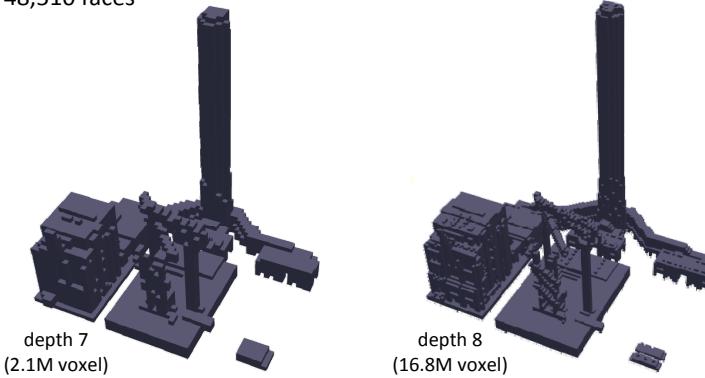
- **spacetrees**
 - example: fully detailed BREP model of a power plant with 12,748,510 faces¹

¹ <http://gamma.cs.unc.edu/POWERPLANT/>

Geometric and Physical Modelling

- **spacetrees**

- example: fully detailed BREP model of a power plant with 12,748,510 faces¹

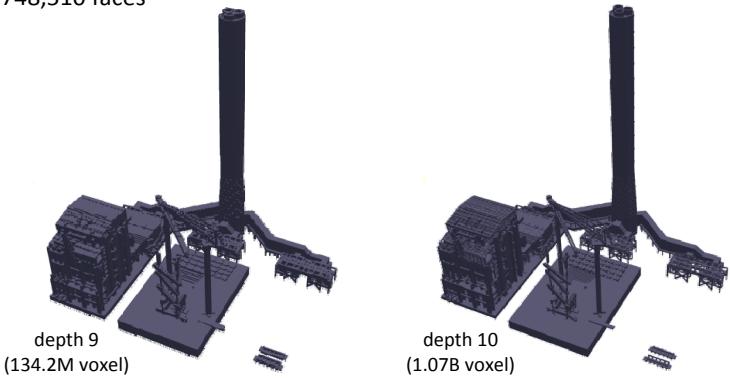


¹ <http://gamma.cs.unc.edu/POWERPLANT/>

Geometric and Physical Modelling

- **spacetrees**

- example: fully detailed BREP model of a power plant with 12,748,510 faces¹

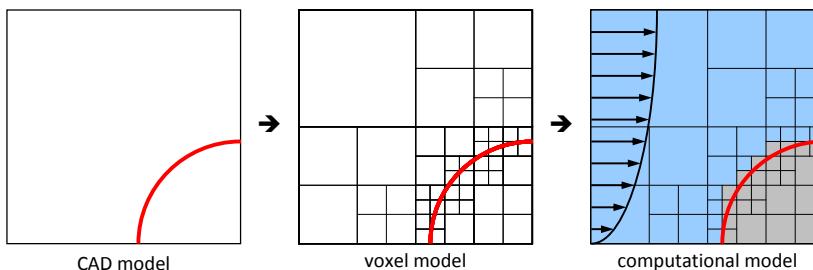


¹ <http://gamma.cs.unc.edu/POWERPLANT/>

Geometric and Physical Modelling

- **generation of computational domain**

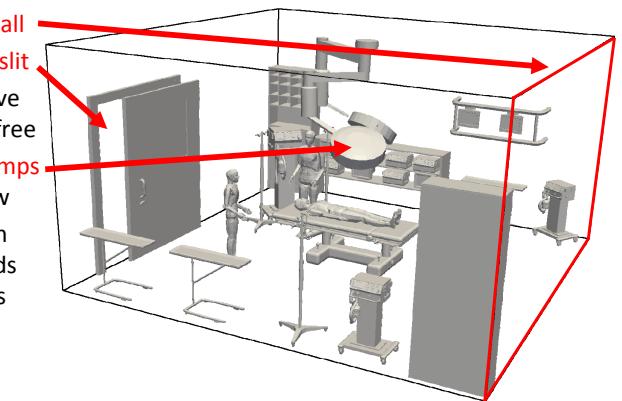
1. discretisation of computational domain
2. tree balancing (1:2) to avoid numerical instabilities
3. setting cell attributes (fluid / obstacle)
4. setting boundary conditions (inflow / outflow / wall / ...)



Geometric and Physical Modelling

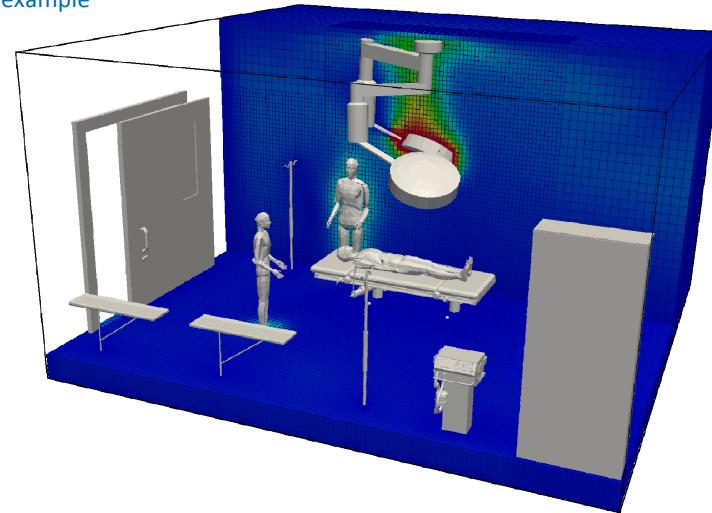
- **complex example**

- operating theatre at 'Klinikum rechts der Isar' (MRI)
- dimensions: 6.30×6.25×3.50 m
- ventilation:
 - inflow: **right wall**
 - outflow: **door slit**
- idea: keep air above patient pollutant-free
- but **hot surgical lamps** influence fluid flow
- thermal simulation using adaptive grids of different depths



