INGRAIN
DIGITAL ROCK PHYSICS LAB

Lattice Boltzmann multi-phase simulations using GPUs

Jonas Tölke

Ingrain

September, 22, 2010

GTC 2010

San Jose, CA

Outline

INGRAIN
DIGITAL ROCK PHYSICS LAB

- Ingrain and Digital Rock Physics
- Multiphase flow in porous media: different scales
- Lattice Boltzmann methods (LBM) on the pore scale
- GPU Programming for LBM
- Examples
- Summary and Conclusions

Stages of oil recovery

INGRAIN
DIGITAL ROCK PHYSICS LAB

- Primary oil recovery (15%): oil and gas are drawn out automatically by the underground natural mechanical energy.
- Secondary oil recovery (15–20%): water is pumped in the underground to maintain the energy level.
- Tertiary oil recovery (EOR) (10-30%):
  - Gas injection
  - Chemical injection
  - Ultrasonic stimulation
  - Microbial injection
  - Thermal recovery

Ingrain

INGRAIN
DIGITAL ROCK PHYSICS LAB

Leading oil and gas companies worldwide use Ingrain's digital rock physics labs

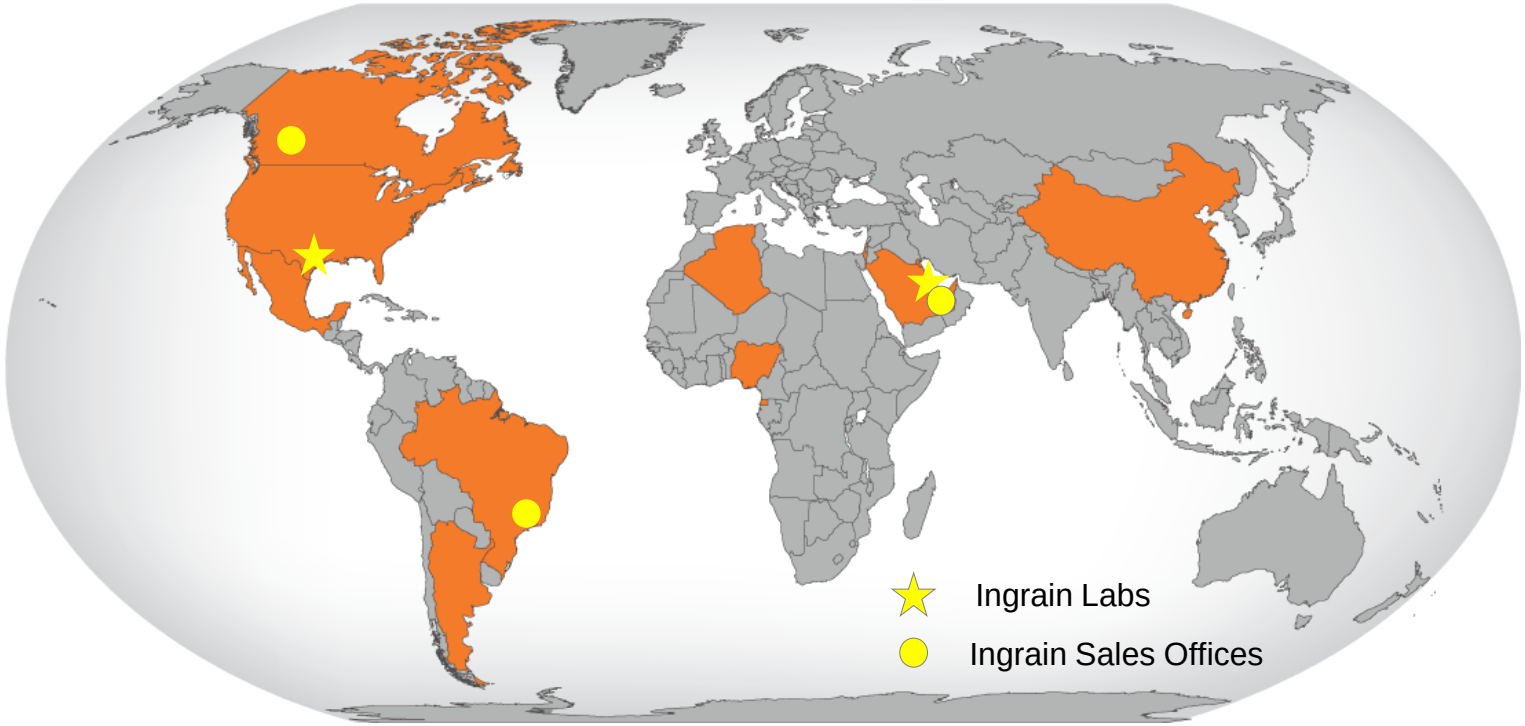

- 25 US and international independents
- National oil companies in 6 countries
- 4 of 5 supermajors

Over 4000 rock samples processed and 125 commercial jobs completed in the past 18 months

Descending Scale systems

INGRAIN
DIGITAL ROCK PHYSICS LAB

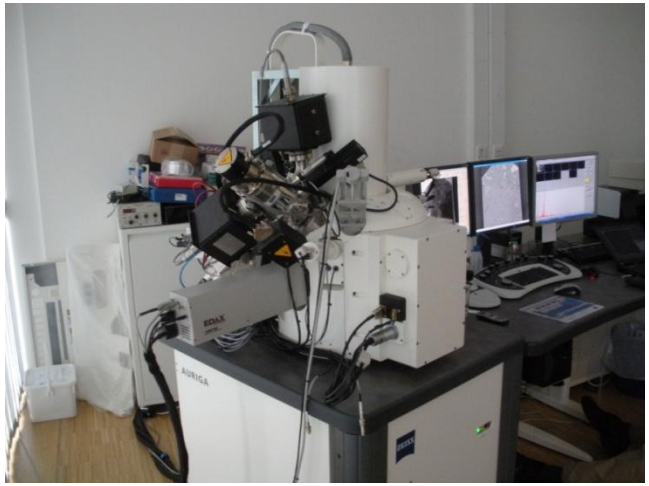

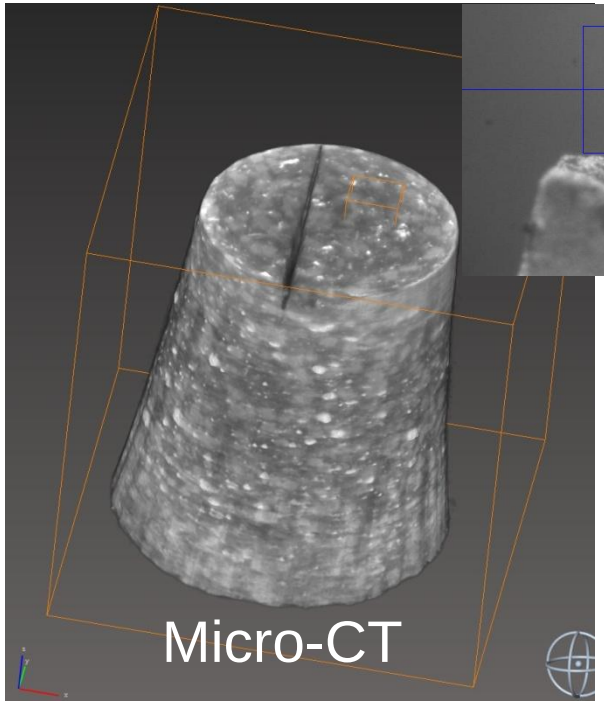

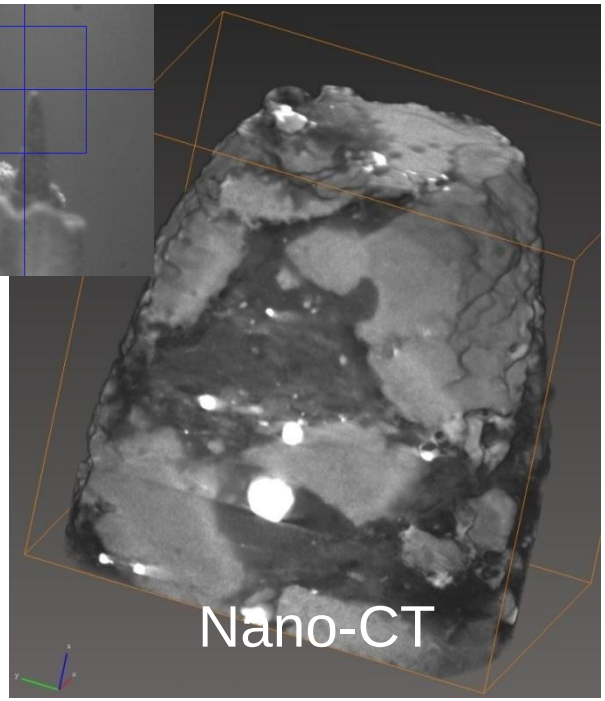

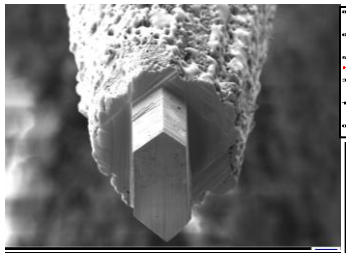

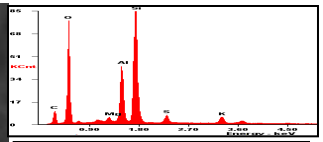

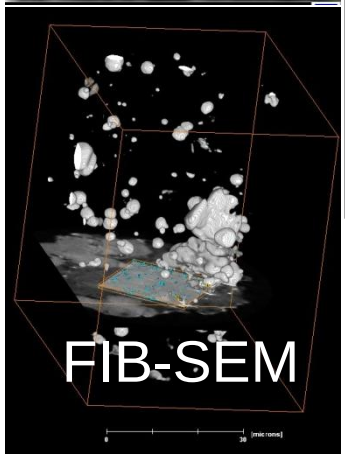

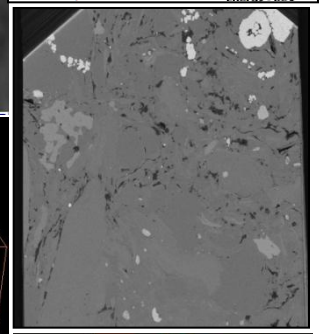

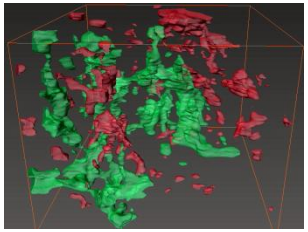

Ingrain images reservoir rock samples at the resolution needed to compute basic and advanced rock properties

INGRAIN
DIGITAL ROCK PHYSICS LAB

Macro 625
625 microns/voxel

reservoir

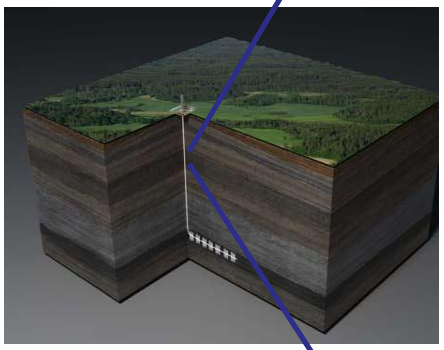

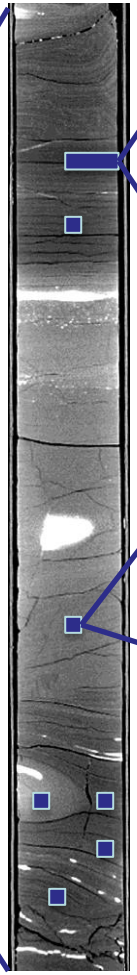

Whole core (1m)

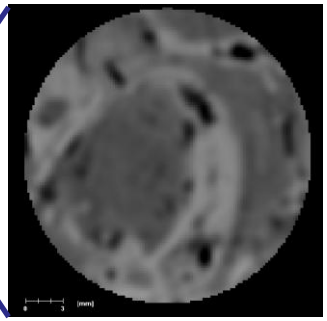

Macro 40
40 microns/voxel

Core plug (5cm)

When sampling and imaging, Ingrain uses the largest volume and lowest resolution (and therefore largest field of view) that will permit accurate computation of rock properties and fluid flow.

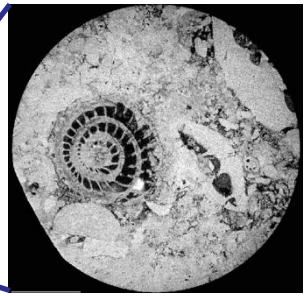

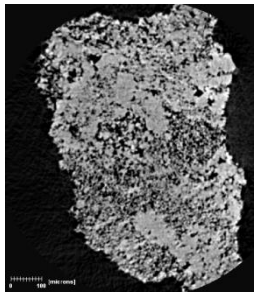

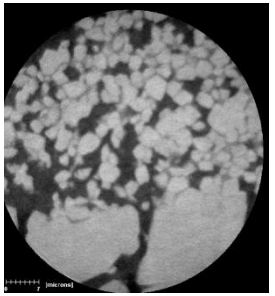

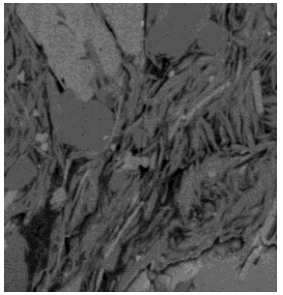

Micro 1/2/4
1-4 microns/voxel

Subsample size 1-4mm

Nano 500
500 nanometers/ voxel

Subsample size .5mm

Nano 65
65 nanometers/ voxel

Subsample size .1mm

Nano 3
3 nanometers/ voxel

Subsample size .1mm

vRock® Basic and Advanced Rock Properties

INGRAIN
DIGITAL ROCK PHYSICS LAB

- Physical properties quickly and accurately computed from vRock® digital reservoir rocks:
  - Porosity
  - Absolute permeability in X, Y and Z directions
- Advanced rock properties computations include:
  - Formation factor in X, Y and Z directions
  - Elastic properties
  - Grain statistics
  - Two-phase relative permeability
  - Capillary pressure
- Digital subsampling capability: trend of values that reflect the heterogeneity within each rock sample
- Ingrain permanently stores each vRock® in a secure digital rock physics database

