Mitglied der Helmholtz-Gemeinschaft

JÜLICH APPLICATION LAB NVIDIA

JÜLICH FORSCHUNGSZENTRUM

From brain research to high-energy physics: GPU-accelerated applications in Jülich

Dirk Pleiter | Jülich Supercomputing Centre (JSC) | SC13

NVIDIA Application Lab at Jülich

Collaboration between JSC and NVIDIA since July 2012

- Enable scientific applications for GPU-based architectures
- Provide support for their optimization
- Investigate performance and scaling

Andrew Adinetz

Work focus

- Application requirements analysis
- Current GPU architecture and CUDA feature analysis
- Parallelization on many GPUs
- Collaboration with performance tools developers
- Training

Jiri Kraus

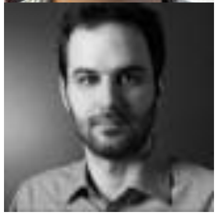

HPC at Jülich Supercomputing Centre

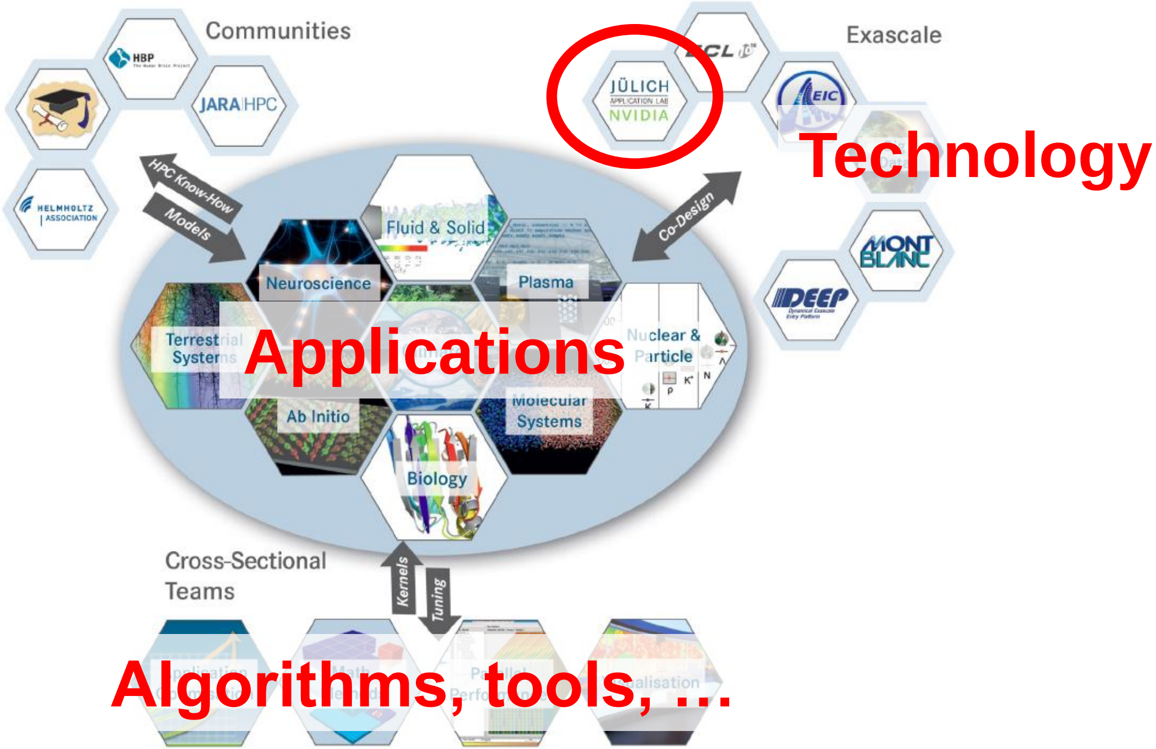

Human Brain Project Application: JuBrain

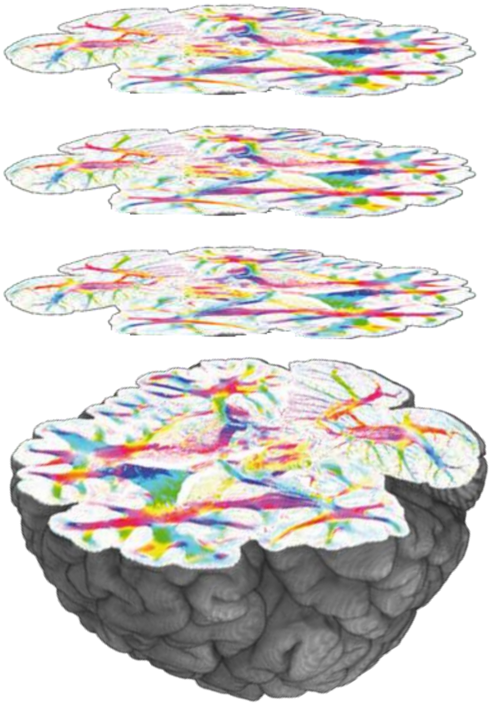

Katrin Amunts, Markus Axer,
Marcel Huysegoms

Human Brain Project

Research goal