Geometric and Physical Modelling

- complex example

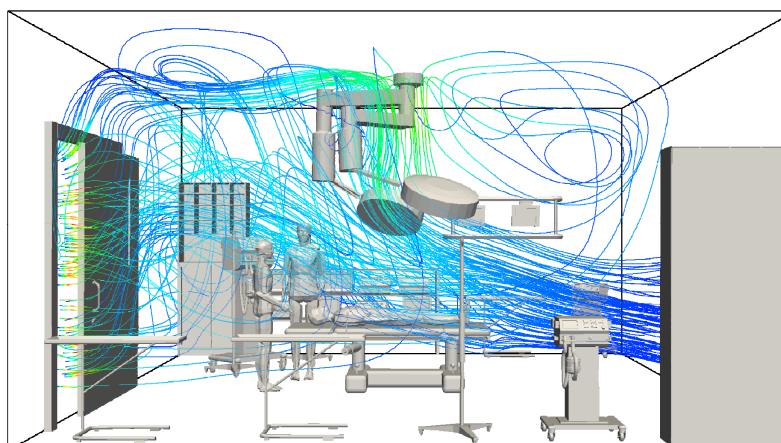


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

13

Geometric and Physical Modelling

- complex example

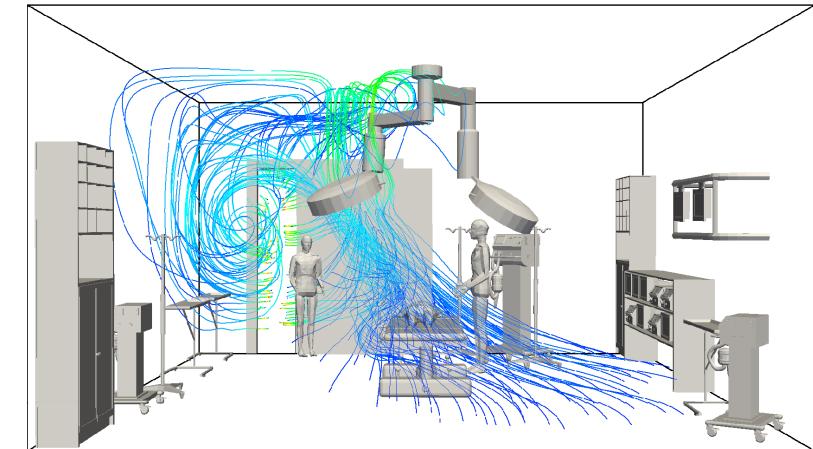


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

15

Geometric and Physical Modelling

- complex example

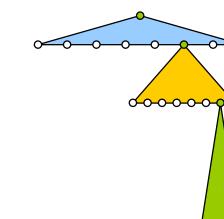


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

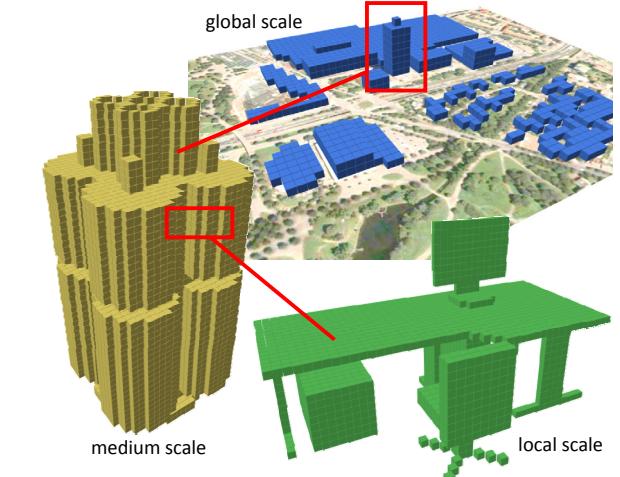
14

Geometric and Physical Modelling

- level of detail concepts



BMW „Vierzyylinder“, Munich

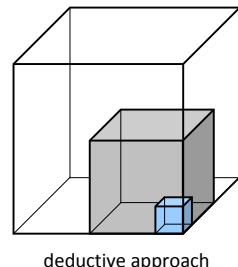


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

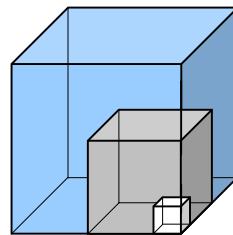
16

Geometric and Physical Modelling

- level of detail concepts
 - towards multiscale simulations



deductive approach



vs.

inductive approach

e.g. flood scenarios / natural disasters

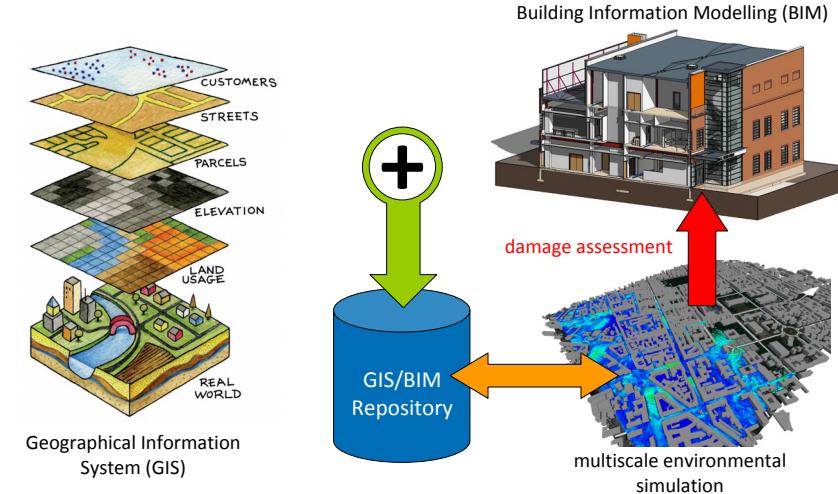
→ local damage assessment

e.g. viral outbreak / artificial disasters

→ global damage assessment

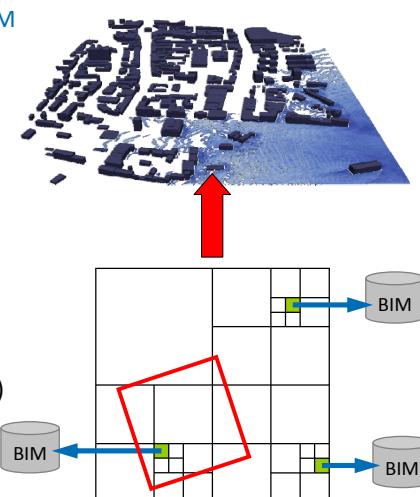
Geometric and Physical Modelling

- bridging worlds and scales: GIS and BIM



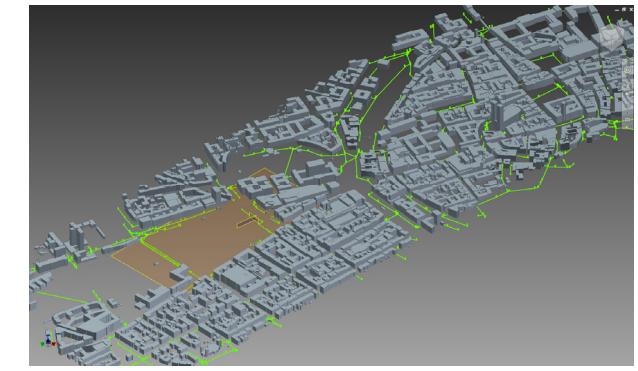
Geometric and Physical Modelling

- bridging worlds and scales: GIS and BIM
 - global scale (GIS)
 - height fields
 - low-fidelity geometries
 - sewerage system
 - local scale (BIM)
 - high-fidelity product models
 - context information
 - GIS/BIM repository (spacetrees)
 - location awareness (proximity)
 - LoD decisions / abstractions
 - selecting region of interest



Geometric and Physical Modelling

- bridging worlds and scales: GIS and BIM
 - coupling with city's sewerage system
 - 3D fluid flows ↔ 1D fluid flow
 - water head from 3D simulation as BC for 1D simulation



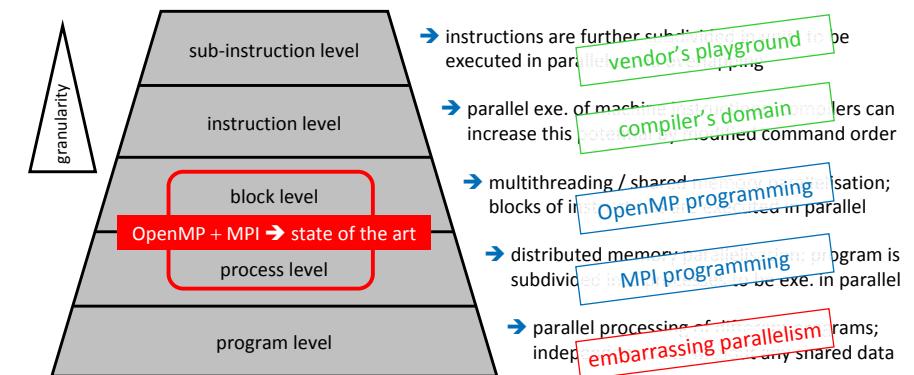
Munich city centre plus sewerage

- overview

- geometric and physical modelling
- foundations / parallel architectures
- multigrid methods
- towards massive parallel HPC...
- interactive visual data exploration

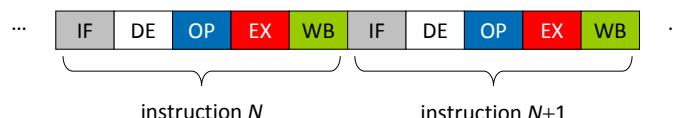
Foundations / Parallel Architectures

- levels of parallelism



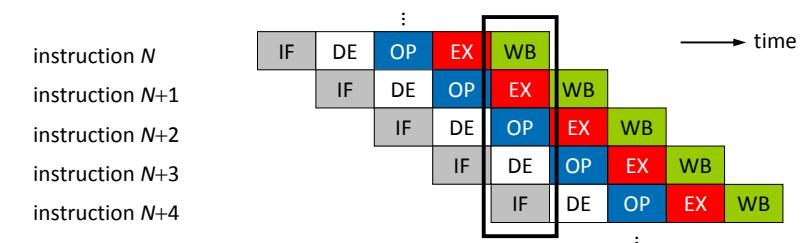
Foundations / Parallel Architectures

- a brief history of time: instruction pipelining
 - instruction execution involves several operations
 1. instruction fetch (IF)
 2. decode (DE)
 3. fetch operands (OP)
 4. execute (EX)
 5. write back (WB)
 - which are *executed successively*
- hence, only one part of CPU works at a given moment



Foundations / Parallel Architectures

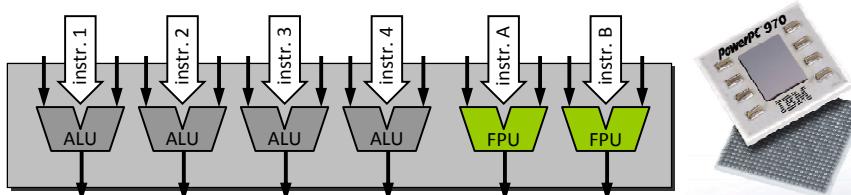
- a brief history of time: instruction pipelining
 - observation: while processing particular stage of instruction, other stages are idle
 - hence, multiple instructions to be overlapped in execution ➔ instruction pipelining (similar to assembly lines)
 - advantage: no additional hardware necessary



Foundations / Parallel Architectures

- a brief history of time: superscalar

- faster CPU throughput due to simultaneously execution of instructions within one clock cycle via redundant functional units (ALU, multiplier, ...)
- dispatcher decides (during runtime) which instructions read from memory can be executed in parallel and dispatches them to different functional units
- for instance, PowerPC 970 ($4 \times$ ALU, $2 \times$ FPU)