vRock®

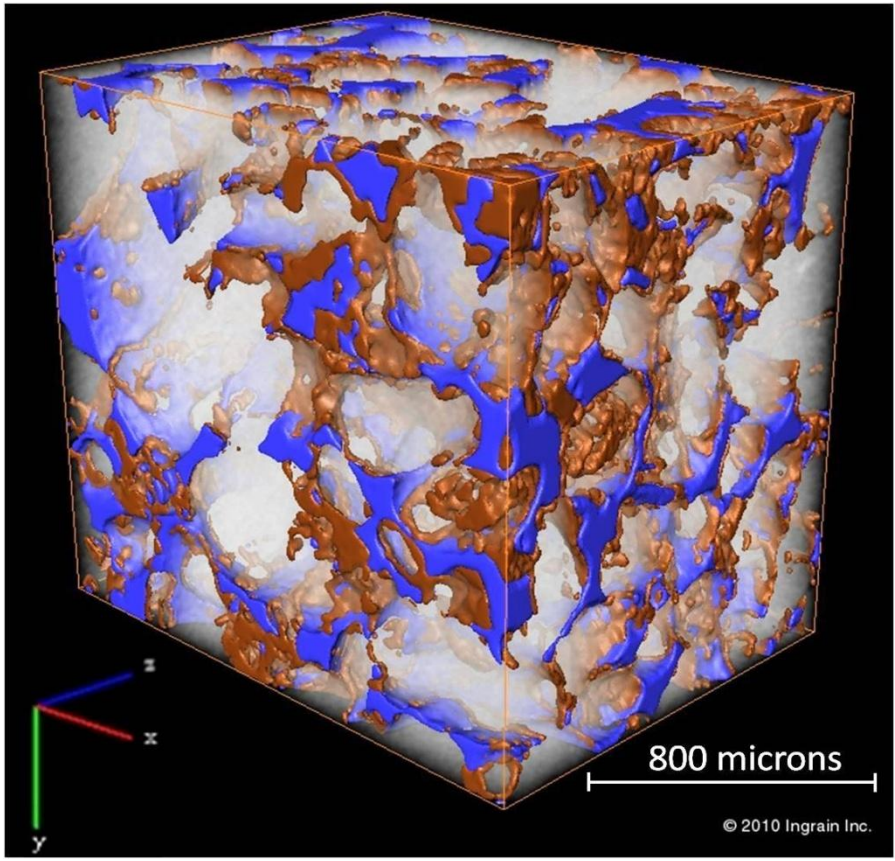

Work Flow to obtain rock properties on the pore scale

INGRAIN
DIGITAL ROCK PHYSICS LAB

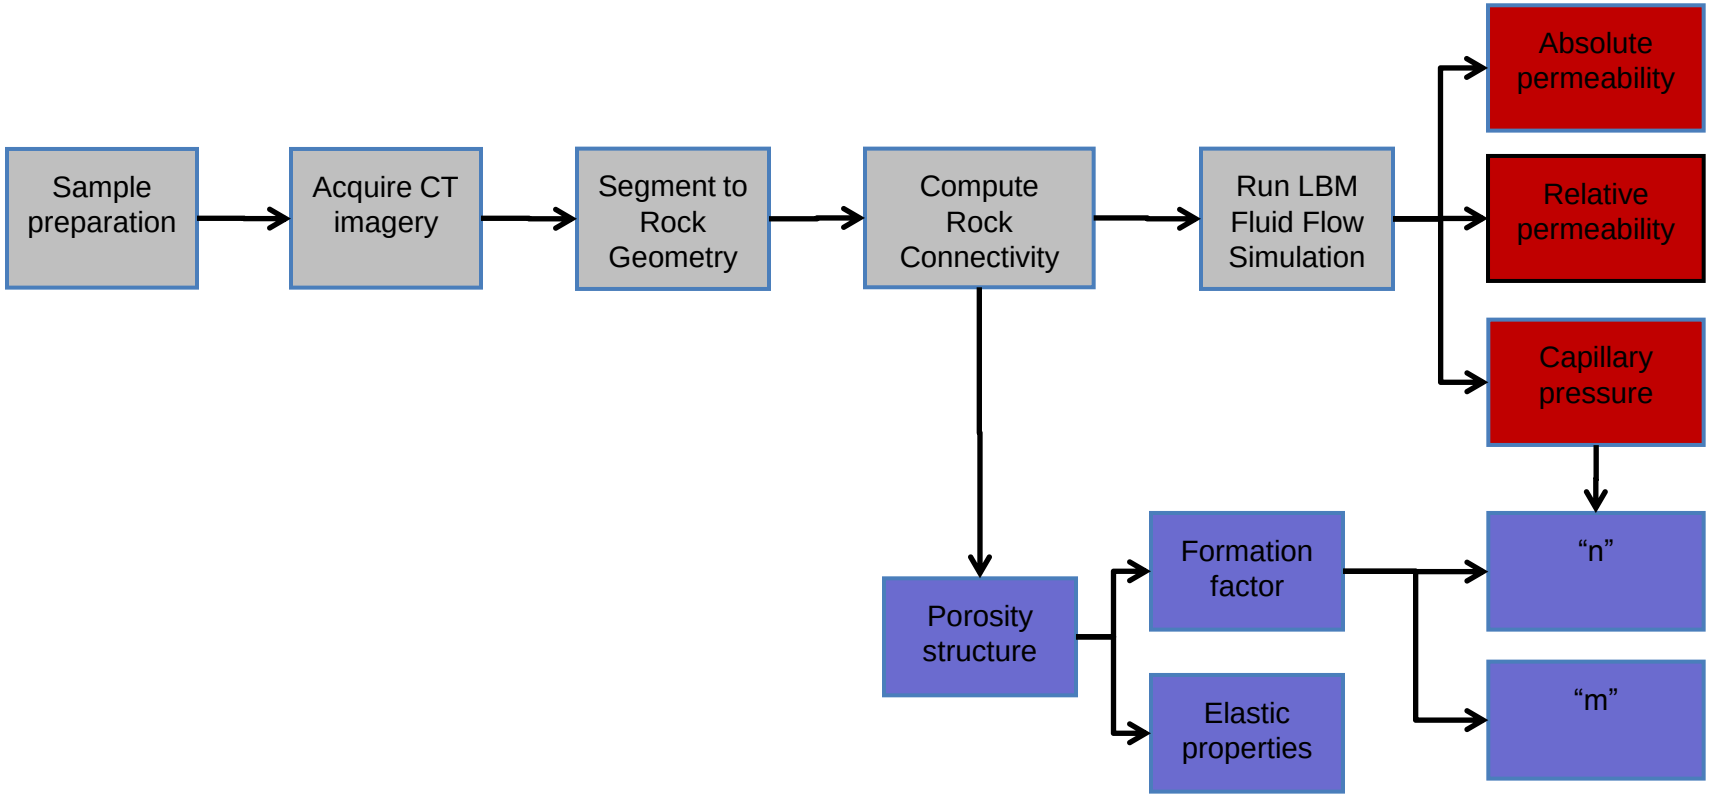

Imaging and Upscaling from Pore to Core

INGRAIN
DIGITAL ROCK PHYSICS LAB

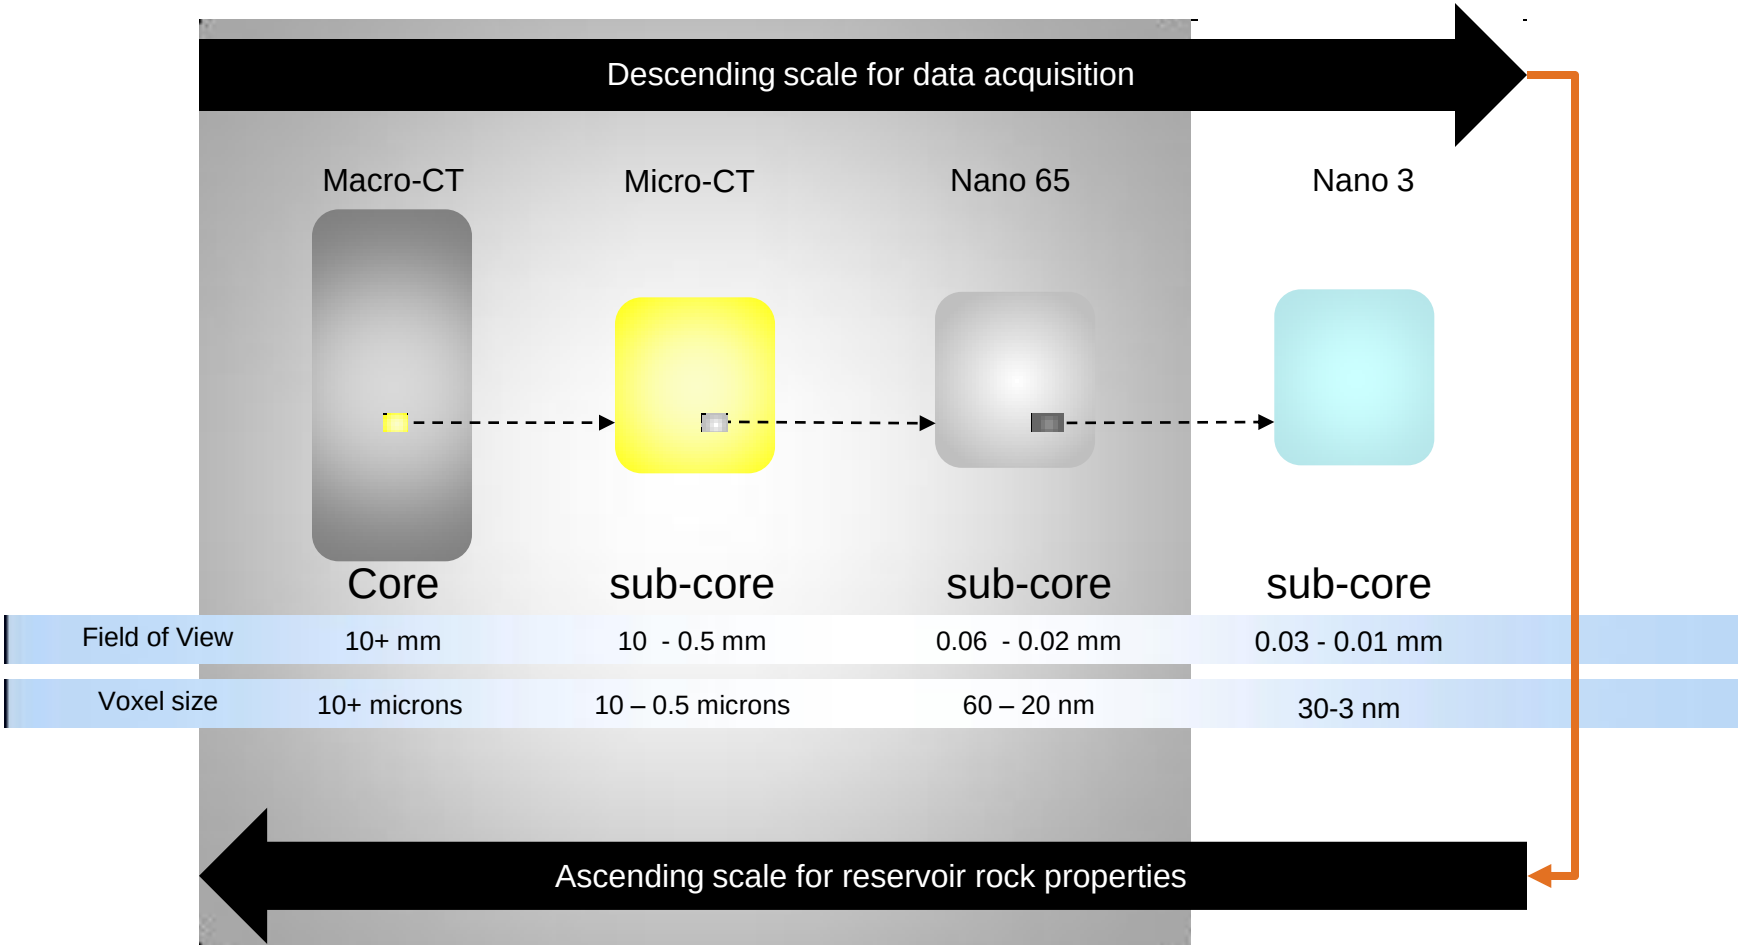

Flow in porous media– different scales

INGRAIN
DIGITAL ROCK PHYSICS LAB

Pore-scale:
- Multiphase (Navier-)Stokes

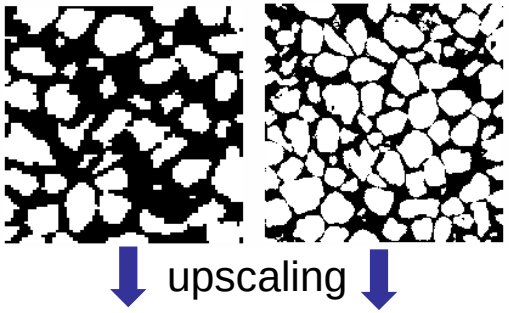

Darcy-Scala:
- Multiphase Darcy

upscaling

Homogenous Darcy-Scale

Heterogenous Darcy-Scale

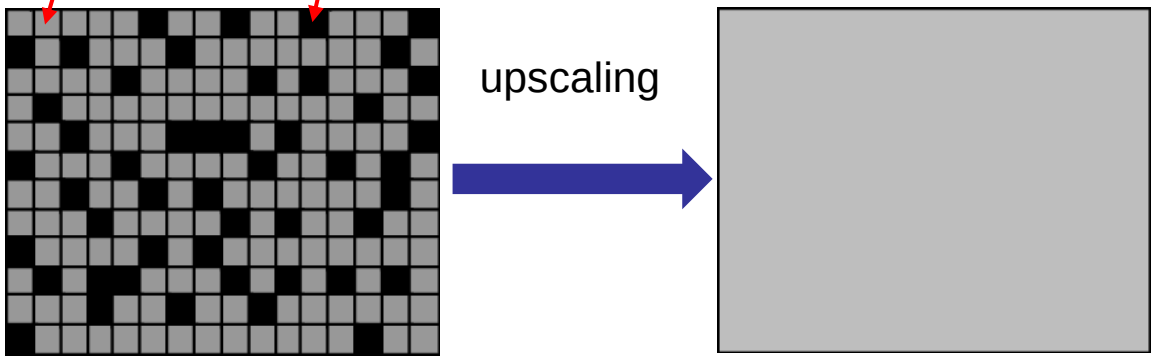

Ingrain's multi-scale "pore to core" approach to rock typing

INGRAIN
DIGITAL ROCK PHYSICS LAB

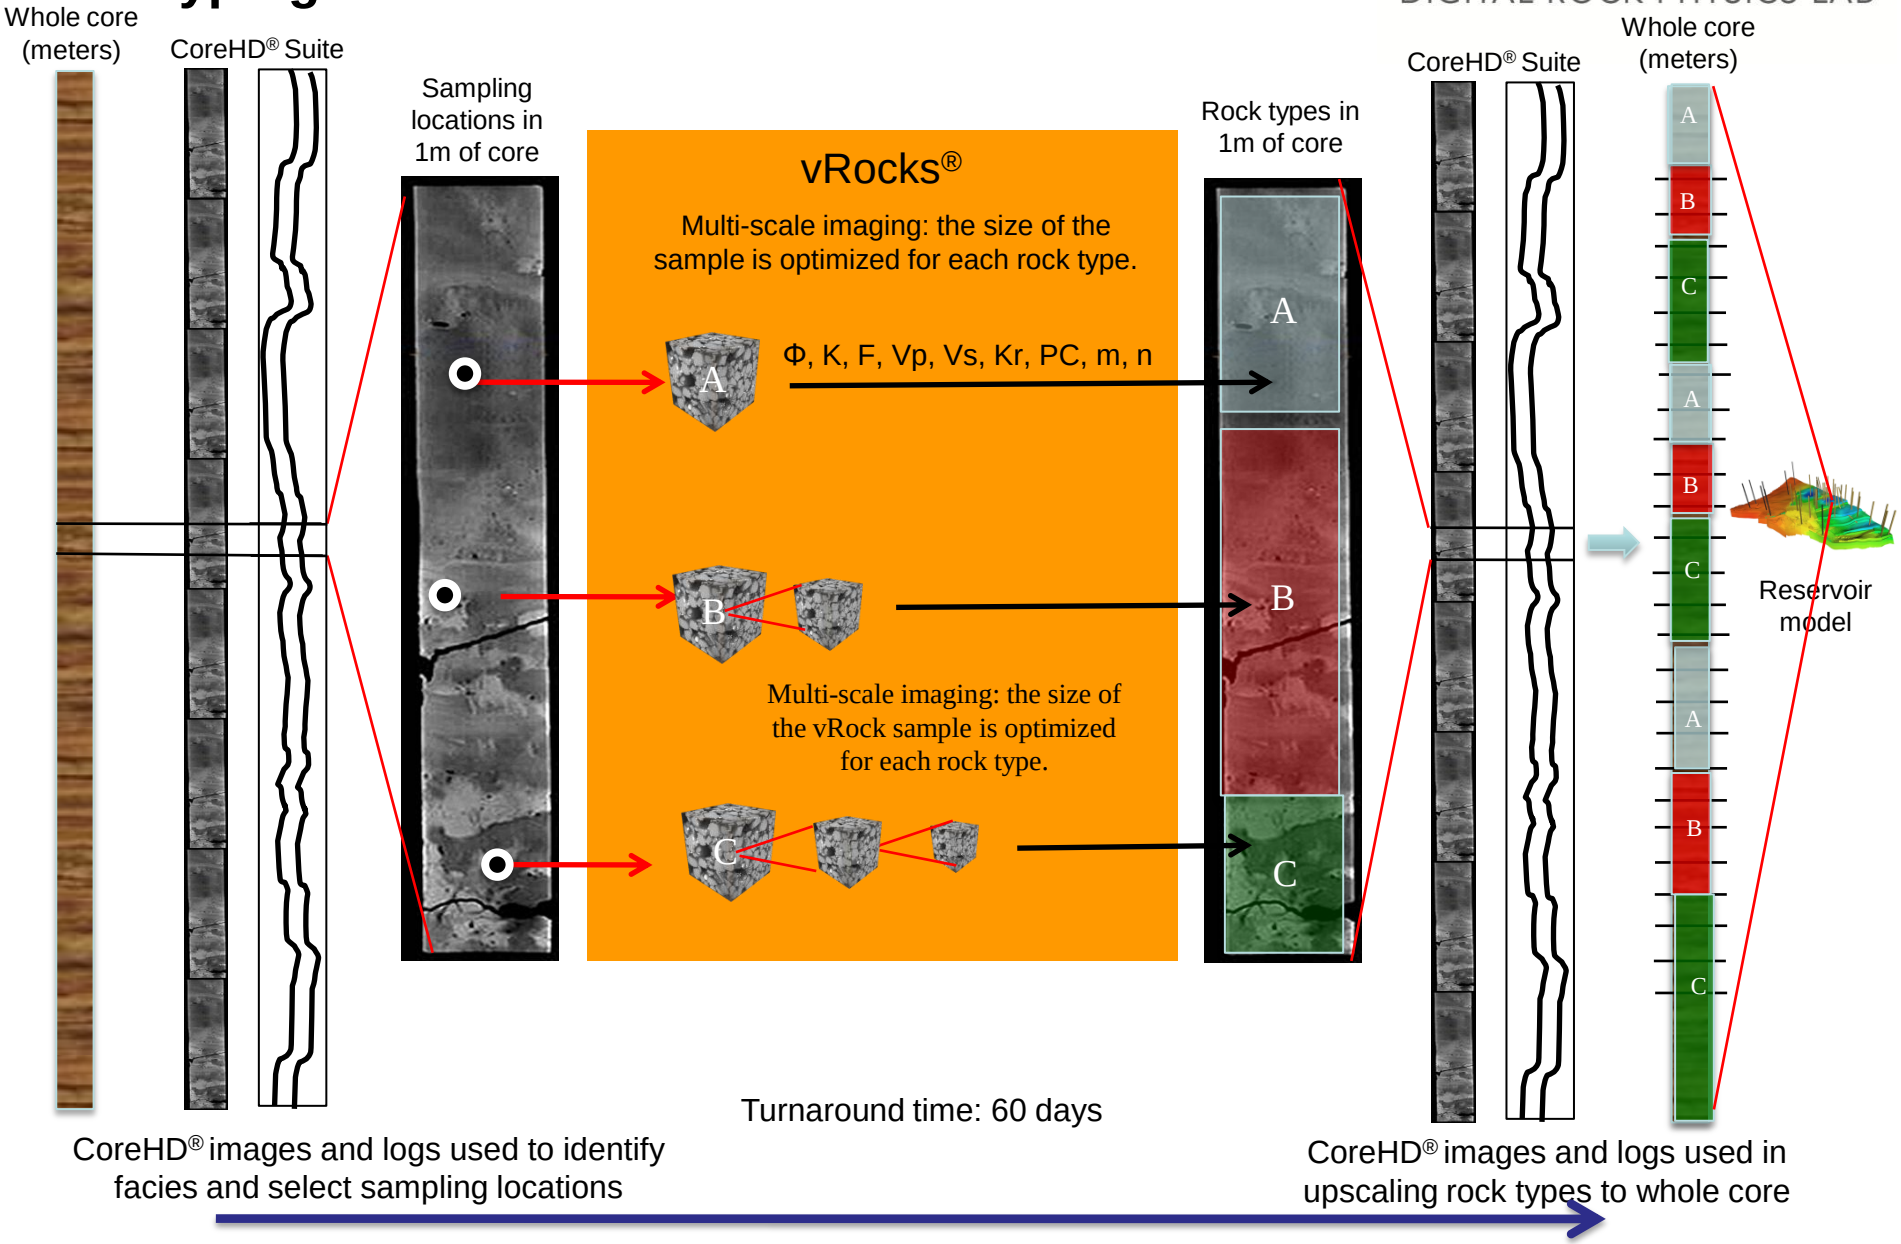

Fluid (Continuum)

$$\mathrm{p}, \nu, \vec{\mathrm{u}}(\vec{\mathrm{r}}, \mathrm{t}), \rho$$

Particle Ensemble

$$\mathrm{m_i}, \tau, \vec{\mathrm{v}}_\mathrm{i}(\vec{\mathrm{r}}, \mathrm{t})$$

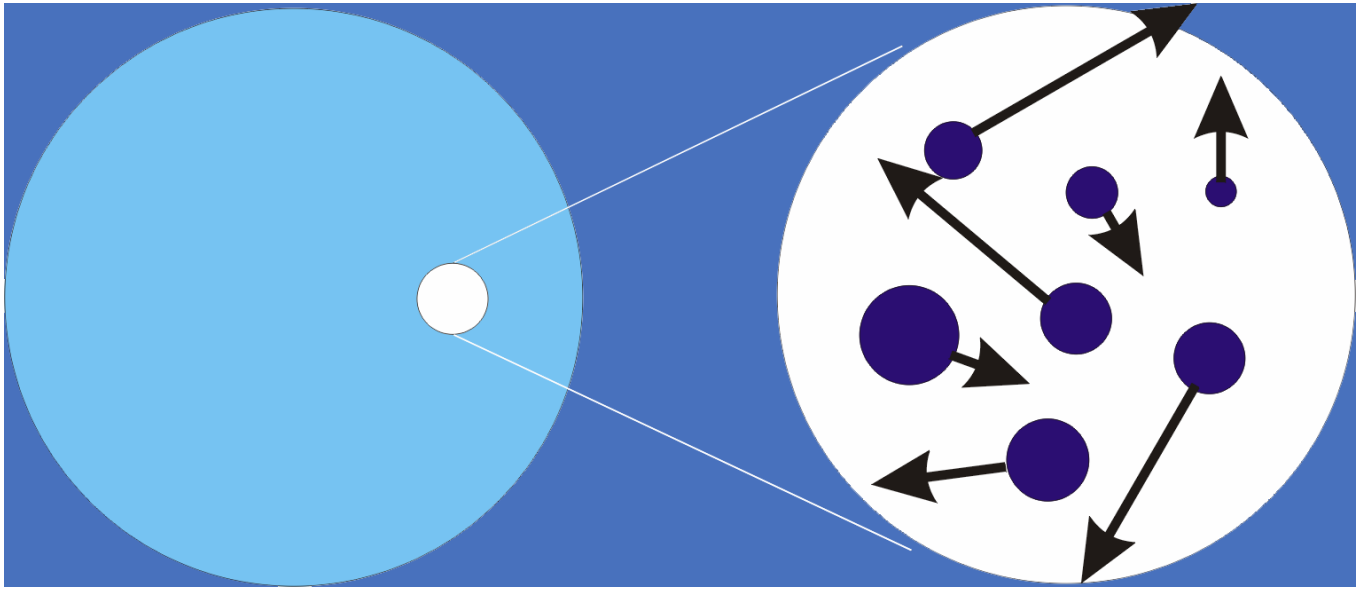

$$\frac{\partial \vec{u}}{\partial t} + \vec{u}\nabla \otimes \vec{u} = -\nabla p^* + \nu \Delta \vec{u}$$

$$m_i \frac{\partial v_i}{\partial t} = \sum_{i \neq j} \frac{\partial V(\|\vec{r}_i - \vec{r}_j\|)}{\partial \vec{r}_j}$$

Lattice-Boltzmann method 1

INGRAIN
DIGITAL ROCK PHYSICS LAB

Wolfram 86, Hasslacher 86

Lattice Boltzmann Equation (LBE)

$$f_i(t+\Delta t, \mathbf{x}+\mathbf{e}_i\Delta t) = f_i(t,\mathbf{x})+\Omega_i, \quad i=0,\ldots,b-1$$

f	Mass fractions
e	Microscopic velocity of the particles
t	Time
x	Space

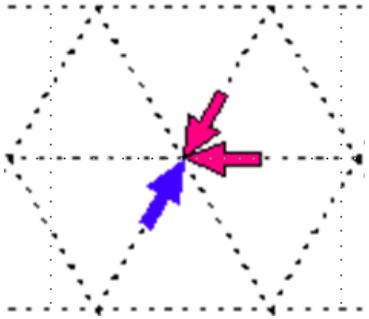

d2q6-Model

Wolfram 86, Hasslacher 86

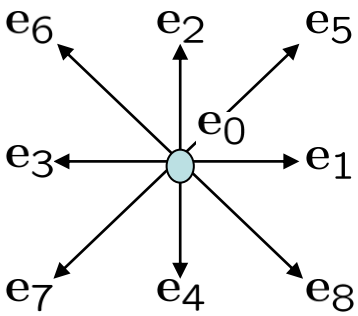

d2q9-Model

Qian 92

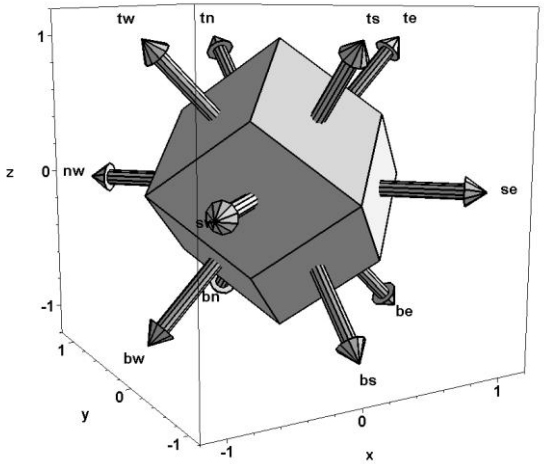

d3q13-Model

d'Humieres 01

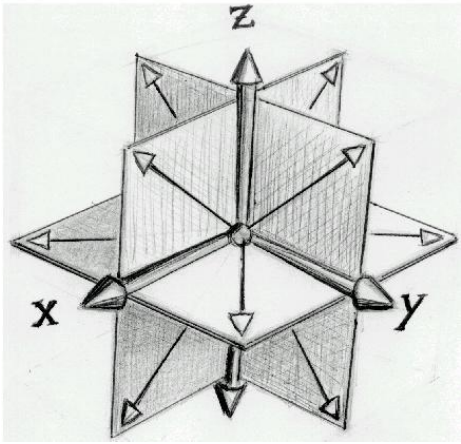

