

GP-GPU-Computing for the cluster analysis of fiber tracts: Replacing a \$15000 high end PC with a \$500 graphics card

C. Ros¹, R. Tandetzky¹, D. Güllmar¹, and J. R. Reichenbach¹

¹Medical Physics Group, Department of Diagnostic and Interventional Radiology I, Jena University Hospital, Jena, Thuringia, Germany

Introduction - Whole brain tractography (WBT), either performed with deterministic or probabilistic approaches, reconstructs a multitude of fiber tracts, where the number of tracts in typical data sets can easily exceed several thousand up to more than a million fiber tracts. Even though fiber tracts are just an abstraction of the underlying microscopic fiber configuration [1, 2], they do capture the inherent complexity of white matter structures such as axonal pathways at a very detailed level [1]. Henceforth, fiber tracts are eminently suitable to analyze neural connectivity or structural changes in the underlying white matter, by employing quantitative tractography-based analysis [3-6]. The bundling of fiber tracts, which is essential for quantitative tract-based analysis of white matter bundles, can easily be performed fully automated and unsupervised with cluster analysis (CA) methods. However, due to the huge amount of raw data that needs to be processed as well as the high computational complexity of CA techniques, the application of CA methods is restricted to relatively small data sets. Otherwise, long processing times on the order of hours or even days must be accepted. Even though the processing time can be dramatically reduced by employing advanced CA methods [7], the computationally most demanding part of the CA – the calculation of similarities between different fiber tracts – remains and limits the application of CA severely.

To facilitate the use of cluster analysis for exploratory data analysis, the quantitative analysis in multi-subject imaging studies and the applicability of CA for large tractography data sets, we focused on parallel computation of fiber tract similarities in this study. The potential of Symmetric-Multi-Processing systems (SMPs) and General Purpose Graphics Processing Units (GP-GPU) in speeding up fiber tract similarity computation was investigated and analyzed. One SMP system as well as three GPU systems were employed for the calculation of the Hausdorff similarity metric [8] and the minimum Euclidean tract distance metric [8].

Materials and Methods - DTI datasets of a healthy volunteer were measured on a clinical 3 T whole body MR-Scanner (Magnetom Tim Trio, Siemens Healthcare, Erlangen, Germany) by employing a conventional twice refocused Echo Planar Imaging (EPI) sequence [9]. A 12 channel phased array matrix head coil was employed and the following parameters were used: TE=113ms, TR=7900ms, $\alpha=90^\circ$, iPAT=2, matrix of 96×96 , 55 slices with a thickness of 2.5mm resulting in a voxel size of $2.5 \times 2.5 \times 2.5 \text{mm}^3$. Five b_0 images without diffusion weighting as well as 70 diffusion weighted images sampled with different gradient directions ($b=1000 \text{s/mm}^2$) were acquired. In-plane interpolation was performed on the MR-scanner, resulting in a nominal voxel size of $1.25 \times 1.25 \times 2.5 \text{mm}^3$. The Diffusion Toolkit [10] was then utilized to perform whole brain fiber tractography. Tracts with lengths less than 30 mm were subsequently removed from the data set. Different SMP- and GPU-systems were used for this study. The SMP systems are composed of various processing cores, which perform computations independently from each other. Contrarily, GPU-systems are hybrid systems, consisting of several independent Multi-Processors (MPs). Every MP is composed of lower clocked cores that share the same instruction decoder. Thus, the cores of one MP are not independent and they are bound to perform the same instructions at a clock cycle. One SMP as well as three GPU-systems were used: • SMP system with 32 cores (AMD Opteron 8350) • GPU system with 64 cores (NVIDIA Quadro FX) • GPU system with 240 cores (NVIDIA Tesla C1060) • GPU system with 448 cores (NVIDIA Tesla C2050). Two distance measures were used for this evaluation – the minimum Euclidean distance (MD) between fiber tracts and the Hausdorff distance (HD). Due to significant differences between CPU and GPU architectures, different implementations for the computation of the similarity measures were used. For the CPU, a straightforward implementation of the MD and HD was employed, as well as more sophisticated methods to determine the HD and an approximated MD (max error < 1mm). Total similarities were partitioned into N sets that were processed in parallel by N CPU cores ($N=1, 2, \dots, 32$). For the GPU, optimized implementations of the HD and MD were used to maximize performance by occupying all cores.