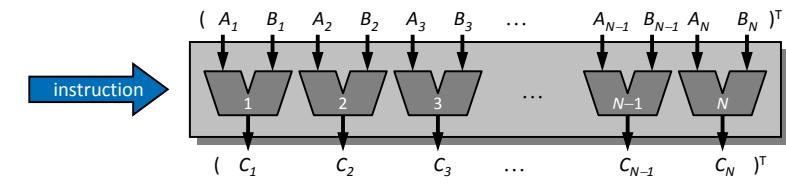


- but, performance improvement is limited (intrinsic parallelism)

Foundations / Parallel Architectures

- a brief history of time: vector units

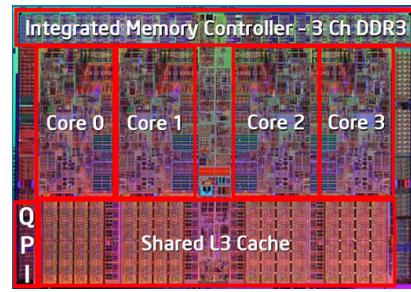
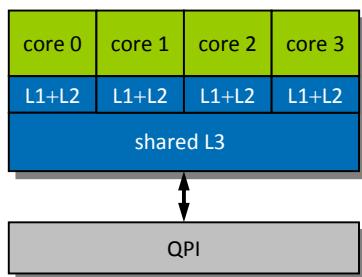
- simultaneously execution of *one instruction* on a *one-dimensional array of data* (=vector)
- VU first appeared in 1970s and were the basis of most supercomputers in the 1980s and 1990s



- specialised hardware → very expensive
- limited application areas (mostly Computational Fluid Dynamics, Computational Structures Dynamics, ...)

Foundations / Parallel Architectures

- INTEL Nehalem Core i7



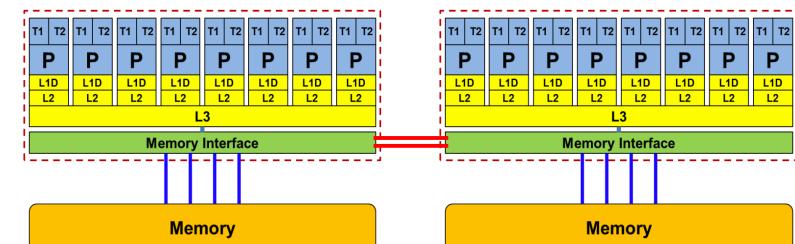
source: www.samrathacks.com

QPI: QuickPath Interconnect replaces FSB (QPI is a point-to-point interconnection – with a memory controller now on-die – in order to allow both reduced latency and higher bandwidth → up to (theoretically) 25.6 GByte/s data transfer, i.e. $2 \times$ FSB)

Foundations / Parallel Architectures

- Intel E5-2600 Sandy-Bridge Series

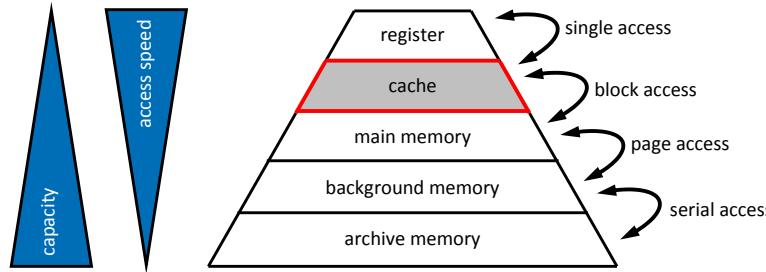
- 2 CPUs connected by **2 QPIs** (Intel Quick Path Interconnect)
- Quick Path Interconnect (1 sending and 1 receiving port)
 - $8 \text{ GT/s} \cdot 16 \text{ Bit/T payload} \cdot 2 \text{ directions} / 8 \text{ Bit/Byte} = 32 \text{ GB/s}$ max bandwidth per QPI
 - 2 QPI links** → $2 \cdot 32 \text{ GB/s} = 64 \text{ GB/s}$ max bandwidth



source: G. Wellein, RRZE

Foundations / Parallel Architectures

- reminder: memory hierarchy
 - memory hierarchy
 - exploitation of program characteristics such as locality
 - compromise between costs and performance
 - components with different speeds and capacities

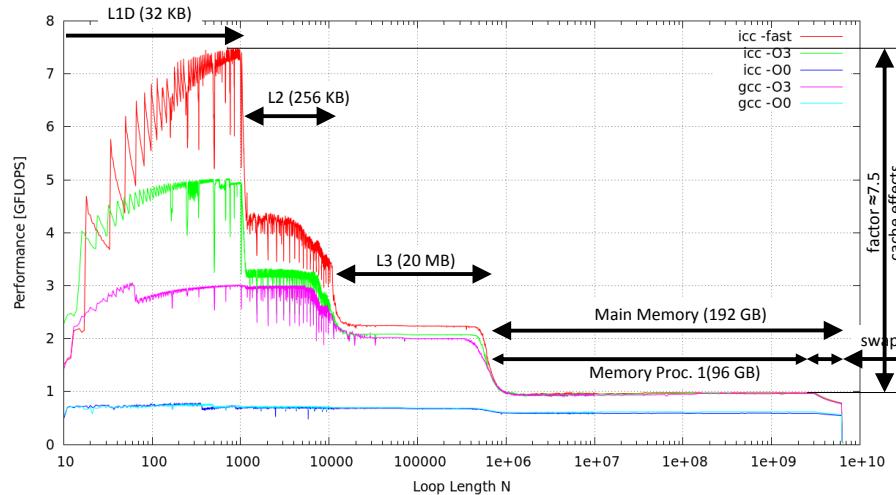


Foundations / Parallel Architectures

- reminder: memory hierarchy
 - example: SCHOENAUER vector triad benchmark
 - main kernel
- ```
double *A, *B, *C, *D
for i ← 0 to N-1 do
 A[i] ← B[i] + C[i] * D[i]
od
```
- report performance for different  $N$
  - kernel is limited by data transfer performance for all memory levels
  - using different compilers on Sandy-Bridge architecture
    - Intel Compiler 13.0.0 (icc)
    - GNU Compiler 4.6.3 (gcc)

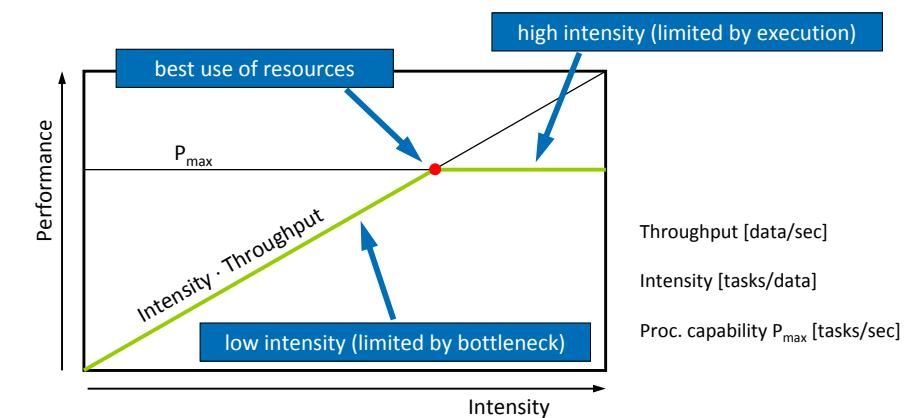
## Foundations / Parallel Architectures

- reminder: memory hierarchy



## Foundations / Parallel Architectures

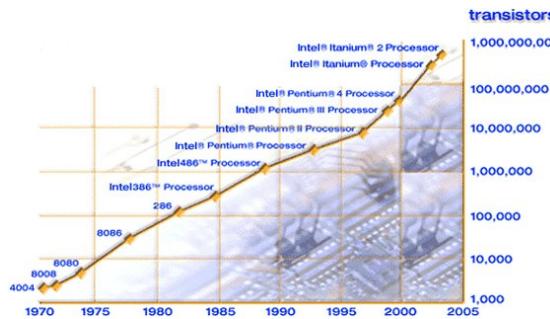
- roofline model
  - an optimistic performance model (for node level optimisation)



## Foundations / Parallel Architectures

- MOORE's law

- observation of Intel co-founder Gordon E. MOORE, describes important trend in history of computer hardware (1965)

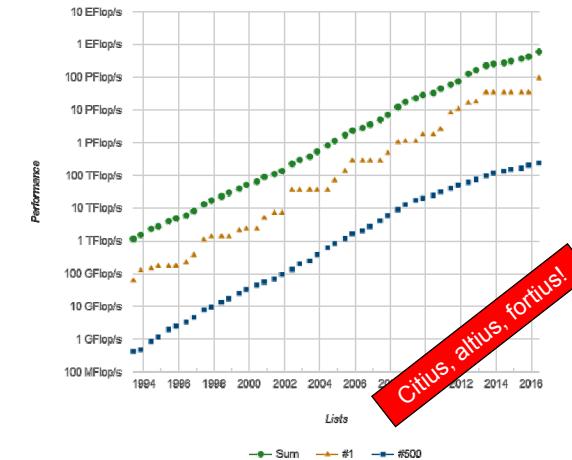


*"number of transistors that can be placed on an integrated circuit is increasing exponentially, doubling approximately every two years"*