d3q19-Model

Qian 92

Lattice Boltzmann method 2

INGRAIN
DIGITAL ROCK PHYSICS LAB

D'Humieres 1992/2003,
Lallemand 2000,
Tölke et al. 2005

Flow Field: Multi-Relaxation-Time model (D3Q19)

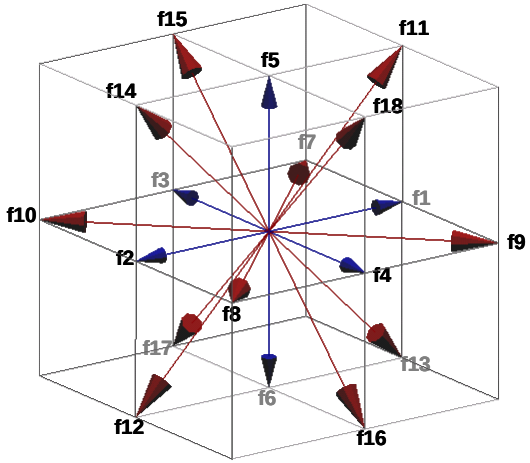

LBE: $$f_i(t+\Delta t,\ \vec{x}+\vec{e}_i\Delta t) = f_i(t,\vec{x}) + \Omega_i(t,\vec{x})$$

collision: $$\Omega = M^{-1}S\left[m - m^{eq}\right]$$

moments: $$m = Mf \qquad \rho = \sum_i f_i \qquad \vec{j} = \rho\vec{u} = \sum_i \vec{e}_i f_i$$

$$\boldsymbol{m} = (\rho, e, \epsilon, j_x, q_x, j_y, q_y, j_z, q_z, 3p_{xx}, 3\pi_{xx}, p_{ww}, \pi_{ww}, p_{xy}, p_{yz}, p_{xz}, m_x, m_y, m_z).$$

relaxation rates:

$$\left\{(s_9, s_{11}, s_{13}, s_{14}, s_{15}\right\} = -\frac{\Delta t}{\tau} \qquad \tau = 3\frac{\nu}{c^2} + \frac{1}{2}\Delta t$$

$$\left\{(s_1), (s_2), (s_4, s_6, s_8), (s_{10}, s_{12}), (s_{16}, s_{17}, s_{18})\right\} \in [-2; 0]$$

equilibrium moments:

$$m_0^{eq} = \rho$$
$$m_1^{eq} = e^{eq} = \rho\,(3\,T + \vec{u}^2 - 1)$$
$$m_2^{eq} = \rho\,(-\frac{9}{5}T + 1)$$
$$m_3^{eq} = \rho\,u_x$$
$$m_5^{eq} = \rho\,u_y$$
$$m_7^{eq} = \rho\,u_z$$

$$m_9^{eq} = 3p_{xx}^{eq} = \rho\,(2u_x^2 - u_y^2 - u_z^2)$$
$$m_{11}^{eq} = p_{zz}^{eq} = \rho\,(u_y^2 - u_z^2)$$
$$m_{13}^{eq} = p_{xy}^{eq} = \rho\,u_x\,u_y$$
$$m_{14}^{eq} = p_{yz}^{eq} = \rho\,u_y\,u_z$$
$$m_{15}^{eq} = p_{xz}^{eq} = \rho\,u_x\,u_z$$

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

$$\frac{\partial \vec{u}}{\partial t} + \vec{u}\nabla \otimes \vec{u} + \nabla p = \nu\Delta\vec{u}$$

LB Multiphase Model 1

INGRAIN
DIGITAL ROCK PHYSICS LAB

two-phase model:

(Rothmann-Keller 1991, Gunstensen 1993)

1. two phases (red/blue)
2. generation of surface tension (capillary forces)
3. separation of phases (immiscible)

