


# Simulation of Blood Flow in the Human Brain using the Cell Processor

Jan Götz, Markus Stürmer

University Erlangen-Nuremberg – System Simulation


July 18th 2007 @ ICMMES



- Introduction
- The Cell Processor
- Implementation
- Results
- Outlook

## Aneurysms


- dilatation (local ballooning) of the blood vessel
- localized mostly at larger arteries in soft tissue (e. g. aorta, brain)
- can be diagnosed by modern imaging techniques
- can be treated by clips, coils, etc



<sup>1</sup>from The Toronto Brain Vascular Malfunction Study Group

## Aneurysms

- dilatation (local ballooning) of the blood vessel
- localized mostly at larger arteries in soft tissue (e. g. aorta, brain)
- can be diagnosed by modern imaging techniques
- can be treated by clips, coils, etc



## Stenosis

- abnormal narrowing in a blood vessel

<sup>1</sup>from The Toronto Brain Vascular Malfunction Study Group

## Motivation

- aneurysms are a major public health issue in every developed nation
- the flow situation could be crucial for further treatment of the patient

## Motivation

- aneurysms are a major public health issue in every developed nation
- the flow situation could be crucial for further treatment of the patient

## Goals

- to help in understanding the development of aneurysms
- to support planning of therapy

## Motivation

- aneurysms are a major public health issue in every developed nation
- the flow situation could be crucial for further treatment of the patient

## Goals


- to help in understanding the development of aneurysms
- to support planning of therapy

## Challenges

- current imaging techniques result in data sets of  $512^3$  and more
- long run times on desktop PCs and workstations
- for intra-surgery planing the algorithm should perform quasi real-time

# The Cell Processor

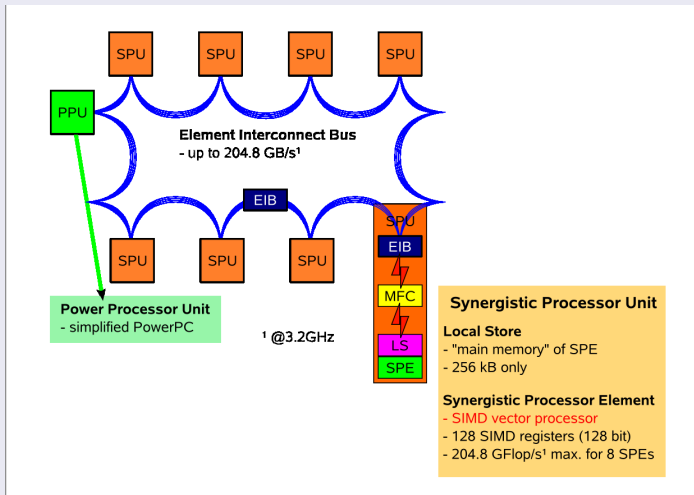
## A short overview of the Cell processor