## Foundations / Parallel Architectures

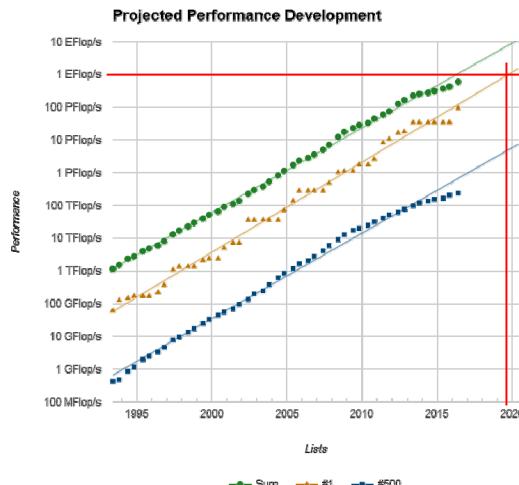
- some numbers: Top500 (as of June 2016)

Performance Development



## Foundations / Parallel Architectures

- some numbers: Top500 (as of June 2016)



## Foundations / Parallel Architectures

- the 10 fastest supercomputers in the world (as of June 2016)

| Rank | Site                                                            | System                                                                                                            | Cores      | Rmax<br>(TFlop/s) | Rpeak<br>(TFlop/s) | Power<br>(kW) |
|------|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|------------|-------------------|--------------------|---------------|
| 1    | National Supercomputing Center Wuhan China                      | Sunway TaihuLight - Sunway MPP, Sunway SW26010 250C, 1.45GHz, Sunway NRPC                                         | 10,649,600 | 93,014.6          | 125,435.9          | 15,371        |
| 2    | National Super Computer Center Guangzhou China                  | Tianhe-2 (MilkyWay-2) TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.20GHz, TH Express-2, Intel Xeon Phi 31S1P NUDI | 3,120,000  | 33,842.7          | 54,902.4           | 17,808        |
| 3    | DOD/SC/Oak Ridge National Laboratory United States              | Titan - Cray XK7, Opteron 6274 16C 2.20GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.                       | 560,640    | 17,590.0          | 27,112.5           | 8,209         |
| 4    | DOD/NNSA/LLNL United States                                     | Sequoia - BlueGene/Q, Power 800 16C 1.60 GHz, Custom IBM                                                          | 1,572,844  | 17,173.2          | 20,132.7           | 7,890         |
| 5    | RIKEN Advanced Institute for Computational Science (AICS) Japan | K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu                                                      | 705,024    | 10,510.0          | 11,280.4           | 12,660        |
| 6    | DOD/SC/Argonne National Laboratory United States                | Mira - BlueGene/Q, Power 800 16C 1.60GHz, Custom IBM                                                              | 786,432    | 8,598.6           | 10,066.3           | 3,945         |
| 7    | DOD/NSA/LANL/SNL United States                                  | Trinity - Cray XC40, Xeon E5-2693 12C 2.30GHz, Aries interconnect, Cray Inc.                                      | 301,056    | 8,100.9           | 11,078.9           |               |
| 8    | Swiss National Supercomputing Centre (CSCS) Switzerland         | Piz Daint - Cray XC30, Xeon E5-2670 8C 2.40GHz, Aries interconnect, NVIDIA K20x Cray Inc.                         | 115,784    | 6,271.0           | 7,788.9            | 2,325         |
| 9    | HLRS - Hochleistungszentrum Stuttgart Germany                   | Hazel Hen - Cray XC40, Xeon E5-2680 12C 2.50GHz, Aries interconnect Cray Inc.                                     | 185,088    | 5,640.2           | 7,403.5            |               |
| 10   | King Abdullah University of Science and Technology Saudi Arabia | Shahriar II - Cray XC40, Xeon E5-2690 8C 2.30GHz, Aries interconnect Cray Inc.                                    | 196,608    | 5,537.0           | 7,239.2            | 2,834         |

Rpeak ≡ theoretical peak performance

Rmax ≡ sustained peak performance

- overview

- geometric and physical modelling
- foundations / parallel architectures
- multigrid methods
- towards massive parallel HPC...
- interactive visual data exploration

## Multigrid Methods

- solvers for linear systems

- many PDEs result in a system of linear equations  $\mathbf{A} \cdot \mathbf{u} = \mathbf{f}$
- solution of such linear systems via
  - direct solvers
  - iterative solvers

- typical iterative solvers

- RICHARDSON method
- JACOBI method
- GAUSS-SEIDEL method
- relaxation methods
- CG and derivatives
- multigrid methods

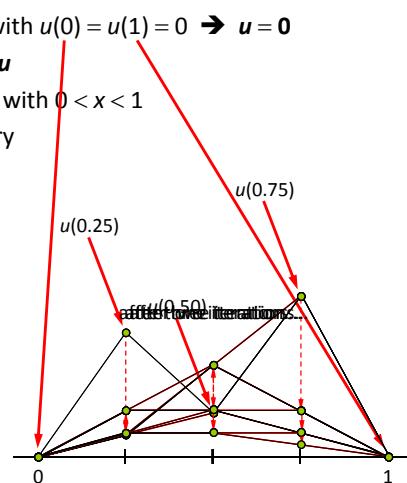
simple / moderate parallelisation effort

most effective and considered to be S.O.T.A. ☺

## Multigrid Methods

- something about smoother

- model BV problem:  $-u'' = 0$  with  $u(0) = u(1) = 0 \rightarrow \mathbf{u} = \mathbf{0}$
- from the above follows  $\mathbf{e} = -\mathbf{u}$
- arbitrary start values for  $u(x)$  with  $0 < x < 1$
- initial error  $\mathbf{e}$  highly oscillatory
- now applying a smoother...



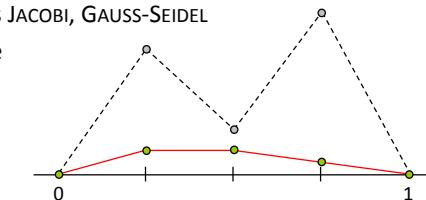
## Multigrid Methods

- something about smoother

- model BV problem:  $-u'' = 0$  with  $u(0) = u(1) = 0 \rightarrow \mathbf{u} = \mathbf{0}$
- from the above follows  $\mathbf{e} = -\mathbf{u}$
- arbitrary start values for  $u(x)$  with  $0 < x < 1$
- initial error  $\mathbf{e}$  highly oscillatory
- now applying a smoother...

- some observations

- high frequency parts of error are *smoothed out* by standard solvers such as JACOBI, GAUSS-SEIDEL
- on smooth functions above solvers become ineffective



## Multigrid Methods

- a more analytical approach
  - one smoothing step to be represented as

$$\hat{\mathbf{u}}_1 = \mathbf{R} \cdot \hat{\mathbf{u}}_0 + \mathbf{g}$$

with  $\mathbf{R}$  denoting the iteration matrix of the smoother; furthermore, the exact solution  $\hat{\mathbf{u}}$  is a fixed-pointed of the iteration, that means

$$\hat{\mathbf{u}} = \mathbf{R} \cdot \hat{\mathbf{u}} + \mathbf{g}$$

- with  $\mathbf{e} = \hat{\mathbf{u}} - \mathbf{u}$  subtracting the last two expressions yields

$$\mathbf{e}_1 = \mathbf{R} \cdot \mathbf{e}_0$$

- repeating this, after  $m$  smoothing steps the error is given by

$$\mathbf{e}_m = \mathbf{R}^m \cdot \mathbf{e}_0$$

- with  $\rho(\mathbf{R}) < 1$ , the error is forced to zero as the iteration proceeds

## Multigrid Methods

- a more analytical approach
  - let  $\mathbf{w}_k$  denoted the  $k$ -th eigenvector of  $\mathbf{R}$ , then it is possible to expand  $\mathbf{e}_0$  as

$$\mathbf{e}_0 = \sum_{k=1}^{n-1} c_k \cdot \mathbf{w}_k$$

with coefficients  $c_k \in \mathbb{R}$  denoting weighting factors for each  $\mathbf{w}_k$  in the error

