The Visual Computer manuscript No.
(will be inserted by the editor)

A GPU framework for parallel segmentation of volumetric images using discrete deformable models

Jérôme Schmid[1], José A. Iglesias Guitián[2], Enrico Gobbetti[2] and Nadia Magnenat-Thalmann[1]

Abstract Although research in image segmentation has been very active during the last decades, it is still a very challenging problem. A lot of difficulties may arise related to, for example, the different image modalities, noise and artifacts of source images, or the shape and appearance variability of the structures to segment. Motivated by problems of image segmentation in the medical field, we present in this paper a GPU framework based on explicit discrete deformable models, implemented over the NVidia CUDA architecture, aimed for the segmentation of volumetric images. The framework supports the segmentation in parallel of different volumetric structures as well as interaction during the segmentation process and real-time visualization of the intermediate results. Promising results in terms of accuracy and speed on a real segmentation experiment have demonstrated the usability of the system.

Keywords Simulation and Modeling, GPU Programming, Segmentation

1 Introduction and related work

Medical image segmentation is nowadays at the core of medical image analysis, is also present in computer vision applications and attracts the interest of the Computer Graphics community. Although research has been very active these last decades, segmentation is still a very challenging problem. The design of efficient segmentation methods can expedite tedious parameter tuning and reduce the limitations of segmentation methods as interactive control is available [5].

[1]MIRALab, University of Geneva, Battelle, Carouge, Switzerland
[2]CRS4 Visual Computing Group, POLARIS Ed. 1, 09010 Pula, Italy

High-level segmentation approaches, such as deformable models, were ported to GPU architectures by considering at first implicit deformable models. Level-sets approaches [4] became particularly popular in the GPU-segmentation community as significant speed-ups and interactive rendering were made available. Geodesic active contours, which are a combination of traditional active contours (snakes) [3] and level-sets evolution, were efficiently implemented in GPU. Nevertheless, little work has been made in implementing explicit discrete deformable models in GPU for segmentation purposes. Methods for implementing active contours based on gradient flow have been proposed [2], but they were limited to the case of 2D images. On the other hand, many works exploited physically-based volumetric deformable models in GPU in other application domains, such as spring mass systems, cloth simulation, volumetric mesh deformation or Finite Element Modeling (FEM).

We present in our paper a GPU framework, implemented using the NVidia's Compute Unified Device Architecture (CUDA), aimed for the segmentation of *volumetric* images based on *discrete physically-based deformable models*. The framework exploits parallelism and performs completely in the GPU being capable of managing real-time interactive segmentation of multiple structures.

2 Segmentation approach

Our segmentation approach first defines generic representations of the anatomical structures to segment and then deform them in a fast way to efficiently capture the patient specific anatomy.

Our segmentation approach is based on physically-based deformable models [6]. The principle is to consider meshes vertices as a set of lumped mass particles with position and velocity subjected to internal and external forces. The concepts are hence similar to any deformable models-based simulation with the particularity that images drive model deformation for a segmentation purpose. Segmentation starts from a topologically identical mesh to the final structure to segment without the necessity of dynamically create or delete vertices.

Our deformable models are represented as a 2-simplex meshes [1]. We use internal forces to regulate the segmentation and external forces to drive it towards the correct result. Internal forces ensure that the model evolution is perturbed as less as possible by image artifacts or possible numerical instabilities. Assumptions are thus made on the model smoothness and shape. External forces are based on the minimization of image-based energies. An image based force is built to attract the vertex towards the optimal target position with the lowest image energy. The evolution of the model is based on the resolution of a discrete differential equations system, which is the result of the Newtonian law of motion applied to the particle system. Given the forces and the particle state, the numerical integration yields a new state of the particle. Various approaches are available for integration (e.g, Explicit/Implicit Euler) depending on stability, accuracy and technical implementation constraints. Our GPU