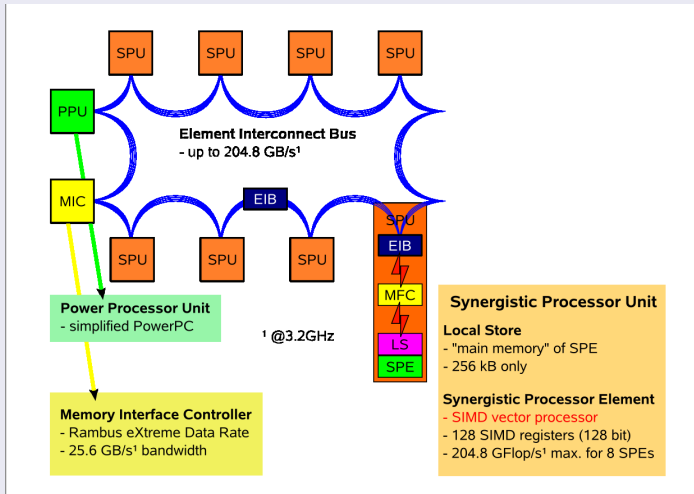
# The Cell Processor

## A short overview of the Cell processor



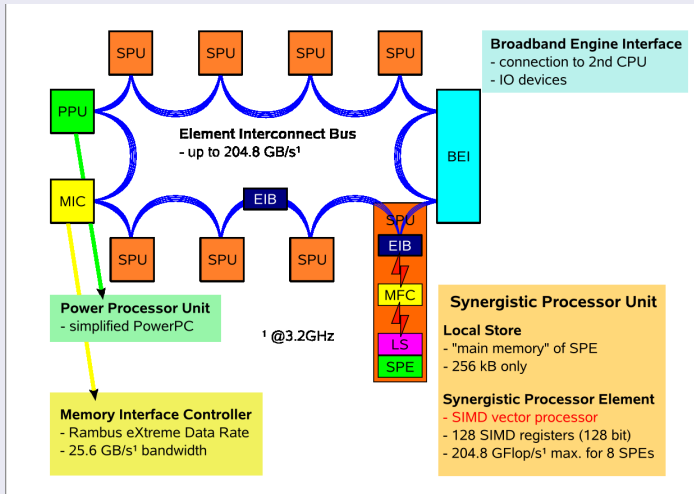
# The Cell Processor

# A short overview of the Cell processor



# The Cell Processor



# A short overview of the Cell processor



# Implementation

## Domain decoupling

Divide whole domain into equally sized patches ( $8 \times 8 \times 8$ ) and only allocate and calculate patches including fluid cells.



## Our Data Structure

- **Patch itself**




- $8 \times 8 \times 8 \times 19$  distribution functions

- **Halo**

- copies of values that need to be exchanged
  - 30 planes and 12 lines
    - streaming becomes quite tricky

- **Remote Halo**


- pointers to the neighbors' copies
- similar data for lattice type information



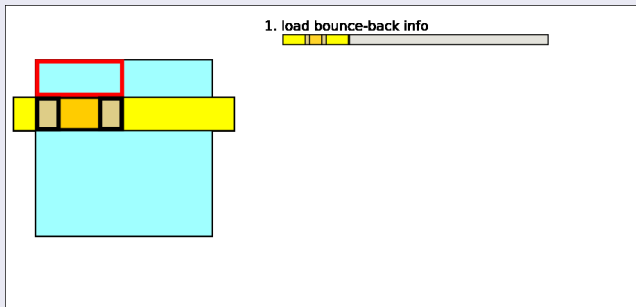
## Updating a Patch on a SPU

- ① fetch...
  - distribution functions
  - LBM-cell type information
  - remote halo pointer structure
- ② fetch...
  - halo planes from neighbors (T, B, N, S, E, W)
  - halo lines from neighbors (TS, NE, BW ...)
- ③ set source / sink
- ④ stream (and bounce) values (including values from halo planes)
- ⑤ correct pressure outflow
- ⑥ calculate collision
- ⑦ prepare new halo structure from calculated values
- ⑧ store patch and halo data

## Example of SIMDized Bounce-Back: Get new north vector

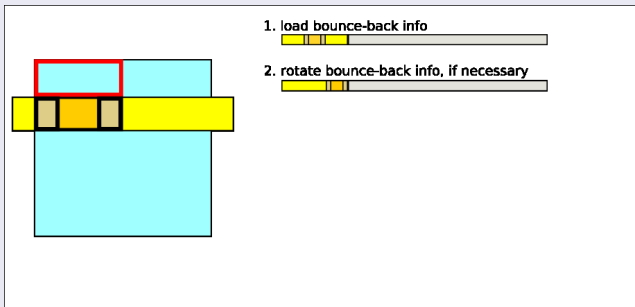


## Example of SIMDized Bounce-Back: Get new north vector

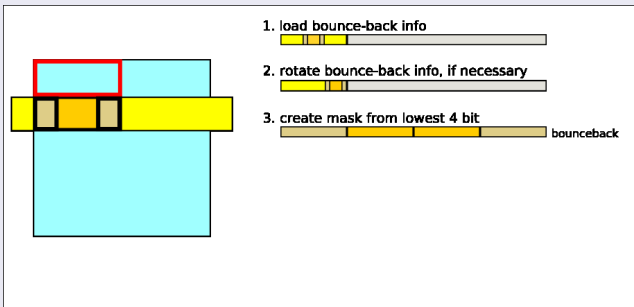




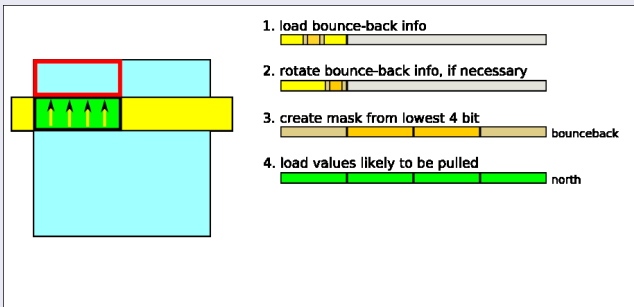
## Example of SIMDized Bounce-Back: Get new north vector