- using

$$\mathbf{e}_m = \mathbf{R}^m \cdot \mathbf{e}_0$$

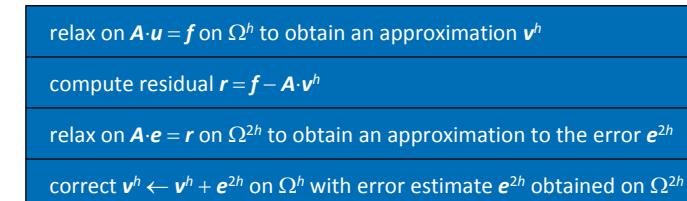
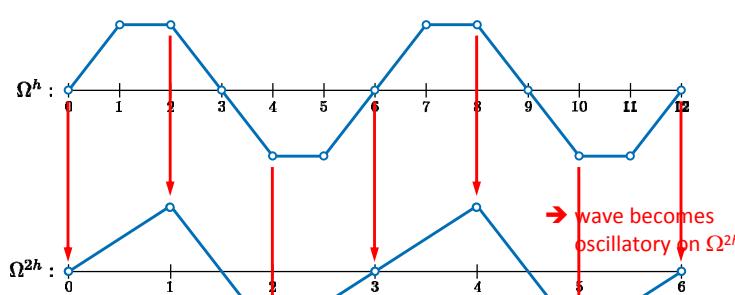
and the eigenvector expansion for  $\mathbf{e}_0$ , we get

$$\mathbf{e}_m = \mathbf{R}^m \cdot \mathbf{e}_0 = \sum_{k=1}^{n-1} c_k \cdot \mathbf{R}^m \cdot \mathbf{w}_k \stackrel{\mathbf{R} \cdot \mathbf{w}_k = \lambda_k(\mathbf{R}) \cdot \mathbf{w}_k}{=} \sum_{k=1}^{n-1} c_k \cdot \lambda_k(\mathbf{R})^m \cdot \mathbf{w}_k$$

- from above expansion we see that small eigenvalues ( $\approx 0$ ) corresponding to high frequency parts of the error diminish faster than large eigenvalues ( $\approx 1$ ) corresponding to low frequency parts of the error

## Multigrid Methods

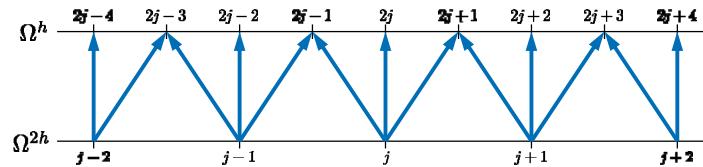
- towards multigrid
  - how do smooth components look like on coarser grids?
  - consider some fine ( $\Omega^h$ ) and coarse ( $\Omega^{2h}$ ) grid with double grid spacing
  - given some smooth wave on  $\Omega^h$  with  $n = 13$  points
  - $\Omega^{2h}$  representation with  $n = 7$  points via direct projection



- question: how to transfer residual  $\mathbf{r}^h$  from  $\Omega^h$  to  $\Omega^{2h}$  (called restriction) and how to transfer the error estimate  $\mathbf{e}^{2h}$  back from  $\Omega^{2h}$  to  $\Omega^h$  (called interpolation or prolongation)?

## Multigrid Methods

- towards multigrid
  - prolongation operator  $\mathbf{I}_{2h}^h$   
→ produces fine-grid vectors from coarse ones according to  $\mathbf{I}_{2h}^h \mathbf{v}^{2h} = \mathbf{v}^h$
  - simplest approach: linear prolongation

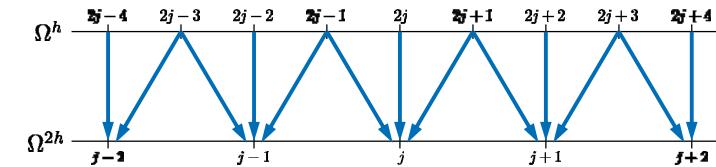


with

$$\begin{aligned} v_{2j}^h &= v_j^{2h} \\ v_{2j+1}^h &= \frac{1}{2} (v_j^{2h} + v_{j+1}^{2h}) \quad , 0 \leq j \leq \frac{n}{2} - 1 \end{aligned}$$

## Multigrid Methods

- towards multigrid
  - restriction operator  $\mathbf{I}_h^{2h}$   
→ produces coarse-grid vectors from fine ones according to  $\mathbf{I}_h^{2h} \mathbf{v}^h = \mathbf{v}^{2h}$
  - typical approach: full weighting



with

$$v_j^{2h} = \frac{1}{4} (v_{2j-1}^h + 2v_{2j}^h + v_{2j+1}^h) \quad , 0 \leq j \leq \frac{n}{2} - 1$$

## Multigrid Methods

- two-grid correction scheme
  - now using well-defined ways to transfer vectors between grids
  - parameters  $\nu_1, \nu_2$  control number of relaxation steps and are in practice often 1, 2, or 3

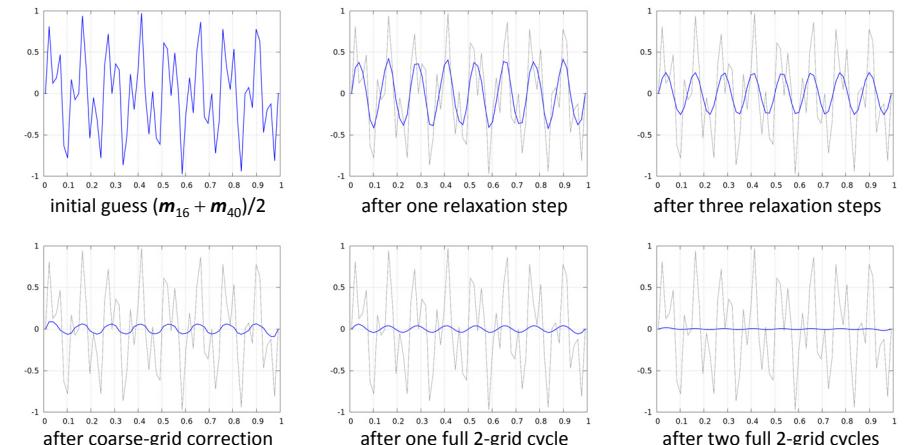
```

relax ν_1 times on $\mathbf{A}^h \cdot \mathbf{v}^h = \mathbf{f}^h$ on Ω^h with initial guess \mathbf{v}^h
compute residual $\mathbf{r}^h = \mathbf{f}^h - \mathbf{A}^h \cdot \mathbf{v}^h$
restrict residual \mathbf{r}^h to coarse grid by $\mathbf{r}^{2h} = \mathbf{I}_h^{2h} \mathbf{r}^h$
solve $\mathbf{A}^{2h} \cdot \mathbf{e}^{2h} = \mathbf{r}^{2h}$ on Ω^{2h}
prolongate coarse-grid error \mathbf{e}^{2h} to fine grid by $\mathbf{e}^h = \mathbf{I}_{2h}^h \mathbf{e}^{2h}$
correct fine-grid approximation $\mathbf{v}^h \leftarrow \mathbf{v}^h + \mathbf{e}^h$
relax ν_2 times on $\mathbf{A}^h \cdot \mathbf{v}^h = \mathbf{f}^h$ on Ω^h with corrected approximation \mathbf{v}^h

```

## Multigrid Methods

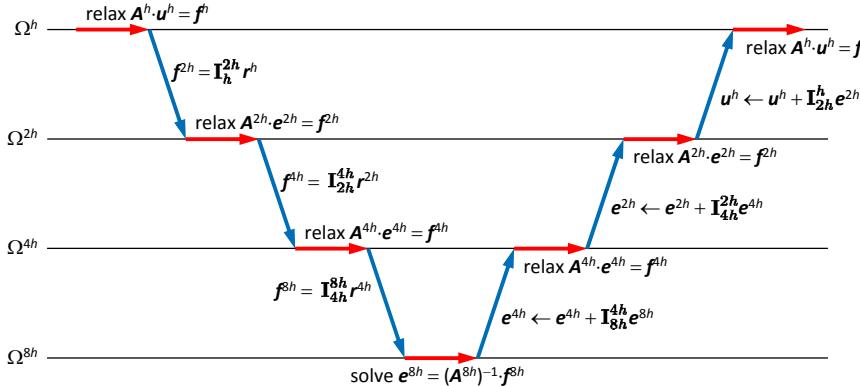
- two-grid correction scheme
  - example ( $\mathbf{u} = \mathbf{0}$ ) with overlay of FOURIER modes  $\mathbf{m}_{16}$  and  $\mathbf{m}_{40}$  as initial guess



## Multigrid Methods

- **V-cycle scheme**

- why restricting approach to two grids only?
- idea: recursive algorithm

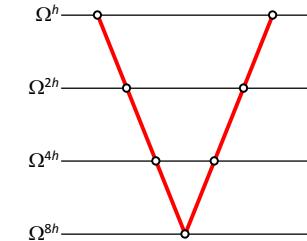


## Multigrid Methods

- **V-cycle scheme**

$$\mathbf{v}^k \leftarrow \text{MG}_V(\mathbf{v}^k, \mathbf{f}^k)$$

1. relax  $\nu_1$  times on  $\mathbf{A}^k \cdot \mathbf{v}^k = \mathbf{f}^k$  with initial guess  $\mathbf{v}^k$
2. if  $\Omega^k$  = coarsest grid, then go to step 4  
else
- $\mathbf{f}^{2k} \leftarrow \mathbf{I}_k^{2k}(\mathbf{f}^k - \mathbf{A}^k \cdot \mathbf{v}^k)$
- $\mathbf{v}^{2k} \leftarrow \mathbf{0}$
- $\mathbf{v}^{2k} \leftarrow \text{MG}_V(\mathbf{v}^{2k}, \mathbf{f}^{2k})$
3. correct  $\mathbf{v}^k \leftarrow \mathbf{v}^k + \mathbf{I}_{2k}^k \mathbf{v}^{2k}$
4. relax  $\nu_2$  times on  $\mathbf{A}^k \cdot \mathbf{v}^k = \mathbf{f}^k$