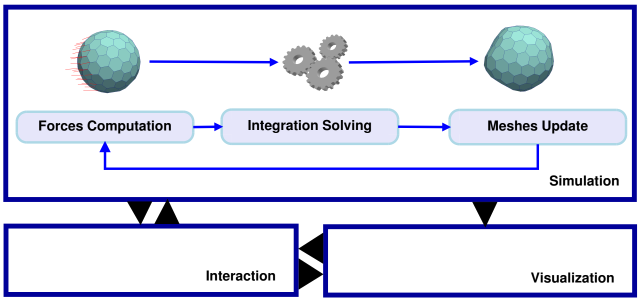

Fig. 1 GPU segmentation: In a simulation step, forces are computed for each particle, whose state is updated in the numerical integration, before updating the meshes parameters. In parallel and asynchronously, meshes and image data are rendered while user can interact with the segmentation.

segmentation framework is implemented on top of NVidia CUDA, the different actions of a simulation step being depicted in Fig. 1. The framework delivers interactive performance and real-time rendering, while providing control to the user. We have designed a simple data access layer to the simplex-mesh data, encoding meshes vertices for a parallel processing within the GPU. The volumetric image information is stored in a raw uncompressed format by using 3D textures taking advantage of spatial locality to provide a fast access for read operations. Different kernels are executed for updating the meshes, computing the different forces and solving the differential equations system driving the deformation of the models (Fig. 1).

3 Experimental results

Our framework was evaluated in the segmentation of hip joint bones from 28 clinical MRI datasets. We compared the accuracy and speed-up of our GPU-framework against a state-of-the-art CPU-based implementation of discrete deformable models [6]. The GPU approach has a similar accuracy and is consistently about $25-70\times$ faster than the CPU version to execute a single time step. In all cases, the time taken by a single GPU iteration is perfectly matching with interactivity constraints. Update frequencies for the GPU are about $47-166$Hz, thus easily supporting the at least 10Hz of refresh rates required for interactivity. On the other hand, the CPU version is unable to fully support interactivity, since update frequencies are about $0.6-6.7$Hz. Having full support to interactivity opens the door to a novel segmentation approach, in which the user is fully able to interact with the segmentation loop.

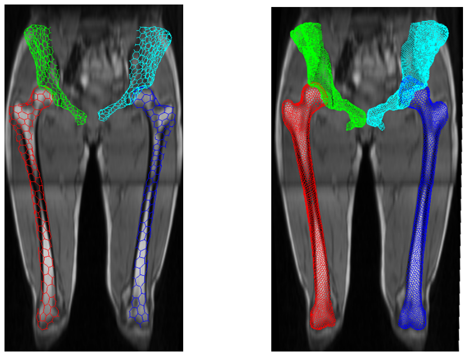

(a) (b)

Fig. 2 GPU-based MRI bones segmentation examples: In a), a coarse mesh is initialized at the beginning of the segmentation, and in b) the final result is shown with meshes at their higher resolution

Acknowledgments: This work is supported by the EU Marie Curie 3D Anatomical Human project (MRTN-CT-2006-035763).

References

1. Delingette H (1999) General object reconstruction based on simplex meshes. Int J Comput Vis 32(2):111–146
2. He Z, Kuester F (2006) GPU-Based Active Contour Segmentation Using Gradient Vector Flow. Advances in Visual Computing 4291:191–201
3. Kass M, Witkin A, Terzopoulos D (1987) Snakes: Active contour models. Int J Comput Vis 1(4):321–331
4. Lefohn AE, Cates JE, Whitaker RT (2003) Interactive, gpu-based level sets for 3d segmentation. In: Proc. MICCAI 2003, pp 564–572
5. Olabarriaga S, Smeulders A (2001) Interaction in the segmentation of medical images: A survey. Med Image Anal 5(2):127–142
6. Schmid J, Magnenat-Thalmann N (2008) Mri bone segmentation using deformable models and shape priors. In: Proc. MICCAI'08, Part I, vol 5241, pp 119–126