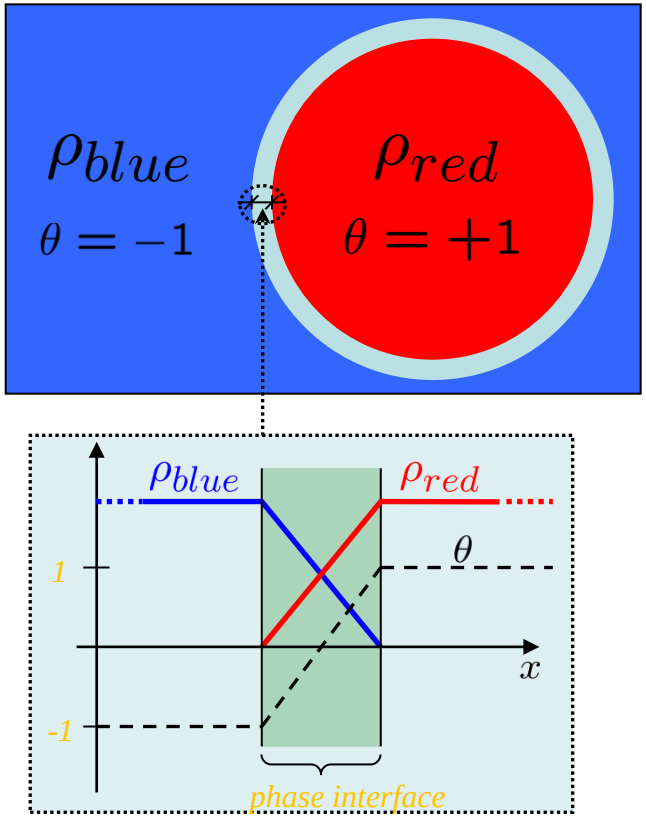

Lattice Boltzmann extensions:

phase field:

$$\theta = \frac{\rho_{red} - \rho_{blue}}{\underbrace{\rho_{red} + \rho_{blue}}_{\rho_\theta}} \in [-1; +1]$$

$$\rho_\theta = \rho_{red} + \rho_{blue}$$

LB Multiphase Model 2

surface tension:

$$m_i^{eq,ST} = m_i^{eq} + m_i^{ST}$$

(Kehrwald, 2004 PhD, Toelke et. al., 2005)

$$\begin{aligned}
m_1^{ST} &= -2\,\sigma\,|\vec{C}|\,(n_x{}^2 + n_y{}^2 + n_z{}^2) \\
m_9^{ST} &= \quad -\sigma\,|\vec{C}|\,(2\,n_x{}^2 - n_y{}^2 - n_z{}^2) \\
m_{11}^{ST} &= \quad -\sigma\,|\vec{C}|\,(n_y{}^2 - n_z{}^2) \\
m_{13}^{ST} &= \quad -\sigma\,|\vec{C}|\,(n_x\,n_y) \\
m_{14}^{ST} &= \quad -\sigma\,|\vec{C}|\,(n_y\,n_z) \\
m_{15}^{ST} &= \quad -\sigma\,|\vec{C}|\,(n_x\,n_z)
\end{aligned}$$

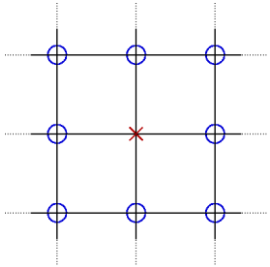

σ : surface tension

$\vec{C}$: gradient of phase field $\quad \vec{C}(t,\vec{x}) = \dfrac{3}{2\,c\,\Delta t}\sum_i w_i\,\dfrac{\vec{e}_i}{c}\,\theta(t,\vec{x}+\vec{e}_i\Delta t)$

n_α : vector normal to phase interface $\quad n_\alpha = \dfrac{C_\alpha}{|\vec{C}|}$

LB Multiphase Model 3

separation:

$$max\ \vec{C}\cdot j^p(after\ separation) = max\ \vec{C}\cdot\sum_i \vec{e}_i f_i^p(after\ separation)$$

maximum separation of phases by recoloring (optimization problem!)

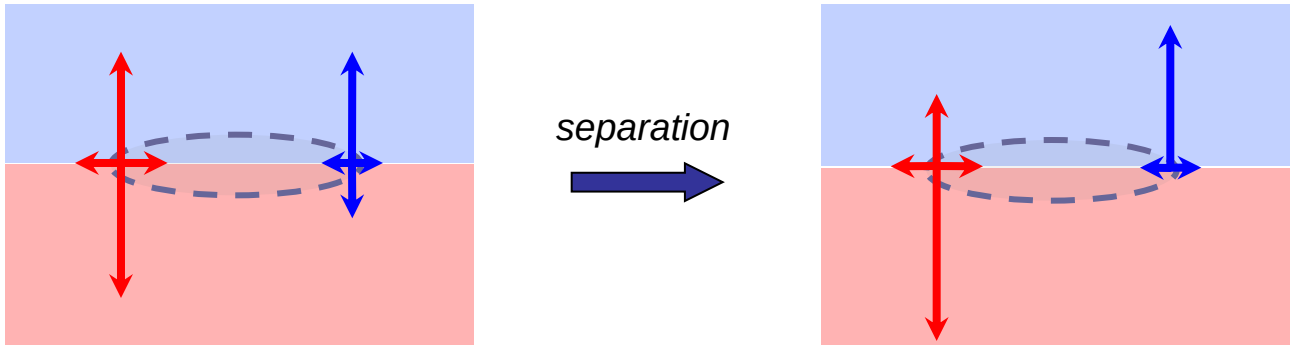

after collision *after separation*

advection of phase field by additional linear LBE:

$$g_i^p(t+\Delta t, \vec{x}+\vec{e}_i\Delta t) = g_i^{p,eq}\Big(\rho_p(t,\vec{x}), \vec{u}(t,\vec{x})\Big)$$

with $$g_i^{p,eq}\Big(\rho_p(t,\vec{x}), \vec{u}(t,\vec{x})\Big) = w_i\,\rho_p(t,\vec{x})\left(1+\frac{3}{c^2}\vec{e}_i\cdot\vec{u}(t,\vec{x})\right)$$

$$p = red,\ blue$$

Porescale: Multiphase-(Navier)-Stokes

Mass conservation:

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

Momentum conservation:

$$\rho(\mathbf{x})\left(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u}\cdot\nabla)\mathbf{u}\right) = -\nabla p + \nabla\cdot\left(\mu(\mathbf{x})(\nabla\mathbf{u} + \nabla\mathbf{u}^T)\right) + \rho(\mathbf{x})\mathbf{g}$$

Interface condition:

$$[\mathbf{u}]\big|_\Gamma = 0, \quad -\left[-p\mathbf{I} + \mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T)\right]\big|_\Gamma \cdot \hat{\mathbf{n}} = \sigma\kappa\hat{\mathbf{n}}$$

capillary pressure

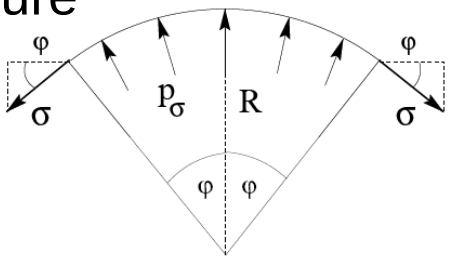

$$p_\sigma = \frac{\sigma}{R} = \sigma\ \kappa$$

contact angle:

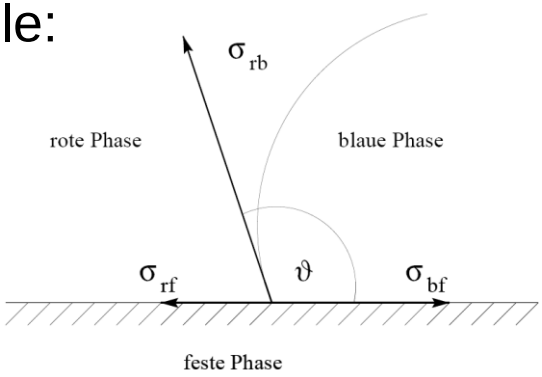

$$\sigma_{rf} - \sigma_{bf} - \sigma_{rb}\cos\vartheta = 0$$

Comparison CPU-GPU

INGRAIN
DIGITAL ROCK PHYSICS LAB

Intel Core i7

nVIDIA Tesla M2070

NEC SX-9A (16 CPUs)

Platform	Memory [MB]	Peak [GFLOPS]	BW [GB/s]	price [$]
Intel Core i7 980 XE (6 core)	24 000	80 (160 SP)	32 (15*)	ca. 4000
NEC SX-9A (16 CPUs, one node)	1 000 000	1600	4 000	expensive
nVIDIA GTX480	1 500	700 (1344 SP)	177 (120*)	ca. 500
nVIDIA Tesla M2070	6 000	515 (1030 SP)	148	ca. 4000

Process Collision and Propagation mapped to 'hardware'

INGRAIN
DIGITAL ROCK PHYSICS LAB

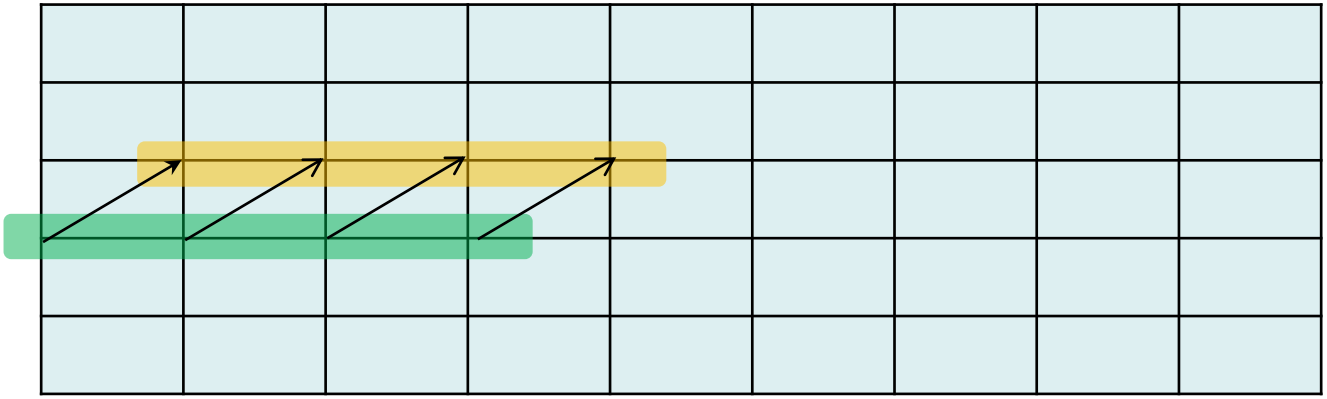

lattice

Load data from memory (Cache Line)
Memory Bandwidth

Collision (Processor)
Peak Performance

Propagation: Write data memory (Cache Line)
Memory Bandwidth

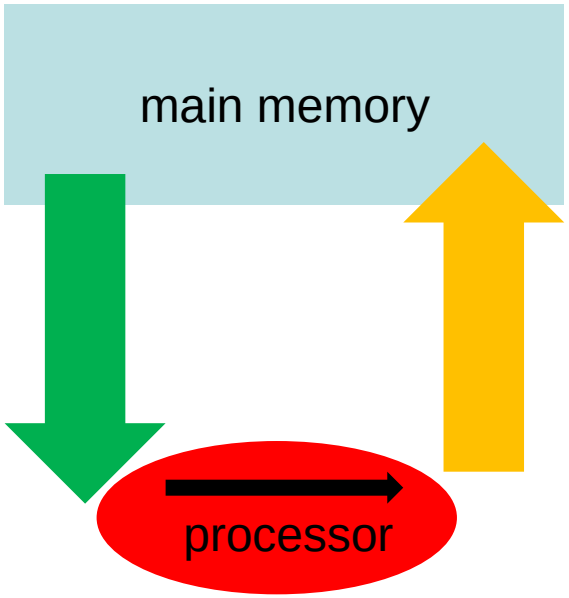

Achievable Performance of LB-Kernels

INGRAIN
DIGITAL ROCK PHYSICS LAB

Wellein 2006

$$P = \min\left\{\frac{MBW}{NB}, \frac{PEAK}{NF}\right\}$$

P	max. perf. in NUPS (nodal updates per second)
MBW	bandwidth bus
NB	number of bytes per cell
PEAK	peak performance
NF	number of FLOP per cell

D3Q13

NB = (14+13) (read+write) * 4 byte (float) = 108 bytes
NF = 260 FLOP

Average Notebook: MBW = 3.5 GB/s, PEAK=8 GFLOPS

- Limitation by ***Memory Bandwidth***:
  max NUPS = 3.5E9 Byte/s / 108 Byte/node
  = 33 E6 NUPS
- Limitation by ***Performance of CPU***:
  max NUPS3 = 8.0E9 FLOPS / 260 FLOP/node
  = 31 E6 NUPS

Reached:

grid	MLUPS	BW	Perf.
8^3	15	-	48 %
64^3	9	27 %	29 %

Common Unified Device Architecture (CUDA)

Programming model: Parallelism on 2 levels

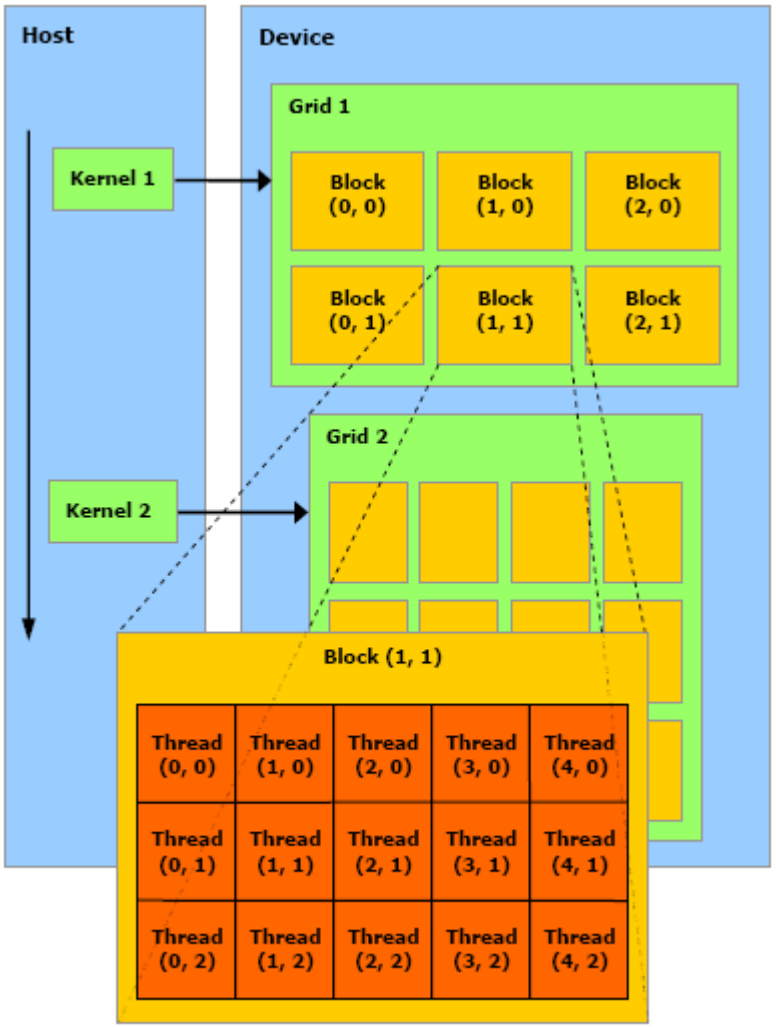

Level 1:
Grid of blocks (no communication)