## Multigrid Methods

- **full multigrid V-cycle**

$$\mathbf{v}^k \leftarrow \text{FMG}(\mathbf{f}^k)$$

1. if  $\Omega^k$  = coarsest grid, set  $\mathbf{v}^k \leftarrow \mathbf{0}$  and go to step 3

else

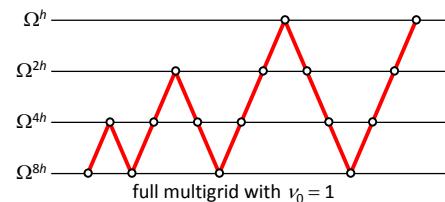
$$\mathbf{f}^{2k} \leftarrow \mathbf{I}_k^{2k}(\mathbf{f}^k)$$

$$\mathbf{v}^{2k} \leftarrow \text{FMG}(\mathbf{f}^{2k})$$

2. correct  $\mathbf{v}^k \leftarrow \mathbf{I}_{2k}^k \mathbf{v}^{2k}$

3.  $\mathbf{v}^k \leftarrow \text{MG}_V(\mathbf{v}^k, \mathbf{f}^k)$   $\nu_0$  times

Here, the idea is to use coarse grids in order to obtain better initial guesses, a strategy called nested iteration.



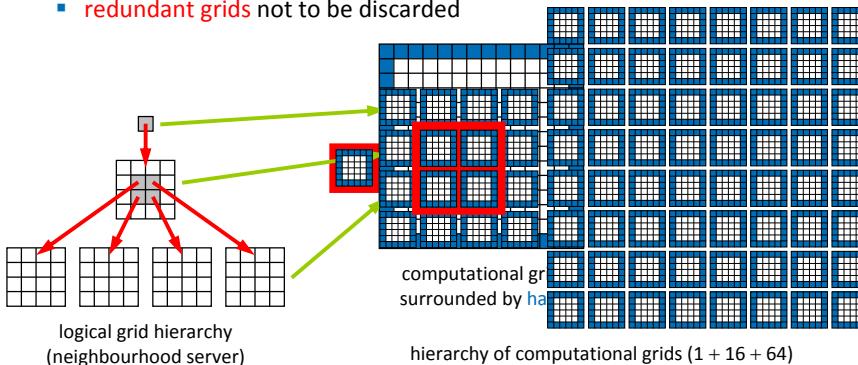
- **overview**

- geometric and physical modelling
- foundations / parallel architectures
- multigrid methods
- towards massive parallel HPC...
- interactive visual data exploration

## Towards Massive Parallel HPC...

- data structure / grid layout

- nested non-overlapping block-structured orthogonal grids
- management (i.e. neighbourhood server) hidden from application
- each logical cell links to a computational grid surrounded by halo
- redundant grids not to be discarded



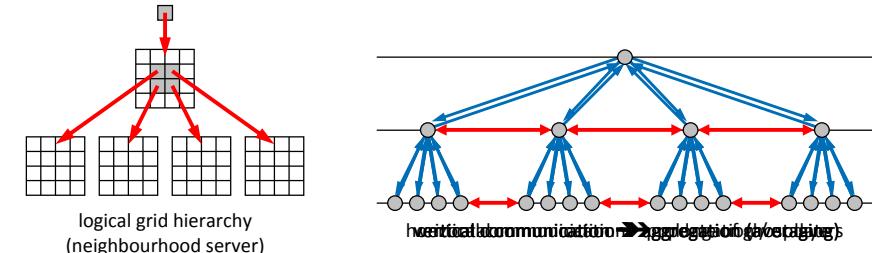
## Towards Massive Parallel HPC...

- data structure / grid layout

- nested non-overlapping block-structured orthogonal grids
- management (i.e. neighbourhood server) hidden from application
- each logical cell links to a computational grid surrounded by halo

- data flow

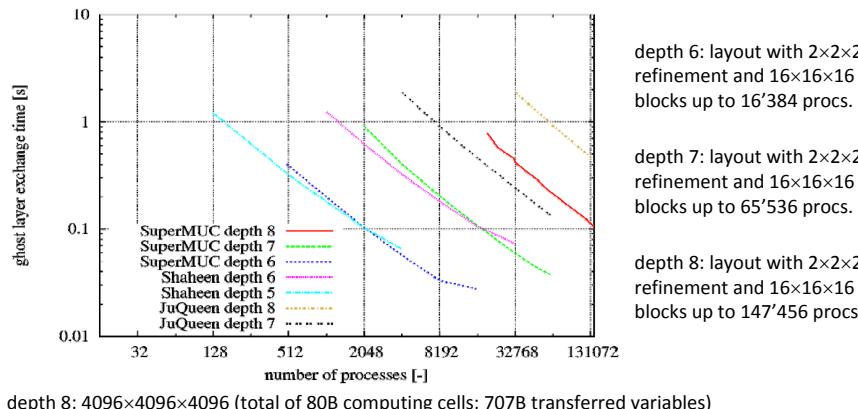
- vertical communication (aggregation / prolongation of values)
- horizontal communication (update of ghost layers)



## Towards Massive Parallel HPC...

- data flow between grids

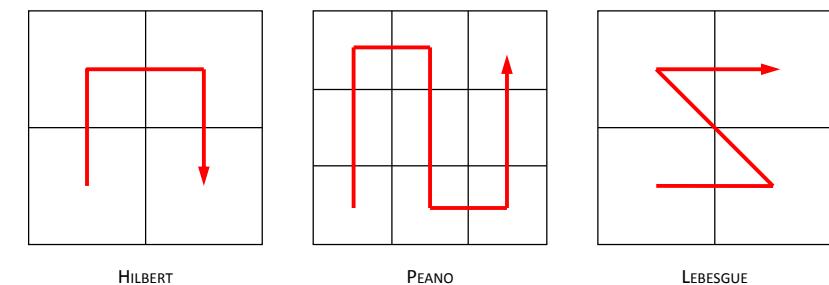
- time for one full processing, i.e. bottom-up + horizontal + top-down communication between all grids (**no computation done**)



## Towards Massive Parallel HPC...

- space-filling curves (SFC)

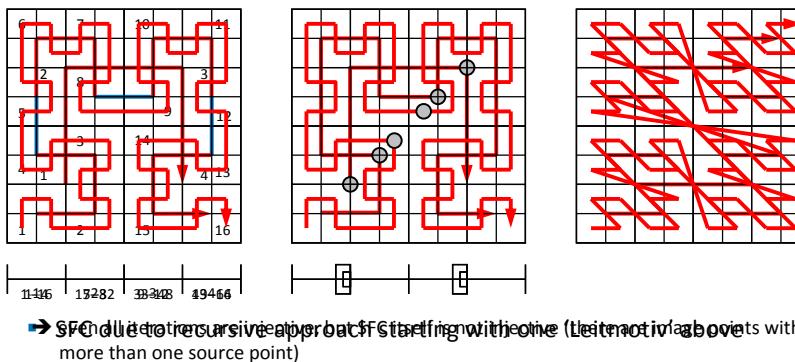
- continuous, surjective mapping  $f: [0, 1] \rightarrow [0, 1]^D$
- advantage: preserving neighbourhood relations
- typical representatives (generator or 'Leitmotiv')



- SFC due to recursive approach starting with one 'Leitmotiv' above

## Towards Massive Parallel HPC...

- space-filling curves (SFC)
  - continuous, surjective mapping  $f: [0, 1] \rightarrow [0, 1]^D$
  - advantage: preserving neighbourhood relations
  - typical representatives (generator or 'Leitmotiv')