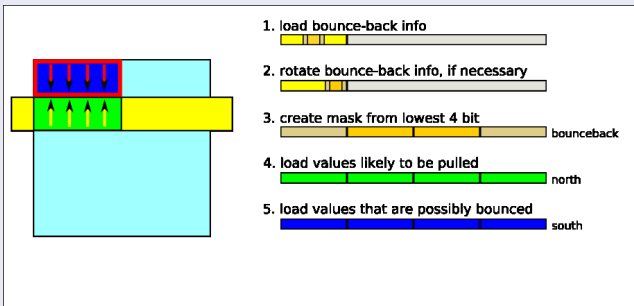
## Example of SIMDized Bounce-Back: Get new north vector



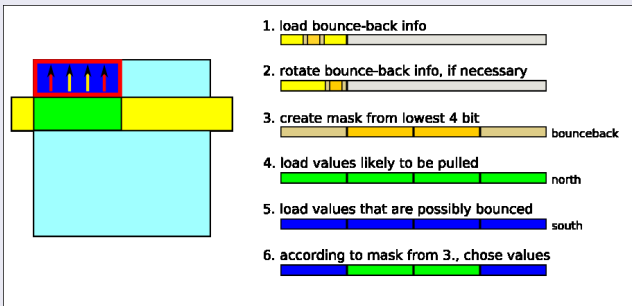
## Example of SIMDized Bounce-Back: Get new north vector



## Example of SIMDized Bounce-Back: Get new north vector

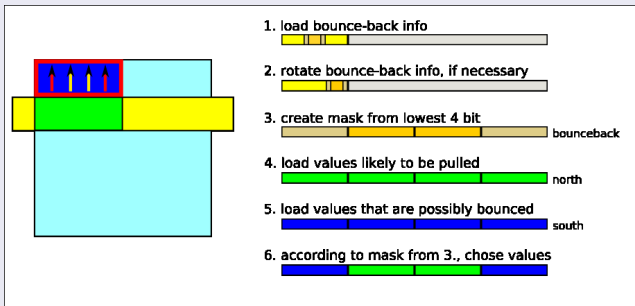


## Example of SIMDized Bounce-Back: Get new north vector



# Implementation

## Example of SIMDized Bounce-Back: Get new north vector



## Remarks on SIMDized Bounce-Back


- streaming must take halo data into account  $\Rightarrow$  complicated
- $18 \times 8 \times 8 \times 8 = 9216$  “decisions” per patch  
 $\Rightarrow$  scalar processing unbearably slow

# Results

Single precision performance and memory for a real vessel geometry with size 250x250x220

	Core 2 Duo(3.0 GHz, both cores) <sup>a</sup>	Cell(3.2 GHz, 8 SPU's)
1 CPU	11.7 MFLUPS	
2 CPU	21.0 MFLUPS	
memory	128 MB	

<sup>a</sup>Measurements on Core 2 Duo (Woodcrest) done by Thomas Zeiser (RRZE) with the code from the Lattice Boltzmann Development Consortium




# Results

Single precision performance and memory for a real vessel geometry with size 250x250x220


	Core 2 Duo(3.0 GHz, both cores) <sup>a</sup>	Cell(3.2 GHz, 8 SPU's)
1 CPU	11.7 MFLUPS	43.4 MFLUPS
2 CPU	21.0 MFLUPS	88.0 MFLUPS
memory	128 MB	106.3 MB

<sup>a</sup>Measurements on Core 2 Duo (Woodcrest) done by Thomas Zeiser (RRZE) with the code from the Lattice Boltzmann Development Consortium






## Single precision performance of Cell vs. Core 2 Duo for real geometry



- Cell (3.2 Ghz, 8 SPU)
- Core 2 Duo (3.0 Ghz, Woodcrest, 2 cores)

# Results

## Single precision performance of Cell vs. Core 2 Duo for real geometry



## Comparison of standard C code to advanced Cell code for a channel

	Standard C Code	Advanced Cell Code
Core 2 Duo	10.2 MFLUPS	
1 SPU	2.0 MFLUPS	38.5 MFLUPS
8 SPU	15.8 MFLUPS	98.1 MFLUPS

## More to do...

- examine influence of rounding mode
- MPI parallelization
- add more difficult model / evaluation

## More to do...

- examine influence of rounding mode
- MPI parallelization
- add more difficult model / evaluation

## Acknowledgements

- thanks to Markus Stürmer for the implementation
- thanks to Thomas Zeiser for performance measurements
- thanks to IBM development center at Böblingen for access to Cell Blade
- thanks to ZAM / FZ Jülich for access to JUICE
- thanks to Professor Dörfler and Doctor Richter for medical advice

Thank you very much for your attention!