Phase Based Volume Registration on the GPU with Application to Quantitative MRI

Anders Eklund* [†], Mats Andersson* [†], Marcel Warntjes[‡] [†], Hans Knutsson* [†]

*Division of Medical Informatics, Department of Biomedical Engineering

[‡]Division of Clinical Physiology, Department of Medicine and Health

[†]Center for Medical Image Science and Visualization (CMIV)

Linköping University, Sweden

Abstract—We present a method for fast phase based registration of volume data for medical applications. As the number of different modalities within medical imaging increases, it becomes more and more important with registration that works for a mixture of modalities. For these applications the phase based registration approach has proven to be superior. Today there seem to be two kinds of groups that work with medical image registration, one that works with refining of the registration algorithms and one that works with implementation of more simple algorithms on graphic cards for speeding up the algorithms. We put the work from these groups together and get the best from both worlds. We achieve a speedup of 10-30 compared to our CPU implementation, which makes fast phase based registration possible for large medical volumes.

I. INTRODUCTION

Image registration is needed in many cases and the goal is to transform a target image to a reference image such that the match is as good as possible. Common clinical application is to easier compare medical images from different modalities, such as MRI and CT, or to compensate for movement between or during scanning sessions. Since the obtained volumes can differ significantly in intensity, especially between different modalities, the local phase, from for example quadrature filters, is better to use since it is invariant to a change in intensity. Phase based image registration has for example been done by Hemmendorff et al. [1] and Mellor et al. [2] but none of them have implemented their algorithm on the graphics processing unit (GPU). GPU based image registration has been done by Bui et al. [3] and Ozcelik et al. [4] but none of them use the phase based approach. Wong et. al. [5] writes about fast phase based volume registration, but does not state any execution times in the article. Pauwels et al. [6] have implemented phase based optical flow on the GPU using Gabor filters. Their implementation is however for 2D and does not perform any registration, since the goal is to calculate motion vectors from video frames. They calculate a motion vector for each pixel separately, while we solve an equation system that is setup globally, to achieve a parameter vector that best describes the transformation between the volumes. We then calculate a motion vector for each voxel from this parameter vector and achieve a movement field that varies smoothly. In this paper we take advantage of the phase based approach to image registration in 3D and use the parallel computing power of graphic cards at the same time.

A. Graphic cards and CUDA

CUDA, Compute Unified Device Architecture, is a parallel computing architecture developed by Nvidia. It enables the user to take advantage of the massive parallel computing power of the GPU. In CUDA, kernels (functions) are launched from the host (the CPU) but is executed on the device (the GPU). Each kernel is launched by a number of blocks, the grid, and a number of threads per block. In cases where we want to perform the same calculations for each pixel or voxel, the computational time can be reduced significanlty since the graphics card performs calculations in parallel.

II. METHOD

A. Quadrature filters and local phase

A quadrature filter is a complex valued filter for combined edge and line detection. The real part of the filter, which is even, detects planes and the imaginary part, which is odd, detects 3D edges. The magnitude of the complex filter response tells us the phase invariant signal intensity and the phase tells us if there is an 3D edge or a plane. We use log-normal quadrature filters Q, which in the Fourier domain are expressed as spherical separable functions with radial function R and directional function D as function of frequency u.

$$Q_k(\boldsymbol{u}) = R(||\boldsymbol{u}||)D_k(\boldsymbol{u}) \tag{1}$$

$$R(||\boldsymbol{u}||) = e^{C \ln^2 \left(\frac{||(\boldsymbol{u})||}{\boldsymbol{u}_0}\right)} \quad C = \frac{-4}{B^2 \ln(2)}$$
(2)