## Towards Massive Parallel HPC...

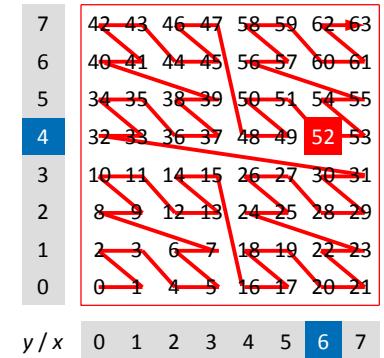
- space-filling curves (SFC)
  - for load distribution inverse function  $f^{-1}: [0, 1]^D \rightarrow [0, 1]$  necessary
  - simple conversion of Z-index in case of LEBESGUE's SFC possible
  - idea: bitwise interleaving of coordinate values

$$x = 6 \rightarrow 110$$

$$y = 4 \rightarrow 100$$

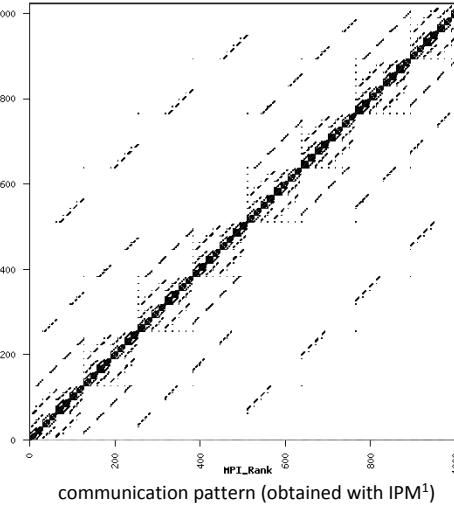
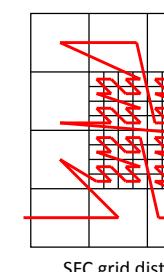
$$110100 \rightarrow 52 = Z$$

→ simple conversion  $(6, 4) \leftrightarrow 52_Z$



## Towards Massive Parallel HPC...

- grid distribution
  - space-filling curves (SFC)
  - neighbourhood relations
  - simple geometric partitioning



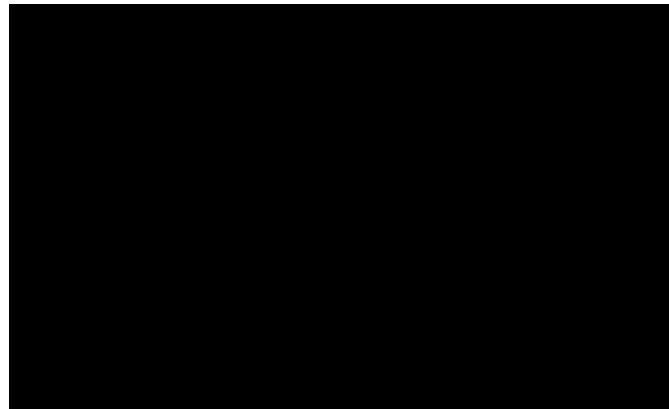
processes

n) answers queries  
Is → grids don't need  
more / processes

<sup>1</sup> Integrated Performance Monitoring, <http://ipm-hpc.sourceforge.net/>

## Towards Massive Parallel HPC...

- grid distribution / load balancing
  - example: temperature distribution – grid migration



## Towards Massive Parallel HPC...

- computational kernel
  - NS equations, FV for spatial, Adams-Bashforth (2<sup>nd</sup> order FD) for temporal discretisation
  - fractional step (Chorin's projection) for solving time-dependent incompressible flow equations, i.e. iterative procedure between velocity and pressure during one time step
  - thermal coupling realised by Boussinesq approximation (**modified body term in NSE momentum equation**)

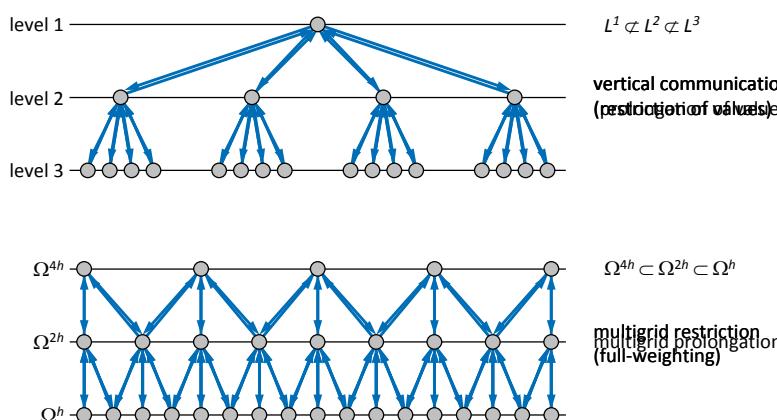
$$\nabla \cdot \vec{u} = 0$$

$$\frac{\partial \rho_\infty u_i}{\partial t} + \nabla \cdot (\rho_\infty u_i \vec{u}) = \nabla \cdot (\mu \nabla u_i) - \nabla \cdot (p \vec{e}_i) - \rho_\infty \cdot \beta \cdot (T - T_\infty) g_i, \text{ with } i \in \{x, y, z\}$$

$$\frac{\partial T}{\partial t} + \nabla \cdot (T \vec{u}) - \nabla \cdot (\alpha \nabla T) - \frac{q_{int}}{\rho_\infty \cdot c_p} = 0$$

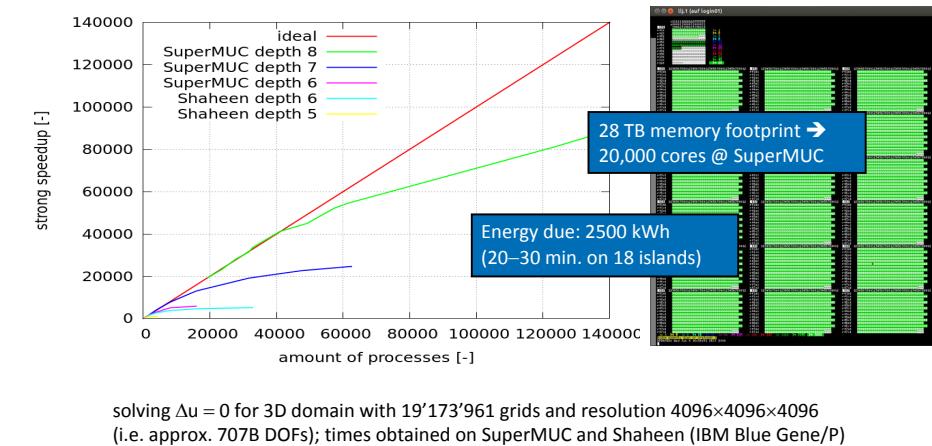
## Towards Massive Parallel HPC...

- parallel multigrid(-like) solver
  - comparison: vertical communication vs. multigrid transfer functions



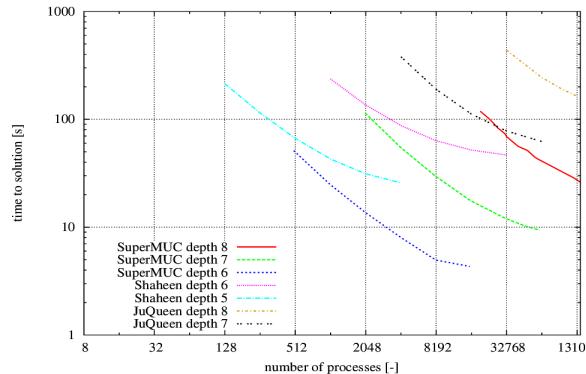
## Towards Massive Parallel HPC...

- parallel multigrid(-like) solver



## Towards Massive Parallel HPC...

- parallel multigrid(-like) solver
  - time to solution for one time step (repeated V-cyles with adaptive relaxation steps (and secret scaling factor ☺) until convergence)



depth 6: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 16'384 procs.

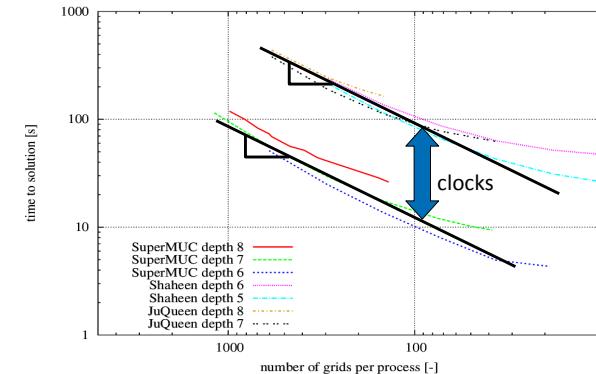
depth 7: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 65'536 procs.

depth 8: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 147'456 procs.

depth 8:  $4096 \times 4096 \times 4096$  (total of 80B computing cells; 707B degrees of freedom)

## Towards Massive Parallel HPC...

- parallel multigrid(-like) solver
  - time to solution for one time step (repeated V-cyles with adaptive relaxation steps (and secret scaling factor ☺) until convergence)