Accurate, highly detailed computer model of the human brain

Computational challenge

- Registration of high resolution images
- Algorithm, e.g., rigid registration → 3 parameters
- Computation of metric based on Shannon entropy

$$H = \sum_i p_i \log \frac{1}{p_i}$$

JuBrain Registration Workflow

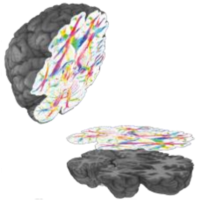

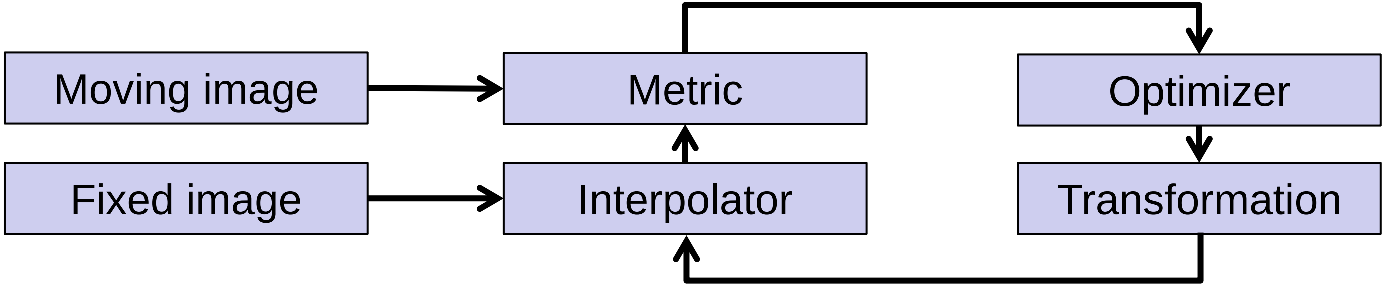

Metric computation → Computing joint histograms for 2 images

```
for(int y = 0; y < fixed_sz_y; y++)
  for(int x = 0; x < fixed_sz_x; x++) {
    int i = bin(fixed[x, y]);
    float x1 = transform_x(x, y);
    float y1 = transform_y(x, y);
    int j = bin(interpolate(moving, x1, y1));
    histogram[i, j]++;  // atomic on GPU
  }
```

L2 atomics performance relevant when computing metric

JuBrain Parallelization Strategies

Simple test bench

- Only rotation

System memory replication

- Device holds local part of fixed image
- Host memory holds full copy of moving image