Results - The processing times for the CPU- and GPU-systems, as well as the various computation methods are shown in Fig. 1. For 1 million similarity computations, a significant reduction of the computation time was observed with the GPUs. Even the weakest GPU system was able to outperform the high-end PC with 32 CPU cores. Despite the fact that the C2050 has almost twice the number of CPU cores (and costs twice as much) compared to the C1060, differences between both systems were quite small. By comparing computation times for the C2050 and the single core CPU implementation, a speed up of ~850 was observed for the straightforward and a speed up of ~90 for the sophisticated implementations. Compared to the processing time for the 32 core implementations, the speed up for the C2050 was only ~40 (straightforward) and ~3.5 (sophisticated), respectively. The results for the computation times of the CPU implementations by using different number of cores (shown in Fig. 2), indicate good speed up, especially if only few cores are employed.

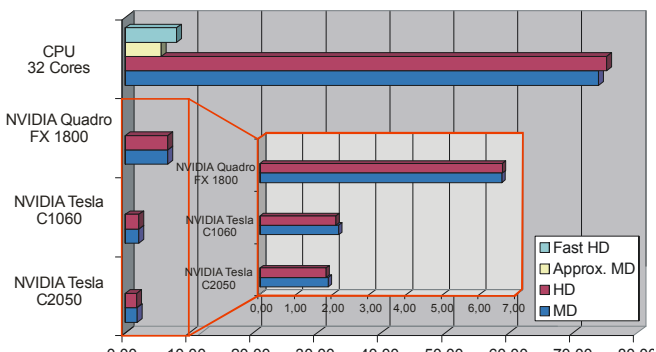


Fig.1: Processing time for 1 million similarity comparisons [s]

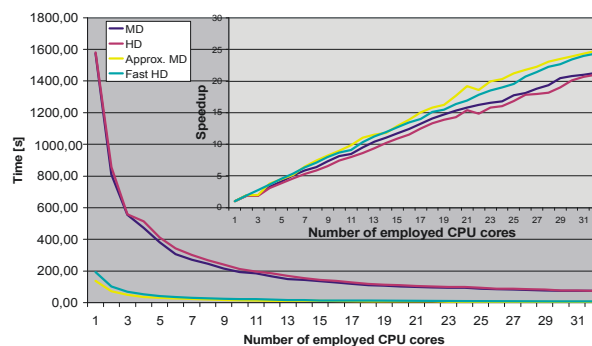


Fig.2 Processing time and speedup for 1 million similarity comparisons with multiple CPU cores

Discussion & Conclusion - Fast computation of fiber tract similarities is crucial for fully automated and unsupervised tractography-based cluster analysis of large data sets and cluster analysis in multi-subject imaging studies. Employing CPU- or GPU-systems with multiple processing cores can reduce the computation times significantly and helps to overcome these obstacles. Even though GPUs offer superior processing capabilities compared to conventional CPU-based systems, the full utilization of the GPU is challenging due to different hardware architectures. To make full use of the GPUs' superior processing capabilities, new fine-grained parallel algorithms have to be developed and old algorithms have to be adapted for the GPU architecture. This, however, is quite challenging because many algorithms are not suitable for parallel, non-independent processing as performed by GPUs. Depending on the specific data that is processed by each GPU core, different execution paths on the hardware can occur, which will lead to sequential execution and thus significant degradation of performance. At this point, smart implementations on CPU-based systems may diminish the advantages of GPUs or even outperform GPUs at all. Nonetheless, GPU computing is not only a promising technique that is able to increase the performance significantly, but also a technique that is more cost efficient than conventional CPU-based systems.

Acknowledgements - This study was supported by the German Federal Ministry of Education and Research (BMBF), project number 01GW0740 and the BMBF Bernstein Group for Computational Neuroscience, project number 01GQ0703. We like to thank Siemens Healthcare, Erlangen, Germany for their support.

References

- [1] Gray, 1973, Anatomy of the human body, 29th edition
- [2] Dauguet et al, 2007, Neuroimage 37, 530–538
- [3] Berman et al, 2005, Neuroimage 27, 862–871
- [4] Wakana et al, 2007, Neuroimage 36, 630–644
- [5] Partridge et al, 2005, J Magn Reson Imaging 22, 467–474
- [6] Voineskos et al, 2009, Neuroimage 45, 370–376
- [7] Ros et al, 2011, Proc Intl Soc Mag Reson Med 19
- [8] Güllmar et al, 2008, Proc Intl Soc Mag Reson Med, 16
- [9] Heid, 2000, Proc Intl Soc Mag Reson Med, 8
- [10] Wang et al, 2007, Proc Intl Soc Mag Reson Med 15, #3720