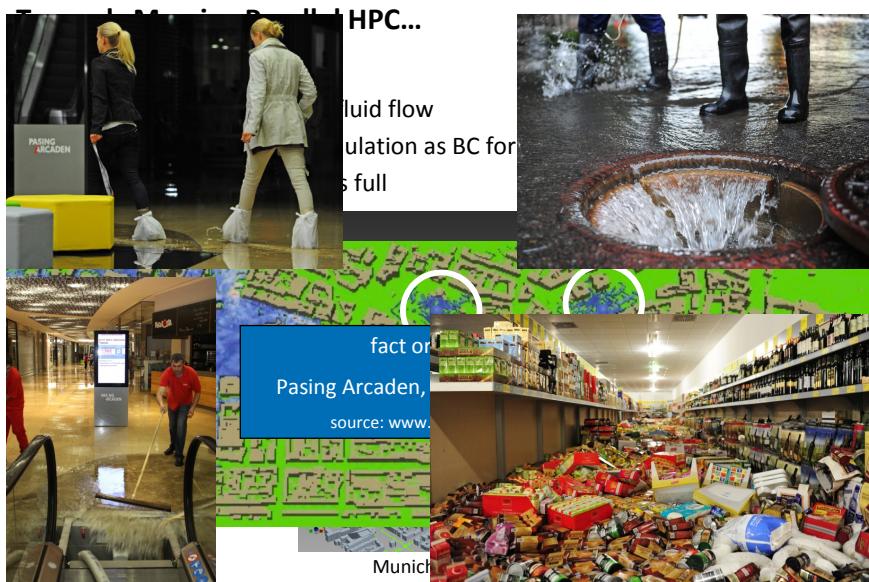


depth 6: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 16'384 procs.

depth 7: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 65'536 procs.

depth 8: layout with  $2 \times 2 \times 2$  refinement and  $16 \times 16 \times 16$  blocks up to 147'456 procs.

depth 8:  $4096 \times 4096 \times 4096$  (total of 80B computing cells; 707B degrees of freedom)



## Towards Massive Parallel HPC...

- multiscale flood simulation



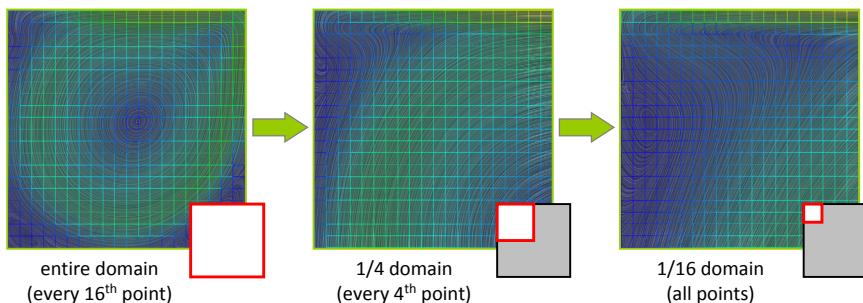
- overview

- geometric and physical modelling
- foundations / parallel architectures
- multigrid methods
- towards massive parallel HPC...
- interactive visual data exploration

## Interactive Visual Data Exploration

- sliding window concept

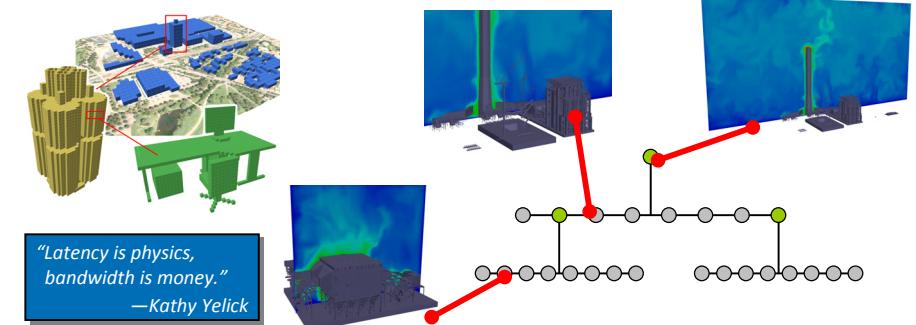
- problem: high resolutions hinder interactive exploration
- solution: user moves / sizes ‘window’ through domain for data exploration  
→ amount of details increases seamlessly
- constant bandwidth of data transmission → simple postprocessing



## Interactive Visual Data Exploration

- sliding window

- idea: online navigation through details

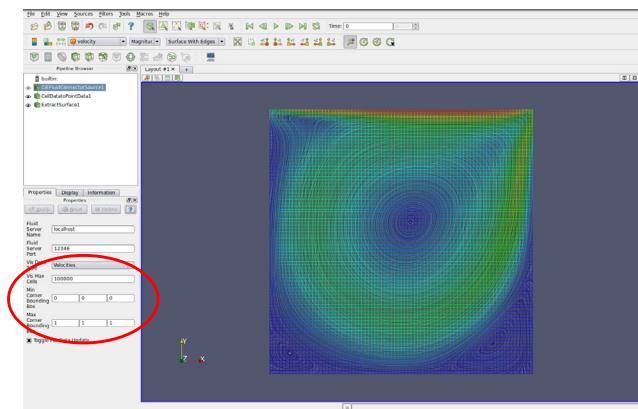


*“High-performance computing must now assume a broader meaning, encompassing not only flops, but also the ability, for example, to efficiently manipulate vast and rapidly increasing quantities of both numerical and non-numerical data.”*

## Interactive Visual Data Exploration

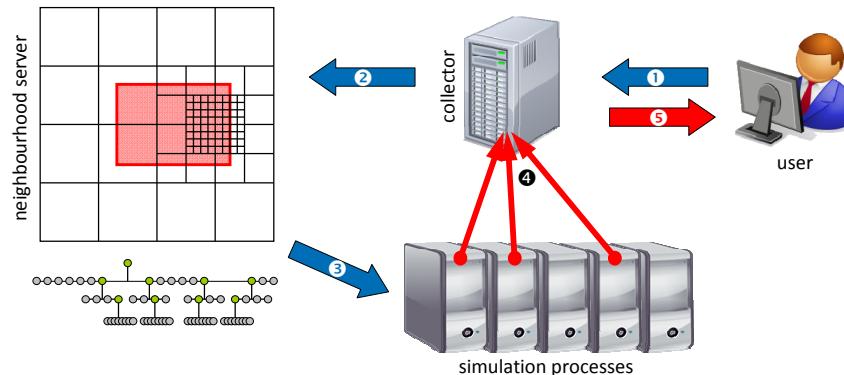
- sliding window concept

- simple part: what happens on the front-end...



## Interactive Visual Data Exploration

- sliding window concept
  - complex part: what happens on the back-end...
  - collector node handles queries and ‘fills’ data stream top-down

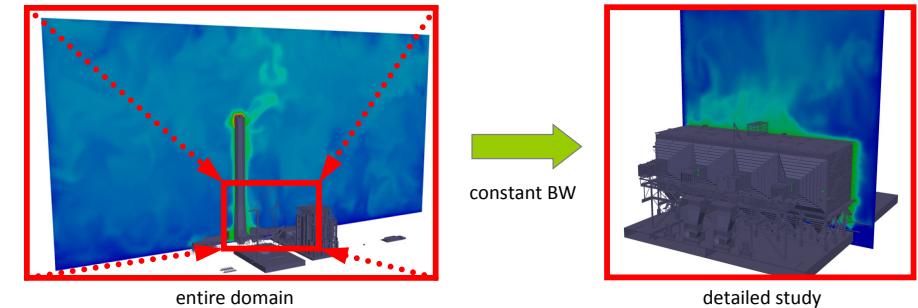


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

73

## Interactive Visual Data Exploration

- sliding window concept
  - geometric model: power plant (BREP with 12,748,510 faces)
  - user selects window for details interactively during runtime

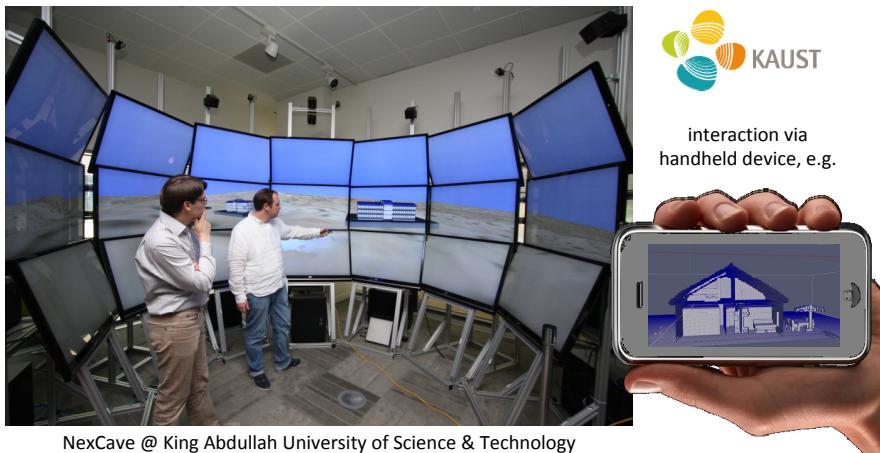


PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

74

## Interactive Visual Data Exploration

- interactive 3D data exploration: size does matter!



NexCave @ King Abdullah University of Science &amp; Technology

PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

75



contact: mundani@tum.de

- acknowledgements



PD Dr. Ralf-Peter Mundani – HPC for Environmental Simulations – SYNASC, Timisoara, Romania, 09/27/2016

76