Available online at www.sciencedirect.com
ScienceDirect
Computerized Medical Imaging and Graphics 32 (2008) 521–530
Computerized Medical Imaging and Graphics
www.elsevier.com/locate/compmedimag

Diffusion tensor fiber tracking on graphics processing units

Adiel Mittmann*, Eros Comunello, Aldo von Wangenheim

Informatics and Statistics Department, Federal University of Santa Catarina, 88040-970 Florianópolis, Brazil

Received 1 November 2007; received in revised form 27 May 2008; accepted 28 May 2008

Abstract

Diffusion tensor magnetic resonance imaging has been successfully applied to the process of fiber tracking, which determines the location of fiber bundles within the human brain. This process, however, can be quite lengthy when run on a regular workstation. We present a means of executing this process by making use of the graphics processing units of computers' video cards, which provide a low-cost parallel execution environment that algorithms like fiber tracking can benefit from. With this method we have achieved performance gains varying from 14 to 40 times on common computers. Because of accuracy issues inherent to current graphics processing units, we define a variation index in order to assess how close the results obtained with our method are to those generated by programs running on the central processing units of computers. This index shows that results produced by our method are acceptable when compared to those of traditional programs.

Keywords: Diffusion tensor imaging; Fiber tracking; Graphics processing units; Fiber tracking results comparison; Performance optimization

1. Introduction

Diffusion tensor magnetic resonance imaging (DT-MRI) is a technique that allows the *in vivo* estimation of a diffusion tensor in every voxel of a volume [1,2]. A diffusion tensor characterizes the amount of water diffusion along any direction in a specific point in space. It can be thought of as an ellipsoid whose shape describes how water diffuses at that point [3].

The DT-MRI process begins with the acquisition of at least seven full magnetic resonance scans. Six of these are composed of diffusion-weighted images that hold diffusion information along six different directions, and one of the baseline images with no diffusion information [4]. By solving a system of six linear equations, a diffusion tensor can be estimated for all voxels. This diffusion tensor has six independent elements, hence the minimum requirement of six diffusion-weighted images.

DT-MRI has been applied with special interest to the human brain, since by using the information provided by the diffusion tensor it is possible to carry on a process known as fiber tracking or tractography, whose aim is to find fiber bundles in the brain's white matter [5–7].

Many algorithms have been proposed for fiber tracking. Examples are the use of tensors as a five-dimensional position orientation space for segmentation [8], stochastic processes [9] and front-evolution algorithms [10,11]. This paper is concerned with a family of algorithms based on the so-called streamline method [5,12], which use the tensor's eigenvectors in order to find fiber tracts. This method is the most traditional one.

Streamline fiber tracking algorithms deal with a large number of tensors, ranging from a few to several hundred thousand, and perform intensive mathematical operations, including interpolation of three-dimensional tensors [7]. A full brain fiber tracking process may become quite lengthy when run on a regular workstation [13], depending on the complexity of the algorithm employed. The duration of such a process can last from a few to several minutes.

This paper shows how the speed of this process can be improved by using graphics processing units (GPUs) [14,15], the processing units of video cards. They were traditionally directly inaccessible to programs running on the central processing unit (CPU) and performed exclusively graphics-related computations, but many of them have been recently extended to allow direct programming. Because programs running on the GPU can do any type of computation, this technology has been termed GPGPU—general-purpose computation on GPUs [14,16].

* Corresponding author. Tel.: +55 48 3721 9516.
E-mail addresses: adiel@inf.ufsc.br (A. Mittmann), eros@telemedicina.ufsc.br (E. Comunello), awangenh@inf.ufsc.br (A. von Wangenheim).

0895-6111/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.compmedimag.2008.05.006

While in theory any problem could be solved by a program running on a GPU, usually a great effort has to be put into adapting the problem to the way GPUs operate. This is especially true for problems unrelated to graphics or image processing. Thus, the main objective of this paper is to show how fiber tracking through the streamline method can be modeled to be executed on a GPU, yielding a performance gain of up to 40 times, depending on the hardware used.

Because GPUs' floating-point operations do not produce exactly the same results as CPUs', a variation index is defined to measure how different two fiber tracking results are, this way allowing the quantification of the difference between results produced by GPUs and those produced by CPUs.

The paper is organized as follows: Section 2 is an overview of streamline fiber tracking algorithms and GPUs; Section 3 details materials used in experiments; Section 4 describes a model for fiber tracking algorithms to be run on GPUs, and also defines an index to measure how much variation exists between two fiber tracking results; Section 5 presents results and discusses them; finally, Section 6 contains our concluding remarks.