List update

- Send local fixed image data and moving image coordinates

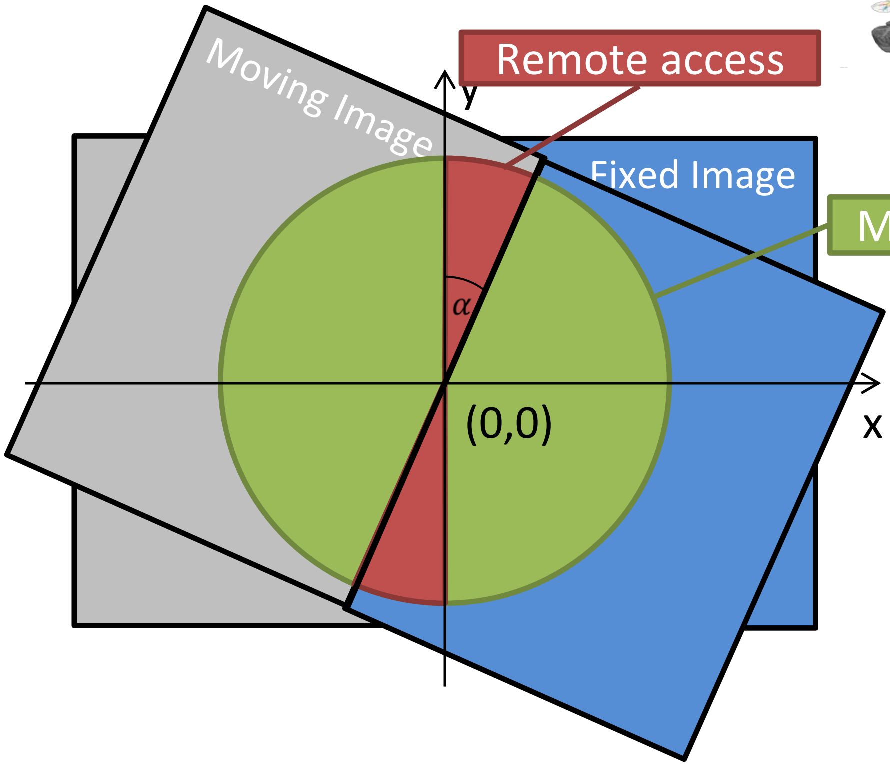

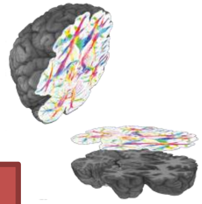

Parallel JuBrain Performance Results

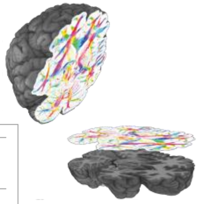

Fermi

- Reasonable scaling for small angles α
- System memory replication faster
- Strong performance degradation for intermediate α ← system memory latency

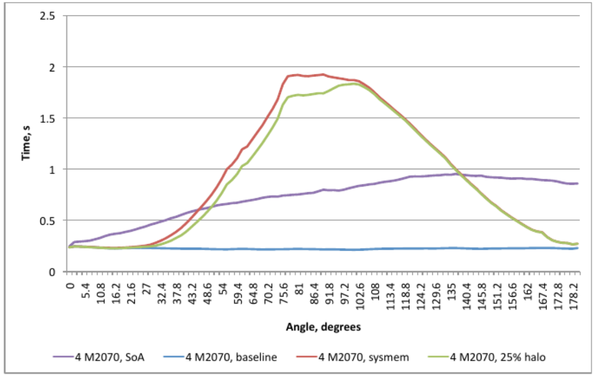

Kepler

- List update strategy faster due to faster L2 atomics

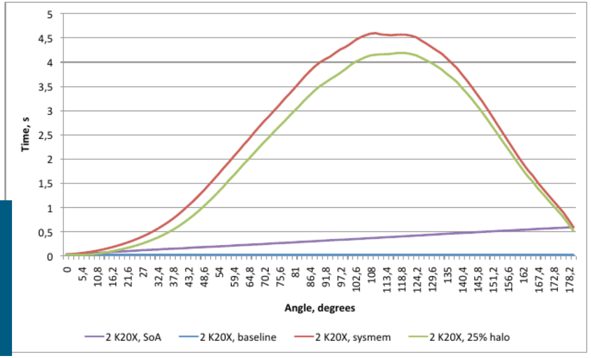

Fine-grained multi-GPU communication potentially tricky

B-CALM: Belgium-California Light Machine

Pierre Wahl

Research goal

- Simulate electromagnetic fields in matter
- Applications
  - Nano-photonics for optical interconnect
  - Optimized photo-voltaic

Finite-difference time-domain (FDTD) method

- 3d grid of E and H fields

Apply method to large systems

- 4000^2x400 grid points → O(250) GBytes

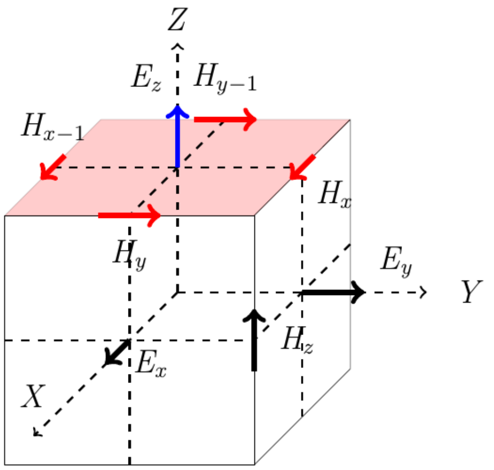

Parallel B-CALM Performance Model

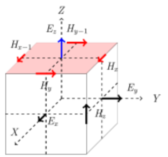

Parallelisation strategies

- 1d domain decomposition z-direction
- Higher dimension decompositions

Simple model ansatz

- Information flow analysis
- Latency-bandwidth model

Comparison model and measurement

- Good agreement for 1d domain decomposition
- No need for higher-dimension decomposition