Since the phase conception is valid only if we define a direction of it, we construct quadrature filters with different direction. The directions are defined such that

$$D_k(\boldsymbol{u}) = \begin{cases} (\boldsymbol{u}^T \hat{n}_k)^2 & \boldsymbol{u}^T \hat{n}_k > 0\\ 0 & \text{otherwise} \end{cases}$$
(3)

where \hat{n}_k is the directional vector for filter k. The complex filter response q is an estimate of a bandpass filtered version of the analytical signal

$$q = A \cdot (\cos(\varphi) + i \cdot \sin(\varphi)) = A \cdot e^{i\varphi}$$
(4)

with magnitude A and phase φ . We use one filter in the x-direction, one in the y-direction and one in the z-direction. The filters we use have a centre frequency $u_o =$

 $\frac{\pi}{3}$, a bandwidth B = 1.7 octaves and a spatial size of 9 x 9 x 9 voxels. To obtain filters with spatial locality and desired frequency response, advanced filter design is necessary [7].

B. Registration algorithm

Our registration algorithm is based on 2 assumptions

I. The motion can locally be described as a movement Δx .

$$I(\boldsymbol{x},t) = I(\boldsymbol{x} + \Delta \boldsymbol{x}, t+1)$$
(5)

II. The image can locally be described as a leaning plane

$$I(\boldsymbol{x} + \boldsymbol{v}(\boldsymbol{x}), t+1) = I(\boldsymbol{x}, t) + \nabla \boldsymbol{I}^T \boldsymbol{v} - \Delta I \quad (6)$$

where

$$\Delta I = I(\boldsymbol{x}, t) - I(\boldsymbol{x}, t+1) \text{ and } \nabla \boldsymbol{I} = [\nabla_x I, \nabla_y I]^T.$$

The first assumption says that the intensity does not change between the two images. The second assumption says that the image locally can be described with a first order Taylor expansion. Combining these assumptions gives us the classical optical flow equation

$$\nabla I^T \boldsymbol{v} - \Delta I = 0 \tag{7}$$

The problem with the normal optical flow is that the assumptions that are needed are not met by many images. The phase of the filter response from quadrature filters is better suited for the assumptions, since it is invariant to a change in intensity and varies more smoothly. We then get the optical flow equation of the phase φ

$$\nabla \varphi^T \boldsymbol{v} - \Delta \varphi = 0 \tag{8}$$

where

$$\Delta \varphi = \varphi_1 - \varphi_2 = \arg(q_1 q_2^*). \tag{9}$$

 q_1 is the filter response from the target volume and q_2 is the filter response from the reference volume and * denotes complex conjugation. The movement field is modelled with the help of a 12-dimensional parametervector according to

$$\boldsymbol{p} = [p_1, p_2, p_3, p_4, p_5, p_6, p_7, p_8, p_9, p_{10}, p_{11}, p_{12}]^T$$

$$v(\boldsymbol{x}) = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix} + \begin{bmatrix} p_4 & p_5 & p_6 \\ p_7 & p_8 & p_9 \\ p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = (10)$$
$$= \underbrace{\begin{bmatrix} 1 & 0 & 0 & x & y & z & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & x & y & z & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & x & y & z \end{bmatrix}}_{B} \boldsymbol{p}$$

i.e. the first three parameters is the translation and the last 9 parameters make up the transformation matrix.

The phase gradients can be estimated with the following expression

$$\begin{bmatrix} \nabla_x \varphi \\ \nabla_y \varphi \\ \nabla_z \varphi \end{bmatrix} = \begin{bmatrix} \arg[q_{1x+}q_{1c}^* + q_{1c}q_{1x-}^* + q_{2x+}q_{2c}^* + q_{2c}q_{2x-}^*] \\ \arg[q_{1y+}q_{1c}^* + q_{1c}q_{1y-}^* + q_{2y+}q_{2c}^* + q_{2c}q_{2y-}^*] \\ \arg[q_{1z+}q_{1c}^* + q_{1c}q_{1z-}^* + q_{2z+}q_{2c}^* + q_{2c}q_{2z-}^*] \\ (11)$$

where

$$q_{x+} = q(x+1, y, z), q_{x-} = q(x-1, y, z)$$
$$q_{y+} = q(x, y+1, z), q_{y-} = q(x, y-1, z)$$
$$q_{z+} = q(x, y, z+1), q_{z-} = q(x, y, z-1)$$