Level 2:
Block of threads (shared memory)

memory access pattern

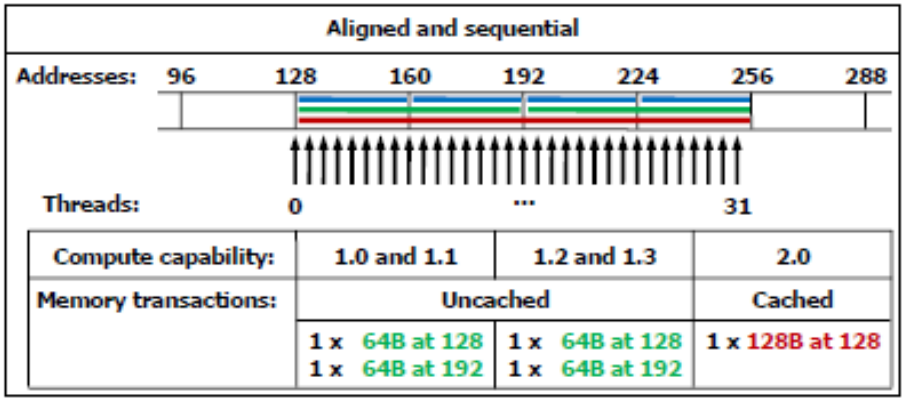

Compute capability:	1.0 and 1.1	1.2 and 1.3	2.0
Memory transactions:	Uncached		Cached
	1 x 64B at 128 1 x 64B at 192	1 x 64B at 128 1 x 64B at 192	1 x 128B at 128

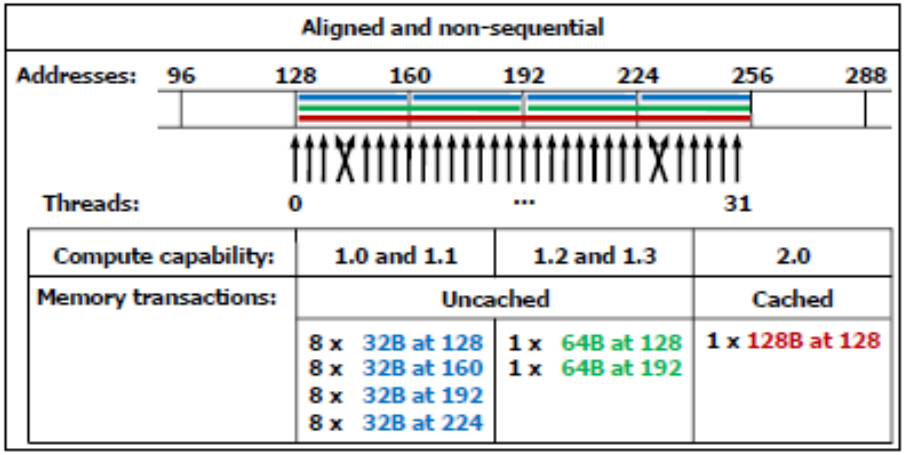

Compute capability:	1.0 and 1.1	1.2 and 1.3	2.0
Memory transactions:	Uncached		Cached
	8 x 32B at 128 8 x 32B at 160 8 x 32B at 192 8 x 32B at 224	1 x 64B at 128 1 x 64B at 192	1 x 128B at 128

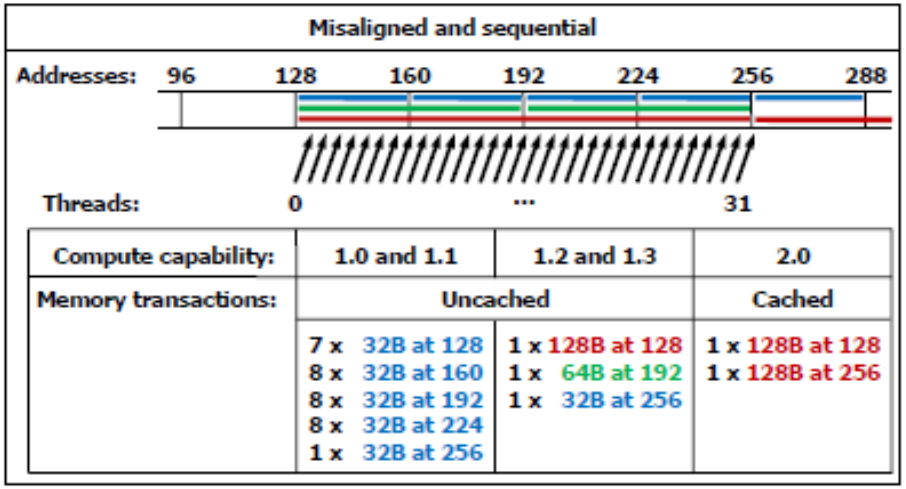

Compute capability:	1.0 and 1.1	1.2 and 1.3	2.0
Memory transactions:	Uncached		Cached
	7 x 32B at 128 8 x 32B at 160 8 x 32B at 192 8 x 32B at 224 1 x 32B at 256	1 x 128B at 128 1 x 64B at 192 1 x 32B at 256	1 x 128B at 128 1 x 128B at 256

INGRAIN
DIGITAL ROCK PHYSICS LAB

CUDA Programming Guide

Mapping *space* location – *memory* location

INGRAIN
DIGITAL ROCK PHYSICS LAB

Toelke et al. 2008

D3Q13 model:
Mapping 3D ***space*** to 1D ***memory***:

$$a = \begin{cases} 0 & \text{if } j \text{ even and } k \text{ even} \\ 0 & \text{if } j \text{ odd and } k \text{ odd} \\ 1 & \text{if } j \text{ odd and } k \text{ even} \\ 1 & \text{if } j \text{ even and } k \text{ odd} \end{cases}$$

$$x = h\,(a + 2\,i)$$

$$y = h\,j$$

$$z = h\,k$$

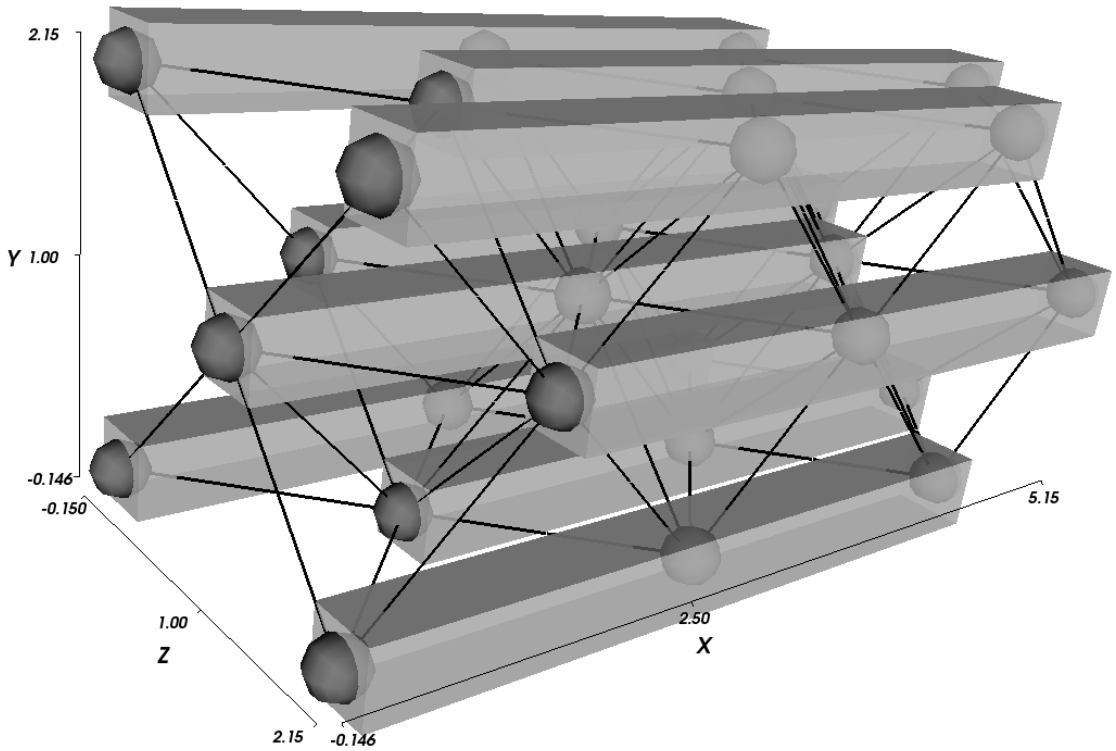

C-Code:

```
int m = nx*(ny*k + j) + i;
float x =  h * ( (j&0x1)^(k&0x1) + i*2 );
float y =  h * j;
float z =  h * k;
```

lattice size:
2 *nx* x *ny* x *nz* (e.g. 128 x 128 x 512)
data structure:
nx x *ny* x *nz* (e.g 64 x 128 x 512)

Mapping of lattice to parallel architecture

INGRAIN
DIGITAL ROCK PHYSICS LAB

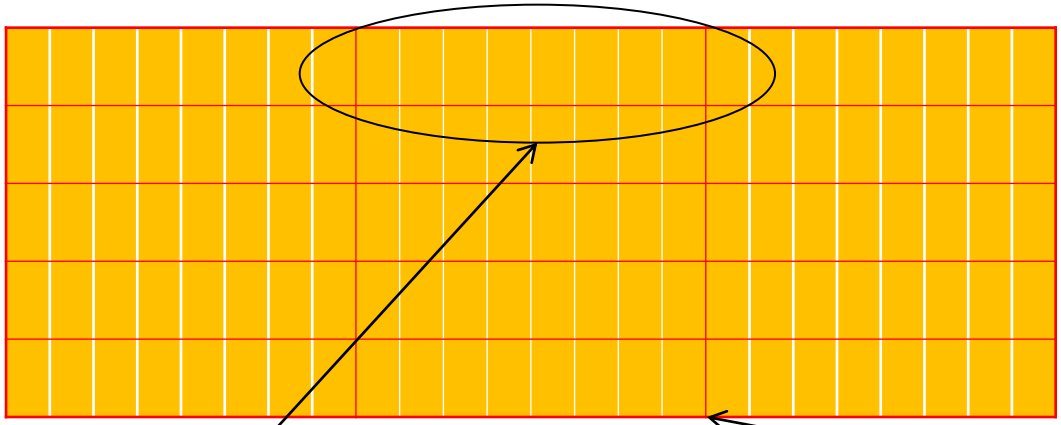

Level 1

Level 2

T1	T2	T3	T4	T5	T6	T7	T8

Block (0,0)	Block (1,0)	Block (2,0)
Block (0,1)	Block (1,1)	Block (2,1)
Block (0,2)	Block (1,2)	Block (2,2)
Block (0,3)	Block (1,3)	Block (2,3)
Block (0,4)	Block (1,4)	Block (2,4)

- block of threads – parallel
- nonlocal operations via shared memory and synchronization

- grid of blocks – parallel
- no communication

LB D2Q9 model with CUDA

INGRAIN
DIGITAL ROCK PHYSICS LAB

Jonas Tölke (2008): ***Implementation of a Lattice Boltzmann kernel using the Compute Unified Device Architecture***, Computing and Visualization in Science, DOI 10.1007/s00791-008-0120-2

- Load streams (9 streams for d2q9)
- Complex computations (collision)
- Write streams (9 streams for d2q9) to correct address (propagation)

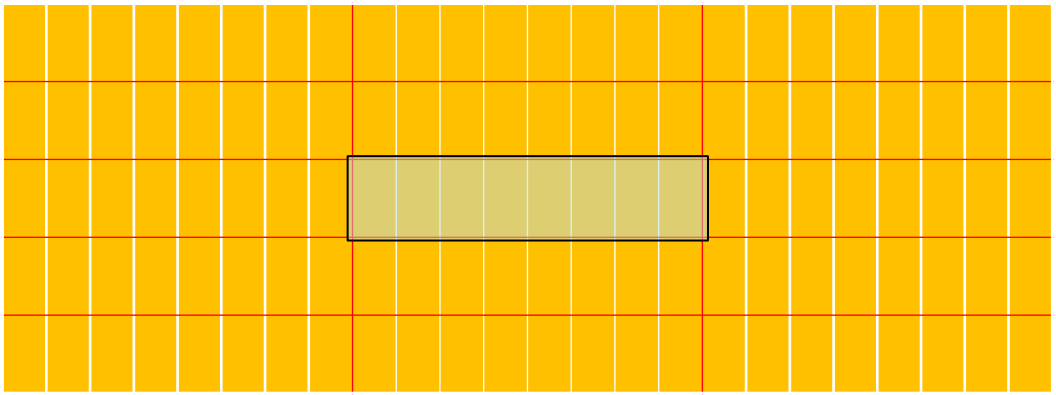

- north
- south
- east
- west
- northeast
- northwest
- southwest
- southeast

LB D3Q13: Mapping of grid to parallel architecture

INGRAIN
DIGITAL ROCK PHYSICS LAB

- x-index mapped to Threads (16-256 (32-512 lattice size))
- y- and z-index mapped to grid (at least 16 in sum)

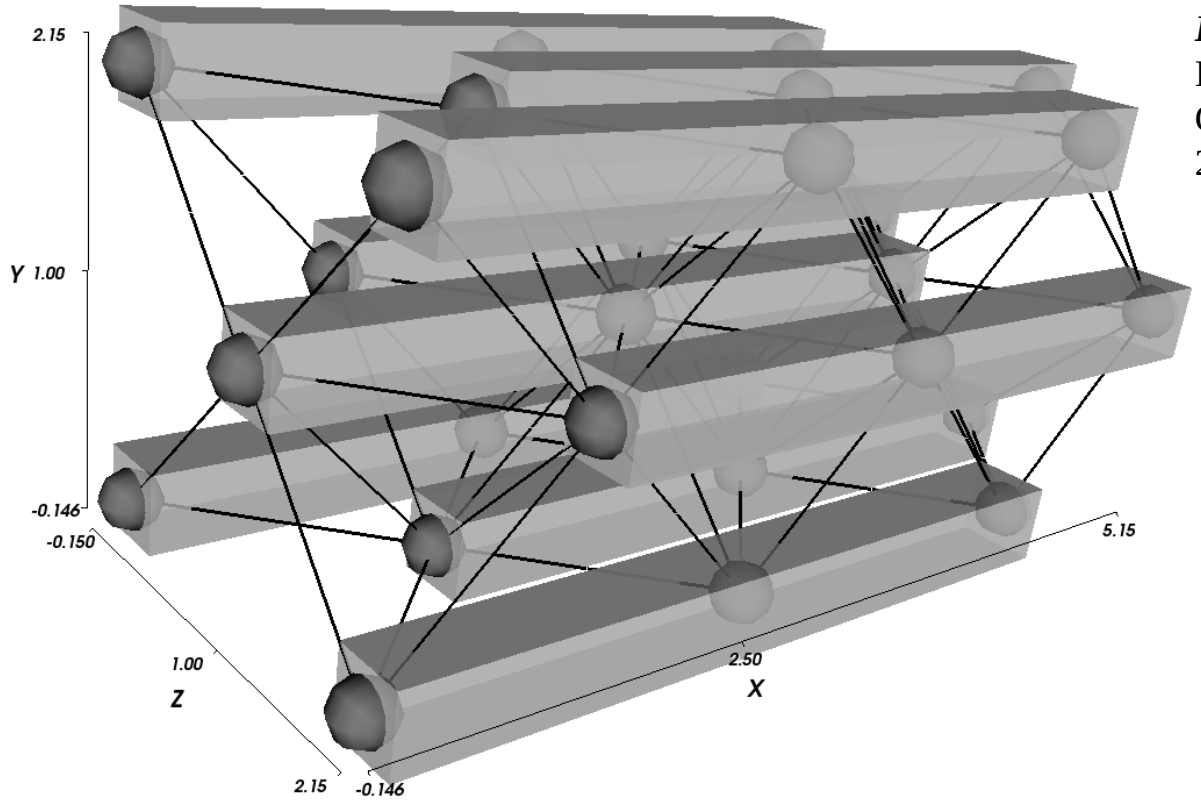

Tölke, J. and Krafczyk, M. (2008): ***TeraFLOP computing on a desktop PC with GPUs for 3D CFD***, International Journal of Computational Fluid Dynamics, 22:7, 443 — 456