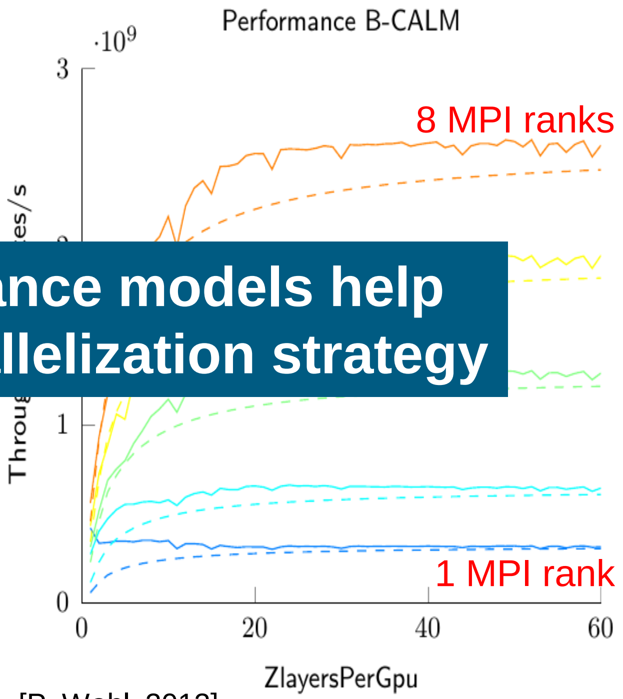

Performance models help fixing parallelization strategy

[P. Wahl, 2013]

GPUMAFIA: Data analysis on GPUs

Sub-space density clustering

- Analysis of high-dimensional data sets
- Find clusters which exist in subsets of dimensions

Applications

- Monte Carlo simulations of protein folding
- Data mining in marketing,
  bio-informatics, medical imaging

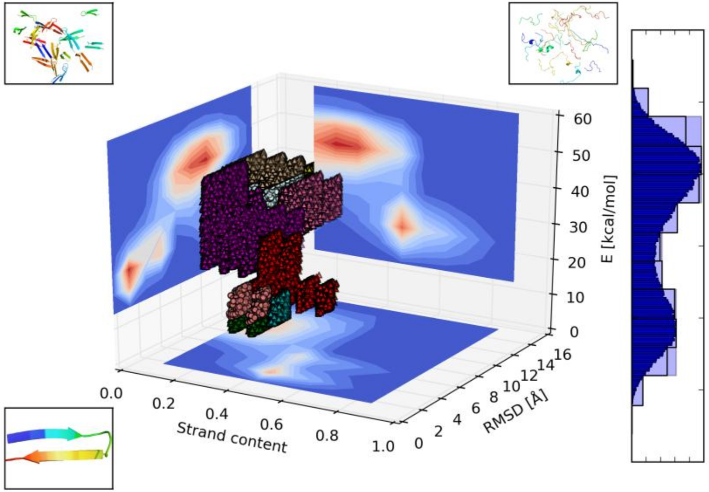

MAFIA = Merging of Adaptive Finite IntervAls

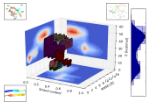

Sub-space clustering

- If a collection of points S is a cluster in a k-dimensional space, then S is also a part of a cluster in any (k-1)-dimensional projection of the space
- Start from constructing histograms in each dimension

Adaptive grid

- Combine bins with similar histogram values

Gradually form higher dimensional clusters

GPUMAFIA Performance Results

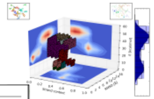

Test setup

- Dual 6-core Xeon
- Single core Xeon + K20x

Synthetic dataset

- 30 dimensions
- 10^5 data points

Observe O(10) speed-up

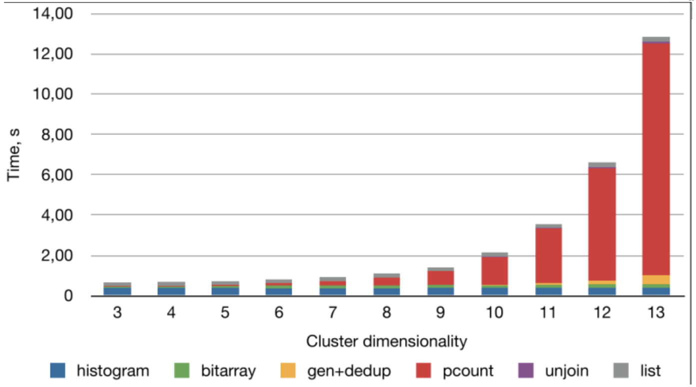

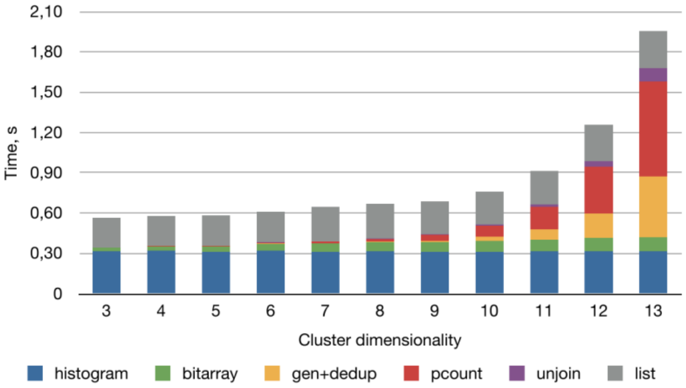

GPUs help getting data analysis to "interactive speed"