2. Background

2.1. *Streamline fiber tracking*

A streamline fiber tracking algorithm generates a series of fiber tracts or trajectories from a set of seed points. Trajectories are determined by repeatedly finding their next point, beginning with a seed point [6,17]. Although algorithms differ in some key features, their overall structure remains the same. The overall structure of such an algorithm is shown by Fig. 1.

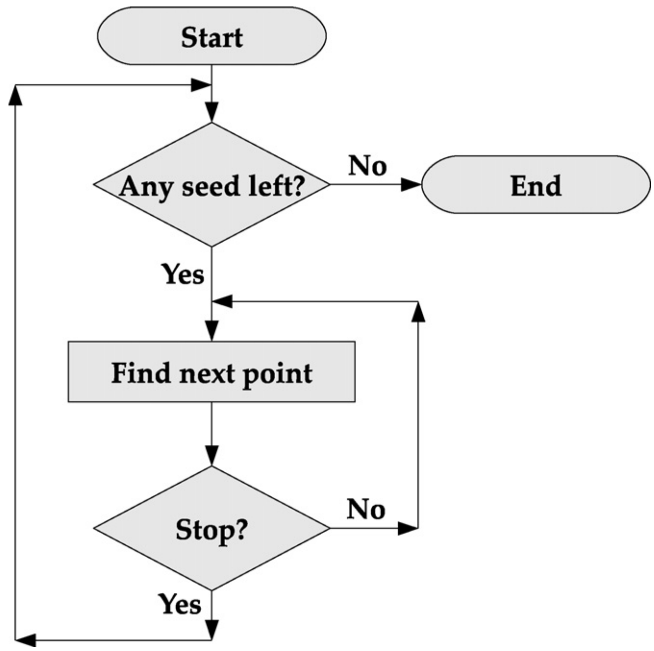

Fig. 1. Overall structure of streamline fiber tracking algorithms. The outer loop iterates through all seeds, and the inner loop, which is the fiber tracking step, finds all points for each trajectory. The two key features of a fiber tracking algorithm are contained within the inner loop: the way it finds the next point of a trajectory and how it decides if a trajectory is finished.

Seed points can be points within a region of interest or the center of all voxels of a volume. For each seed point, a streamline algorithm tries to determine a trajectory, such as the simple one portrayed by Fig. 2. Not all seed points, however, are successful in spawning trajectories; some of them may be eliminated by a validation step because they do not look promising (e.g., they may be located outside of the patient's brain), or the corresponding trajectory may be pruned in a later stage.

In order to decide what is the next point in a trajectory, the algorithm estimates a tensor at the trajectory's previous point. This tensor has information on water diffusion at that point, and the main direction of diffusion is generally used to find the next point for the trajectory. The precise way in which an algorithm determines the next point is a key feature in the algorithm [18]. The process of determining the next point of a trajectory is called the fiber tracking step.

A seed point is used as the initial point for a trajectory, but it must be followed in two opposite directions because tensors provide information about diffusion direction, but not orientation. Algorithms therefore usually repeat the inner loop of Fig. 1 for both directions before starting on a new seed.

Another issue is how a tensor is estimated at arbitrary points, given the discrete nature of magnetic resonance imaging data. Since these data are noisy, approximation (instead of interpolation) is usually desired. This issue is not strictly a part of streamline fiber tracking algorithms, but rather an important dependence. Approximation strategies are the subject of a great many articles [19–22].

The second key feature in which algorithms differ is how they decide when a trajectory is finished, that is, what stop criteria for trajectories are used by that algorithm. Common stop criteria include fractional anisotropy (FA, which indicates how directed diffusion is) and maximum trajectory length. After a trajectory is finished, it may be discarded if it is considered too short.

The final product of a streamline fiber tracking algorithm is a set of trajectories. This set describes the way fiber bundles are organized in the patient's brain, and the quality of the results varies greatly according to the algorithm and the approximation scheme adopted.

2.2. *Graphics processing units*

2.2.1. *Computation on graphics processing units*

A graphics processing unit is the graphics rendering device of video cards. GPUs are made up of microprocessors, just as CPUs, but they are tailored for graphics processing, allowing

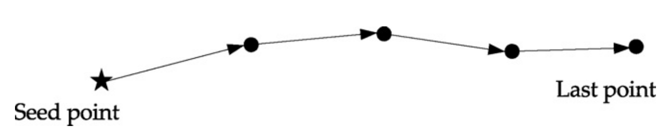

Fig. 2. A fiber tracking trajectory. A trajectory can be seen as a set of interconnected lines. A set of trajectories makes up the result of a fiber tracking algorithm. Real results from an algorithm can be seen in Fig. 10 in Section 5.