Increase in computational speed 1

INGRAIN
DIGITAL ROCK PHYSICS LAB

simple approach with if statements

```
int nThreads = nx;
int nBlocks.x = ny;
int nBlocks.y = nz;
Kernel<<< nBlocks, nThreads >>>(data,geo,...)

__global__ void Kernel(float* data, int* geo, ....)
    int ny = gridDim.x;
    int nz = gridDim.y;
    int x = threadIdx.x;
    int y = blockIdx.x;
    int z = blockIdx.y;

    int m = nx*(ny*z + y) + x;
    if(geo[m]==FLUID) {
        // do average work
    }
    else if(geo[m]==BOUNDARY){
        // do almost nothing
    }
    else if(geo[m]==SOLID){
        // do nothing
    }
    ....
```

Optimize:

- Eliminate if statements
- Launch only elements in parallel which have about the same work load

Increase in computational speed 2

INGRAIN
DIGITAL ROCK PHYSICS LAB

Preprocessing: build a list for each node type

Compute capability >= 1.2

```
vector fluidNodes, boundNodes;
int a,b;   a=0; b=0;
Preprocssing
 for(k=0 ; k<nz ; k++){
   for(j=0 ; j<ny ; j++){
     for(i=0 ; i<nx; i++){
         m = nx*(ny*k + j) + i;
         if(geo[m]==FLUID) {
                fluidNodes[a] = m;
                a++;
         }
         else if(geo[m]==BOUNDARY){
                boundNodes[b] = m;
                b++;
         }
         ......
```

Kernel for each fluid type

```
__global__ void KernelFluid ( int iter, int  nFluid,
                                        list fluidNodes,...){
 int numThreads = blockDim.x;
 int numGrid = gridDim.x;
 int tx = threadIdx.x;
 int bx = blockIdx.x;
 for(int iloop=0; iloop<iter; iloop++){
    l = numThreads*(iloop*numGrid+bx)+tx;
    if(l >= nCompute ) break;
    m =  fluidNodes[l];
    .............

__global__ void KernelBound(.....
```

Kernel calls

```
int  nThreads = 64;
int  nBlocks = 32;
iter =  (nFluid +nThreads*nBlocks-1)/(nThreads*nBlocks)
KernelFluid<<< nBlocks, nThreads >>>(iter, nFluid, fluidNodes,...)
….
KernelBound<<< nBlocks, nThreads >>>(...)
```

Increase in computational speed 3

INGRAIN
DIGITAL ROCK PHYSICS LAB

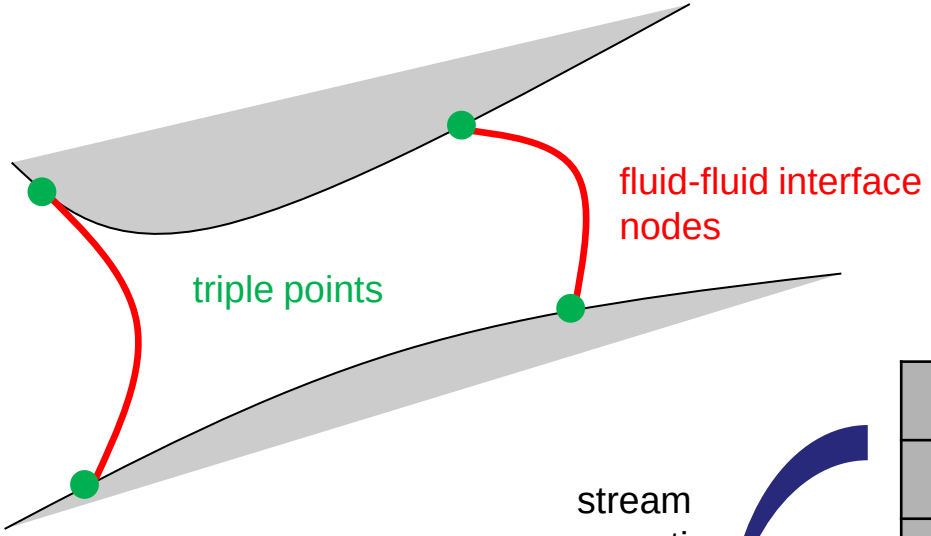

Mark Harris, Shubhabrata Sengupta, and John D. Owens. "Parallel Prefix Sum (Scan) with CUDA". In Hubert Nguyen, editor, GPU Gems 3, chapter 39, pages 851–876. Addison Wesley, August 2007

computational grid

1	**2**	**3**	**4**	**5**	**6**
7	8	9	10	11	12
13	14	15	16	17	18
19	20	21	22	23	24

stream compaction

list of points changes each time step

use (dynamic) stream compaction (cudpp) to generate the lists at each time step

list triple points nodes

2	**4**	**19**	**22**

list interface nodes

8	**9**	**11**	**14**	**17**

CUDPP: CUDA Data Parallel Primitives Library

- parallel prefix-sum (“scan”), sort, reduction.
- stream compaction
- building data structures such as trees and summed-area tables

Reduction of memory consumption 1

INGRAIN
DIGITAL ROCK PHYSICS LAB

sparse data structure 1:

f[0..8, nnodes]:
a list for mass fractions at computational nodes only

memLoc[nx*ny]:
a full matrix of the memory locations for the mass fractions, needed for operations involving neighbors

computeNodes[nnodes]:
a list of the location of compute nodes within memLoc

memory: (9+1) *nnodes + nx*ny =

(10*porosity+1)*n

- for porosities between 20% and 80 % good reduction in memory consumption
- memory consumption goes not to zero if porosity goes to zero
- computing of neighbor index in memLoc is needed for each access

geo-matrix + enumeration (temporary):

1	**2**	**3**	**4**	**5**	**6**
7	8	9	10	11	12
13	14	15	16	17	18
19	20	21	22	23	24

memLoc:

-1	**-1**	**1**	**-1**	**-1**	**-1**
-1	-1	2	3	-1	-1
-1	-1	4	5	-1	-1
-1	-1	-1	6	7	-1

computeNodes:

3	**9**	**10**	**15**	**16**	**22**	**23**

M. Schulz, J. Tölke, M. Krafczyk, E. Rank: Parallel single- and multiphase CFD-applications using lattice Boltzmann methods. *Transactions of the first joint HLRB and KONWIHR status and result workshop*, Springer. 91-103 (2003)

Reduction of memory consumption 2

INGRAIN
DIGITAL ROCK PHYSICS LAB

sparse data structure 2

f[0..8, nnodes]
a list for mass fractions at computational nodes only

neighbors[0..7,nnodes]
a list of all neighbors

memory: (9+8) *nnodes = 17*porosity*n

direct access to neighbors

memory consumption goes to zero if porosity goes to zero

for porosities 50 % and greater no reduction or even increase in memory consumption

(10*porosity+1) = 17*porosity

porosity = 1/7

memLoc (temporary):

-1	-1	1	-1	-1	-1
-1	-1	2	3	-1	-1
-1	-1	4	5	-1	-1
-1	-1	-1	6	7	-1

neighbors:

e	-1	3	-1	5	-1	7	-1
n	-1	1	-1	2	3	5	-1
w	-1	-1	2	-1	4	-1	6
s	2	4	5	-1	6	-1	-1
ne	-1	-1	-1	3	-1	-1	-1
nw	-1	-1	1	-1	2	4	5
sw	-1	-1	4	-1	-1	-1	-1
se	3	5	-1	6	7	-1	-1

Absolute permeability

INGRAIN
DIGITAL ROCK PHYSICS LAB

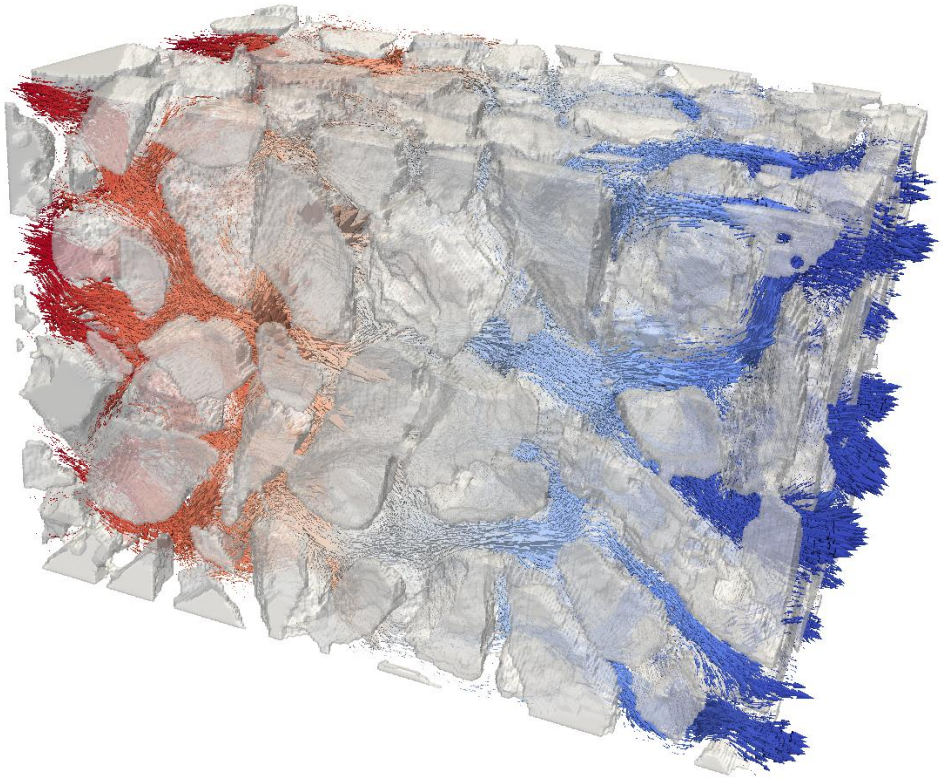

Darcy's law:

$$q = -\frac{k}{\mu}\nabla P$$

- sandstone
- porosity=30%
- sparse data structure 1
- size=200*200*300
- single phase D3Q13 model
- GTX 480, FERMI

	Throughput in GB/s
DeviceToDeviceCopy	118
LB full matrix, only fluid	80
sandstone, sparse	55

Multi-GPU: Teraflop Computing

INGRAIN
DIGITAL ROCK PHYSICS LAB

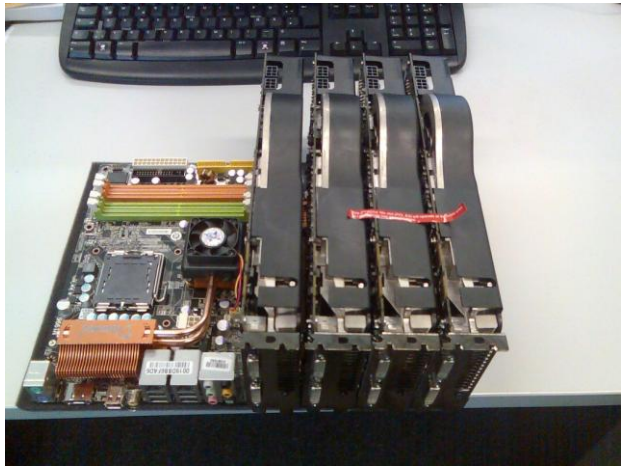

Communication between GPUs:

- PCI-E
- Latency like Front Side Bus (e.g. 266 MHz)
- **PThreads**
- **CUDA**

	Bandwidth [GB/s]	Latency [ns]
PCI-E2.0(x16)	8.0	O(100)
Infiniband QDR 4X	3.2	1 000
10G-Ethernet	1.25	5 000