PANDA Track Reconstruction

Andreas Herten, Marius Mertens, Tobias Stockmanns et al.

PANDA = Next generation hadron physics experiment

- Part of FAIR accelerator in Darmstadt (Germany)

Scientific goal and requirements

- Triggerless track reconstruction
- Sustain data rate of 20 million events/s → 200 GBytes/s
- Achieve O(1000) times data reduction

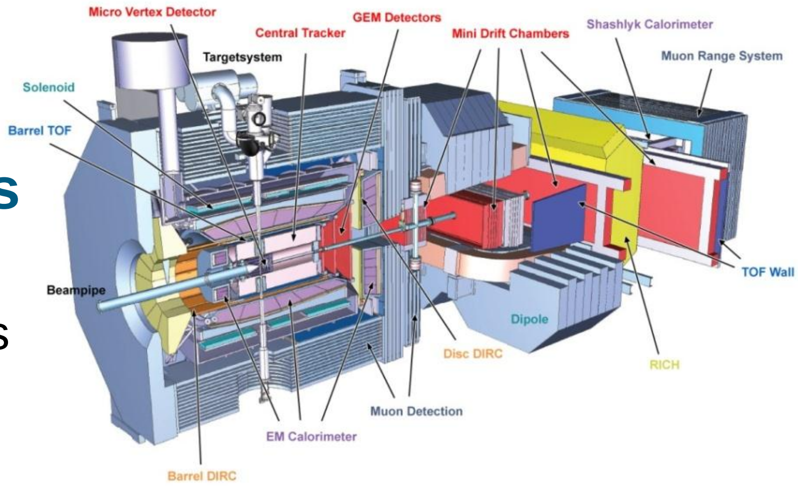

PANDA Track Reconstruction

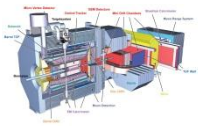

Why using GPUs?

- Easier to program compared to, e.g., FPGAs
- Latencies more predictable than for CPUs

Algorithms

- Hough transformation
- Triplet finder
- Riemann tracker

Close to proof-of-concept for high event-rate processing

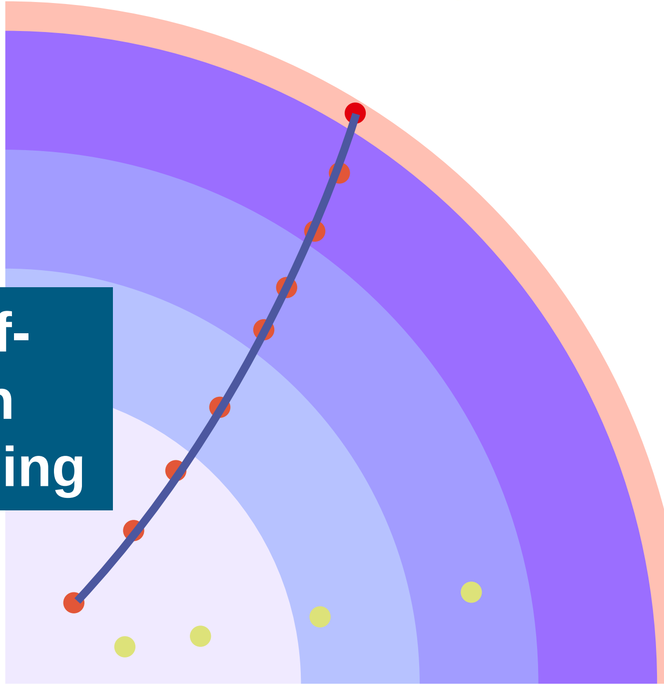

Initial results

- Triplet finder running at rate of <1 μs per hit

Summary

NVIDIA Application Lab at Jülich

- Fruitful model for collaboration

Multi-GPU parallelization

- Required, e.g., due to device memory limitations
- Applications: JuBrain image registration, B-CALM FDTD application

Data-intensive applications on GPUs

- Strongly benefit from improved support of L2 atomics
- Applications: GPUMAFIA clustering, PANDA track recontruction