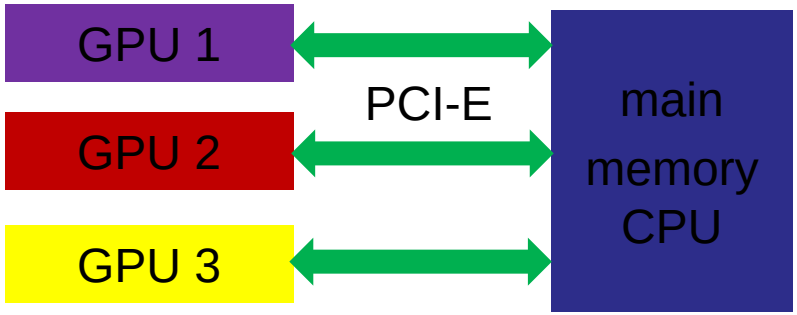

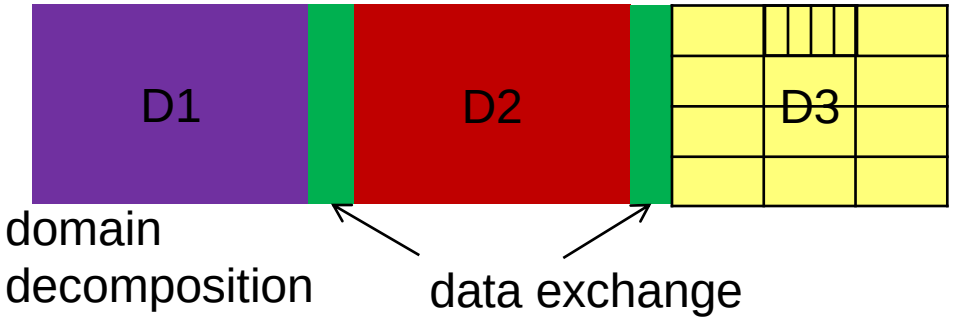

Multi-GPU: communication hiding with streams

INGRAIN
DIGITAL ROCK PHYSICS LAB

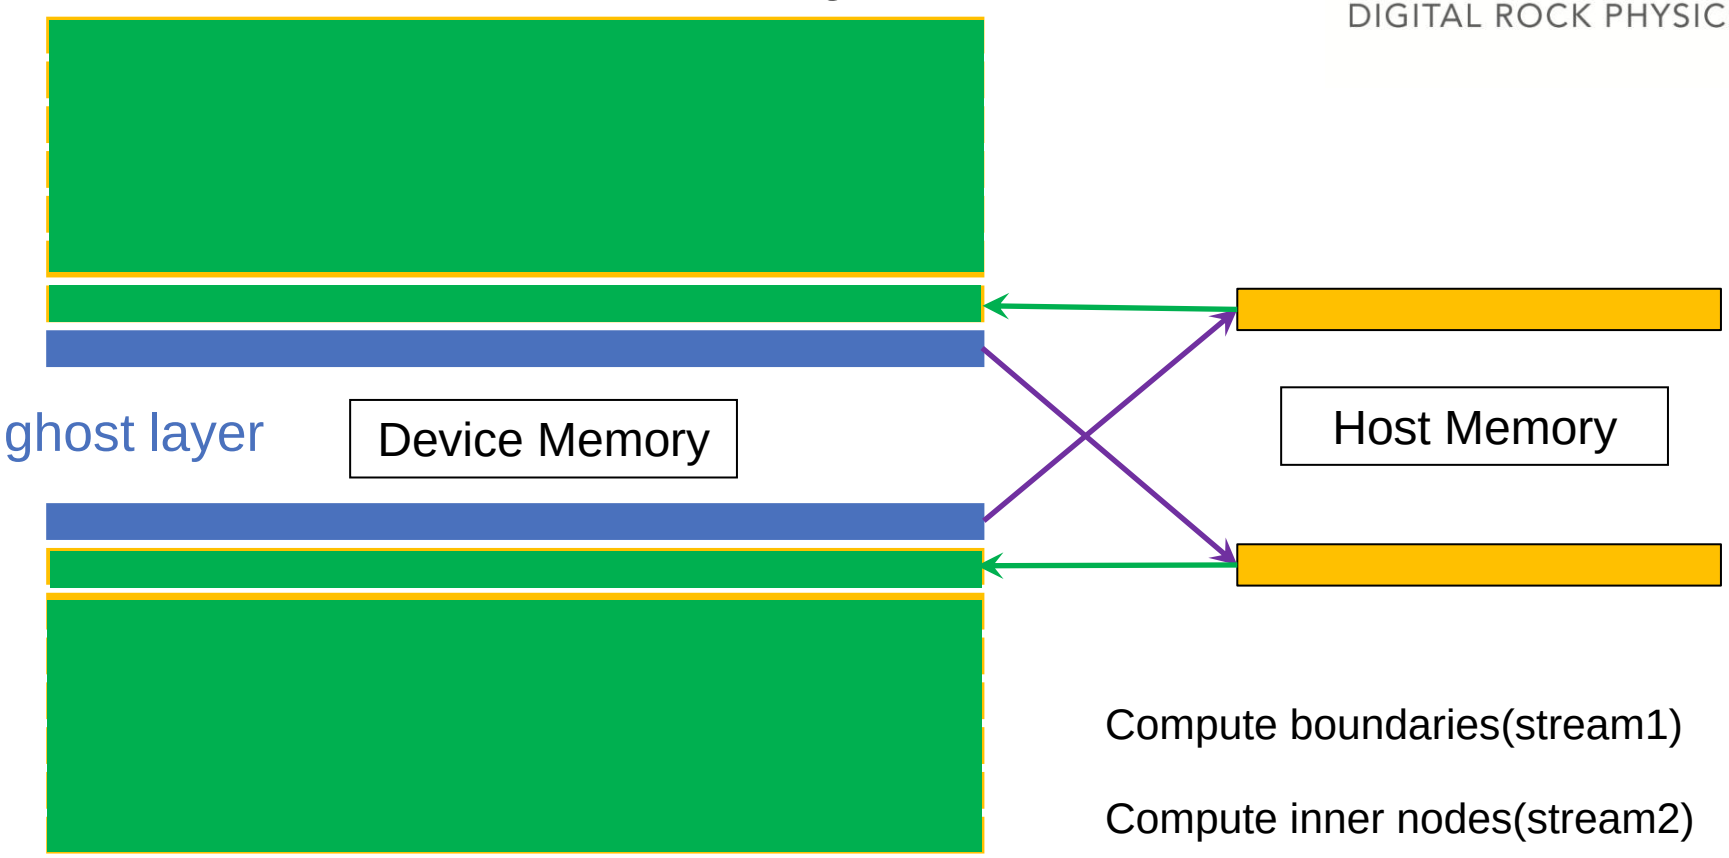

Compute boundaries(stream1)

Compute inner nodes(stream2)

async copy Device to Host(stream1)

synchronize(stream1)

MPI_Sendrecv / barrier Threads

async copy Host Device (stream1)

Multi-GPU: relation parallel efficiency and bandwidth

INGRAIN
DIGITAL ROCK PHYSICS LAB

- Assuming Communication Hiding with Comp. Cap. >= 1.2
- Application performance limited by bandwidth
- Bandwidth GPU: BWg O(100GB/s)
- Bandwidth PCI-E: BWp O(5GB/s)
- Bandwidth Infinband: BWi O(1GB/s)

Within one node: Pthreads+PCI-E

Several nodes: MPI+Infiniband

Domain decomposition: Slices in x

Minimum size of sub-domain in x: $x/2 = BWg/BW\_p,i \rightarrow x=40,200$

Domain decomposition: Cubes

Minimum edge length $n^3/(6*n^2) = BWg/BWp\_p,i \rightarrow n=120,600$

Capillary pressure 1

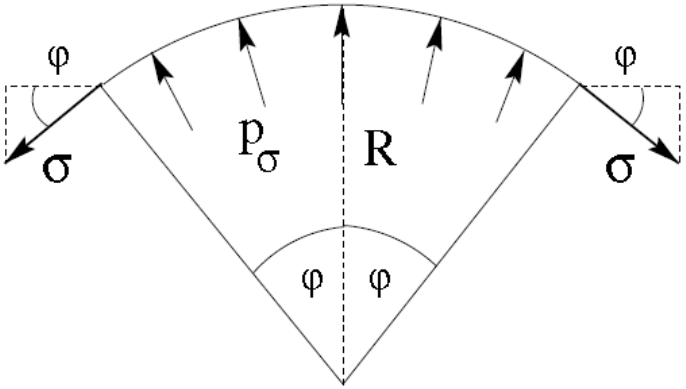

interfacial tension ρ
capillary pressure p = σ/R

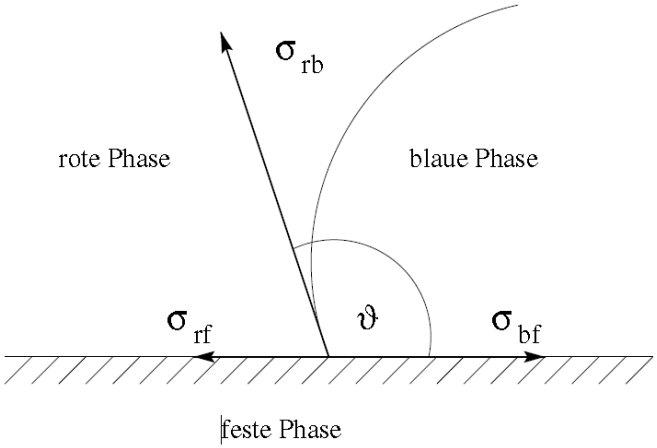

contact angle

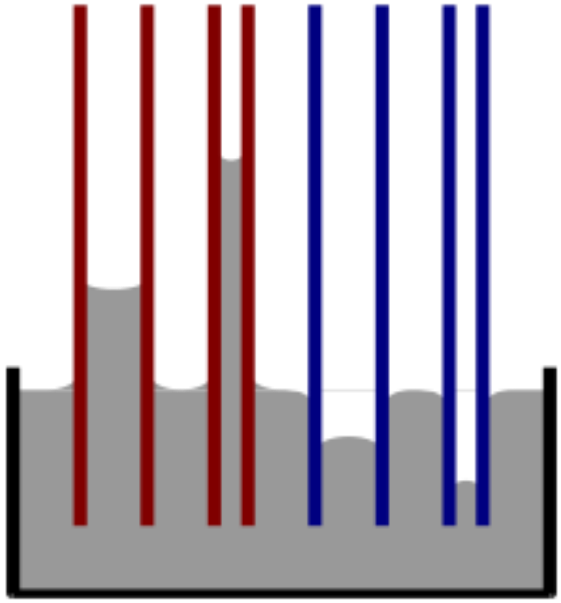

capillary rise

Capillary pressure 2

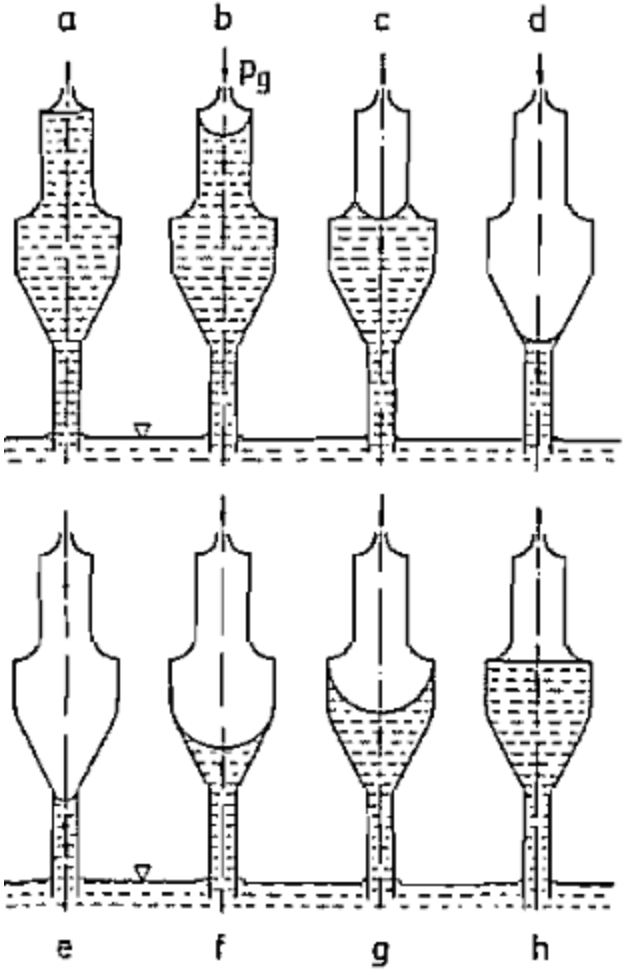

Schubert 1982

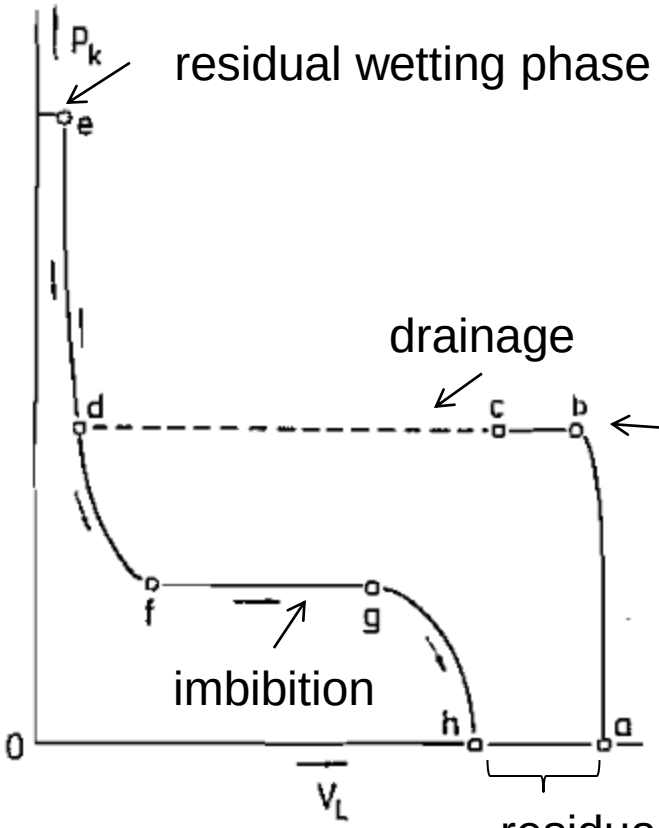

residual wetting phase content

drainage

entry pressure

imbibition

residual non-wetting phase content

Numerical Experiments: Drainage-Imbibition

INGRAIN
DIGITAL ROCK PHYSICS LAB

- sandstone
- LB multiphase-model
- Sparse grids

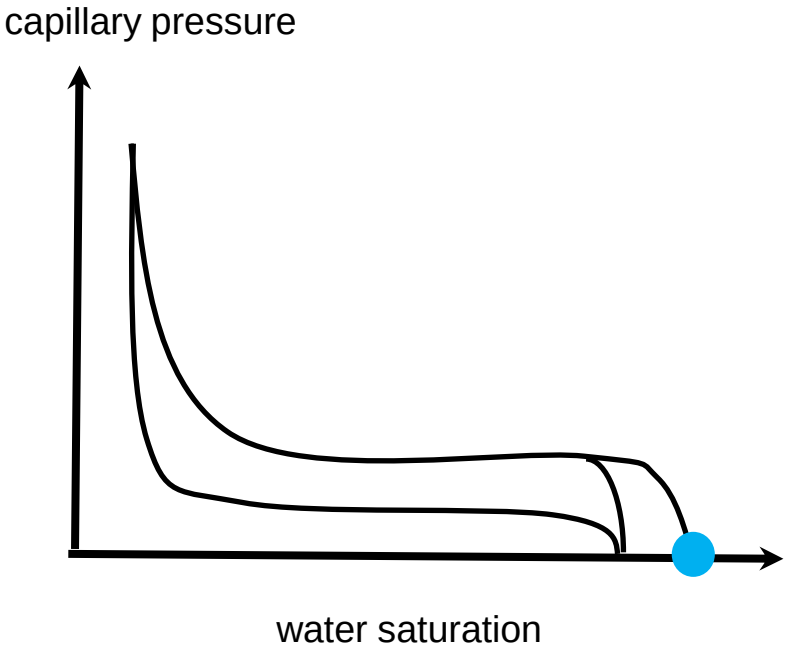

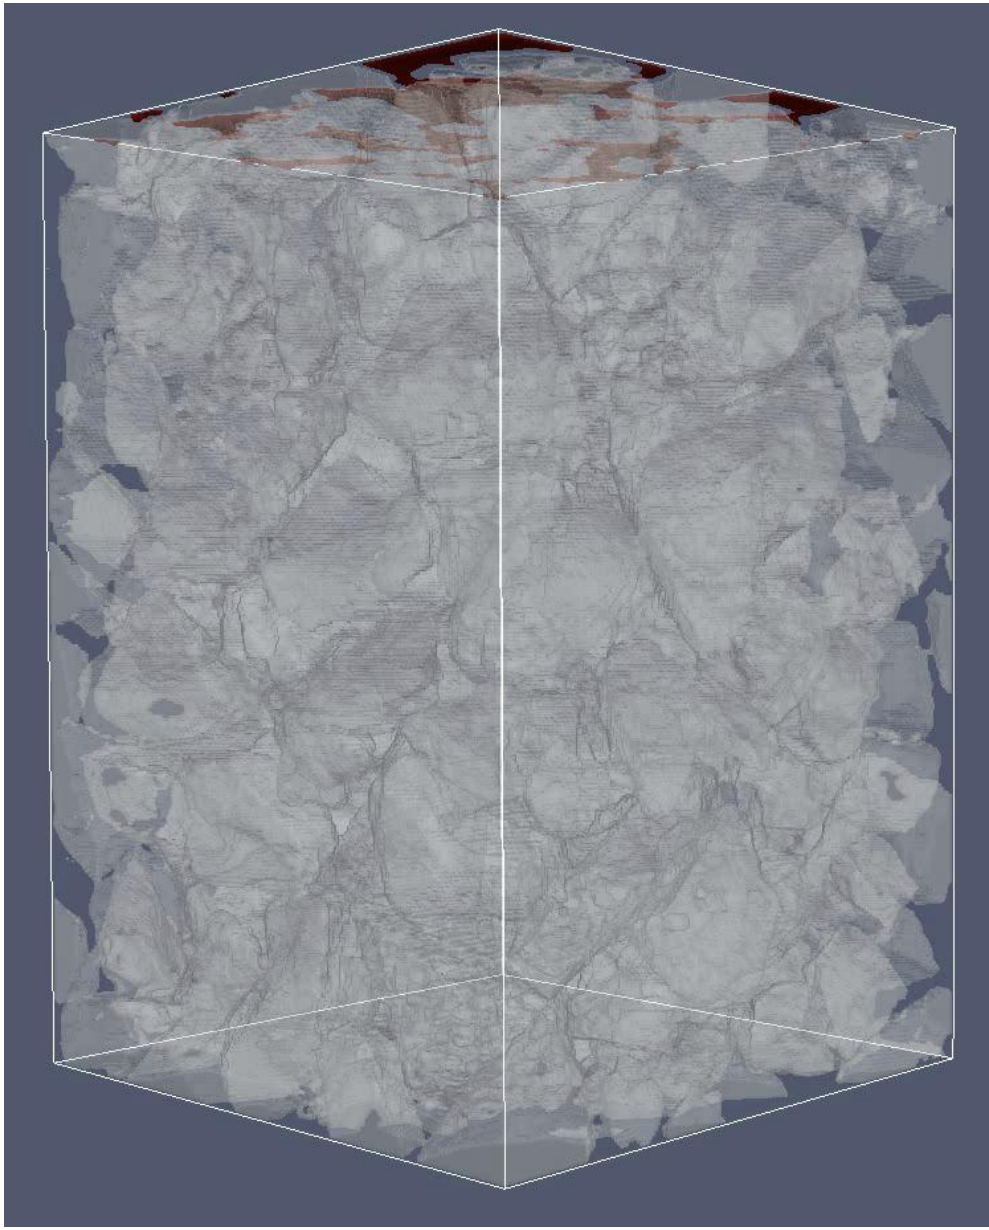

Relative permeabilitiy

INGRAIN
DIGITAL ROCK PHYSICS LAB

- two phases: Darcy's law has to be extended.
- second phase acts like an additional obstacle to the other phase and reduces though the permeability.

$$q = -\frac{k}{\mu}\nabla P$$

$$q_o = -\frac{k\ kr_o(S_w, Ca)}{\mu_o}\nabla P$$

$$q_w = -\frac{k\ kr_w(S_w, Ca)}{\mu_w}\nabla P$$

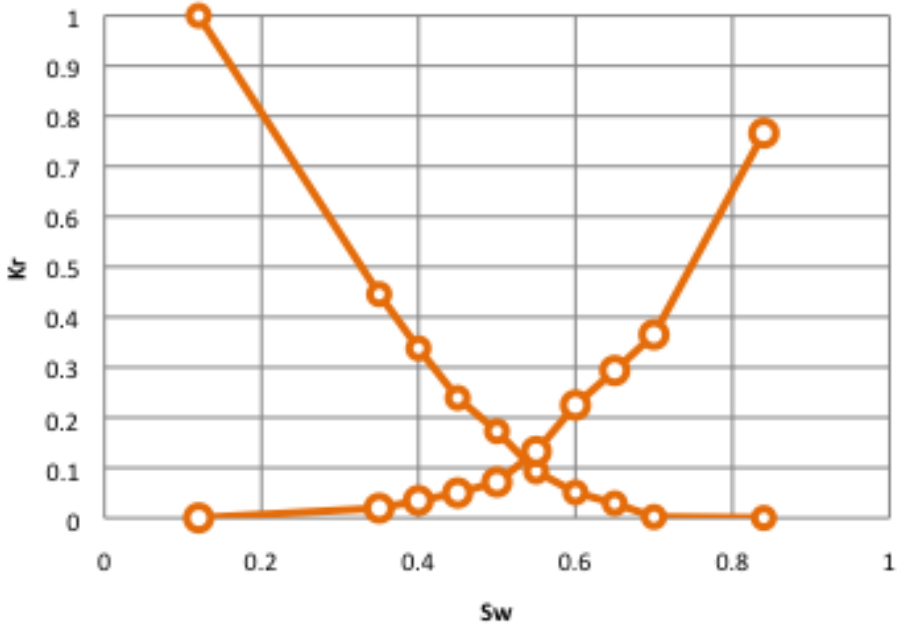

relative permeability sandstone

INGRAIN
DIGITAL ROCK PHYSICS LAB

fractional flow: flow of water and oil simultaneously

- LB multiphase-model
- Sparse grid

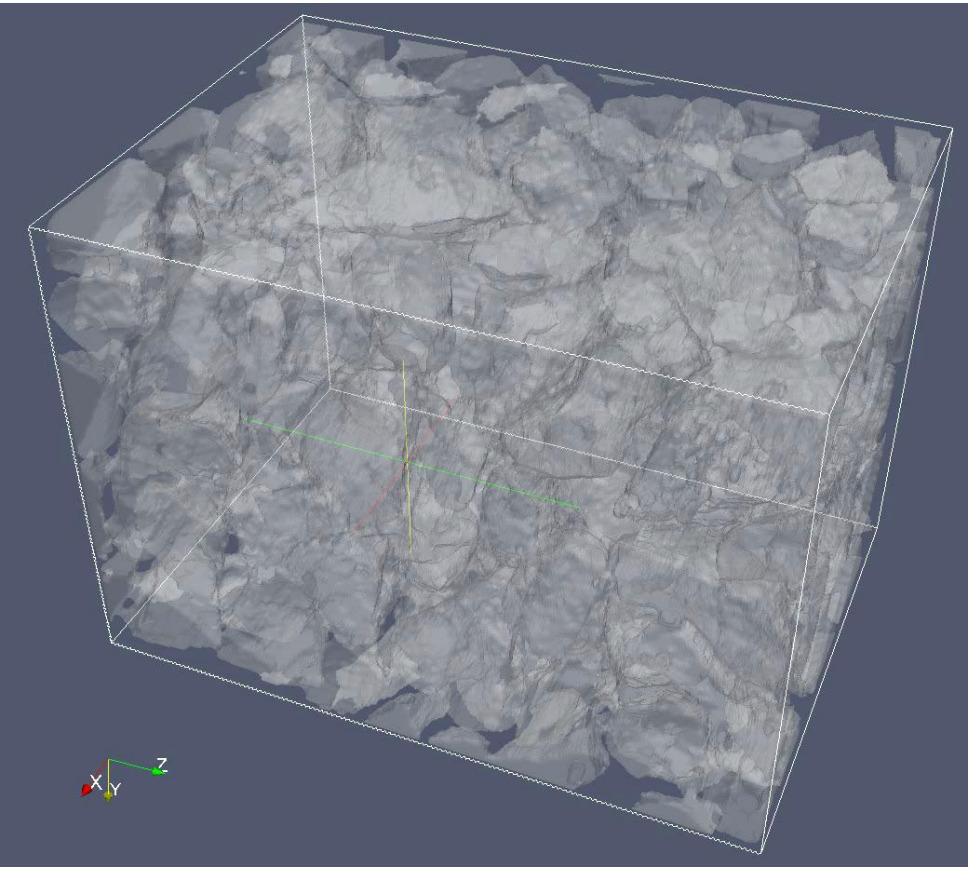

high Capillary number

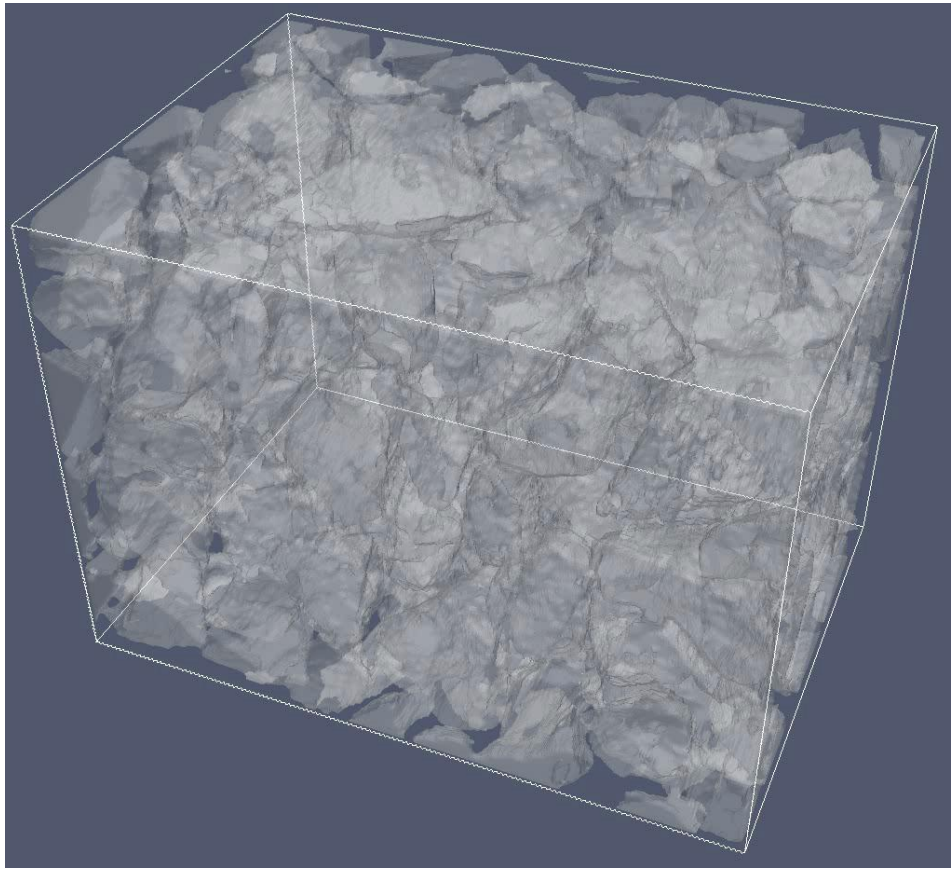

low Capillary number

Results

INGRAIN
DIGITAL ROCK PHYSICS LAB

Ingrain Data and Lab Data

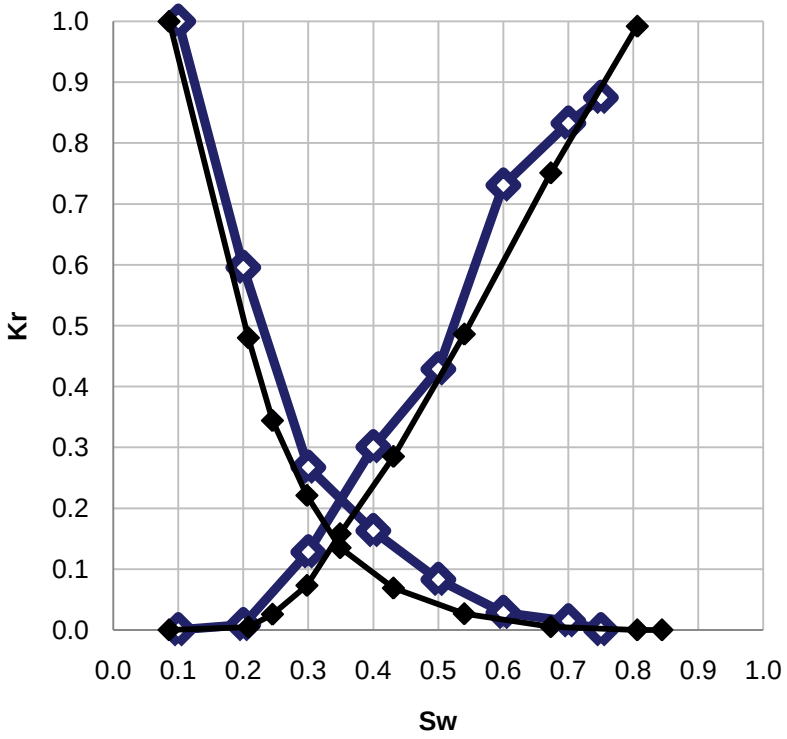

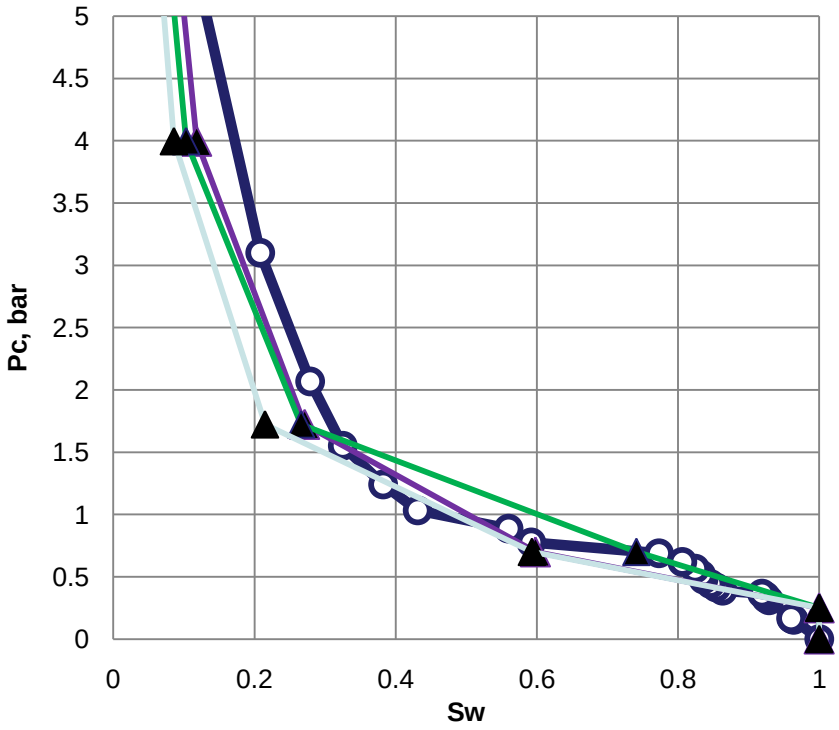

publications

INGRAIN
DIGITAL ROCK PHYSICS LAB

AS Grader; Y. Mu; J. Toelke; C. Baldwin; Q. Fang, G. Carpio, *Ingrain*; B. Stenger; T. Al Dayyani; Z. Kalam, *ADCO*: ***Estimation of relative permeability using the Lattice Boltzmann method for fluid flows in a Cretaceous Formation, Abu Dhabi***, SPE-138591-PP, Abu Dhabi, UAE, November 2010, Abu Dhabi International Petroleum Exhibition & Conference.

Grader, A., Kalam, M. Z., Toelke, J., Mu, Y., Derzhi, N., Baldwin, C., Armbruster, M., Al Dayyani, T., Clark, A., Al Yafei, G. B., and Stenger, B: ***A Comparative study of DRP and laboratory SCAL evaluations of Carbonate cores***, SCA, Halifax, Oct. 2010

Toelke, J., C. Baldwin, Y. Mu, N. Derzhi, Q. Fang. A. Grader, J. Dvorkin: ***Computer simulations of fluid flow in sediment: From images to permeability***, The Leading Edge, (2010) 29, 1040-1048.

AS. Grader, Ingrain, A. B.S. Clark, ADCO, Taha Al-Dayyani, ADCO, and Amos Nur, Ingrain: ***Computations of Porosity and Permeability of Sparic Carbonate Using Multi-scale CT Images***, International Symposium of the Society of Core Analysts, Noordwijk aan Zee, The Netherlands, 27-30 September, 2009

J. Dvorkin, N. Derzhi, Q. Fang, A. Nur, B. Nur, A. Grader, C. Baldwin, H. Tono and E. Diaz: ***From Micro to Reservoir Scale: Permeability from Digital Experiments*****,** The Leading Edge; December 2009; v. 28; no. 12; p. 1446-1452

Summary

INGRAIN
DIGITAL ROCK PHYSICS LAB

- **1 GPUs = 10 powerful Pc's**
- Dollar/FLOP (CHEAP) and WATT/FLOP (Green Computing) very favorable
- No tedious access to supercomputers (data transfer)
- telsa hardware is robust and running very stable.
- L2 cache in Fermi allows for more complicated data structures and irregular access:
  - Nonlocal operations easier
  - mesh refinement

- **Digital rock physics works**
- **Direct solution of two-phase flow in the pore structure**
- **Direct computations – no calibration**
- **Reproducible results**
- **parameter studies**
- **Fast**

Conclusions

INGRAIN
DIGITAL ROCK PHYSICS LAB

- **nVIDIAS GPUs and CUDA helped Ingrain to bring down the cost and time for complex fluid flow simulations in real rocks.**

Abu Dhabi Lab
OpenedJune 2010
10,000 sq feet