 $q_c = q(x, y, z)$

For each voxel and each filter, we also use a certainty that is calculated according to

$$c = \sqrt{|q_1 q_2|} \cos\left(\frac{\Delta\varphi}{2}\right)^2 \tag{12}$$

This certainty measure requires that we have a high magnitude both for the filter response from the reference volume and for the filter response from the target volume, and that the estimated phase does not differ too much.

If we use our model $\boldsymbol{v}(\boldsymbol{x}) = B(\boldsymbol{x})\boldsymbol{p}$ and minimize the least square error

$$\epsilon^{2} = \sum_{k} \sum_{i} c_{ik} (\nabla \varphi_{k}(\boldsymbol{x}_{i})^{T} B(\boldsymbol{x}_{i}) \boldsymbol{p} - \Delta \varphi_{k}(\boldsymbol{x}_{i}))^{2} \quad (13)$$

by setting $\frac{\partial \epsilon^2}{\partial p} = 0$, we get the following equation system

$$\underbrace{\sum_{k} \sum_{i} c_{ik} B_{i}^{T} \nabla \varphi_{ik} \nabla \varphi_{ik}^{T} B_{i}}_{A} \boldsymbol{p} = \underbrace{\sum_{k} \sum_{i} c_{ik} B_{i}^{T} \nabla \varphi_{ik} \Delta \varphi_{ik}}_{\boldsymbol{h}} \underbrace{\boldsymbol{p}}_{\boldsymbol{h}} \underbrace{\boldsymbol{p}}_{\boldsymbol{$$

with the solution

$$\boldsymbol{p} = A^{-1}\boldsymbol{h} \tag{15}$$

Note that the equation system is easy to solve, 12 linear equations, but the cumbersome part is to sum over all voxels i and all filters k. By minimizing a L_2 norm we can *calculate* the parameters that give the best solution. The most common approach is otherwise to maximize a similarity measure by *searching* for the best parameters, using some optimization algorithm. The complete algorithm uses the following steps in each iteration

- Filtering with 3 quadrature filters, that are complex valued in the spatial domain and not cartesian separable.
- Calculating phase differences, phase gradients and certainties for each filter and for each voxel, according to equations 9, 11 and 12 respectively.
- Setting up the equation system, i.e. calculating the A-matrix and the h-vector and solving the equation system to get the parameter vector.
- Calculating a motion vector for each voxel, according to equation 10, and using interpolation to get the value at that position from the modified volume.

To avoid the lowpass filtering effect from repeated interpolation in each iteration we accumulate the parameter vector in each iteration, i.e. $p_{total} = p_{total} + A^{-1}h$, and always interpolate from the original target volume.

III. CUDA IMPLEMENTATION

A. Filtering

The filtering was implemented as multiplication in the frequency domain. Filtering in the frequency domain is faster, than convolution in the spatial domain, if we have filters with many coefficients (729 complex valued in our case) and especially for filtering in 3D and 4D. The volume is first transformed with a 3D FFT, then the transformed volume is multiplied with the transformed filters and finally the filter responses are transformed with a 3D IFFT. We use a CUDA kernel for the complex multiplications between the filters and the volume. Since there is no fftshift function in the CUFFT library, in order to move the origin to the middle of the filter response volume from the corners, we instead perform a change of coordinate system.

B. Phase differences, phase gradients and certainties

Since the calculation of the phase differences, phase gradients and certainties are exactly the same for each voxel, it is ideal for parallel computing on graphic cards. We use one CUDA kernel for calculating the phase differences and certainties and one kernel for calculating the phase gradients.

C. Setting up the equation system

The A-matrix in the equation system is a 12 x 12 matrix and the h-vector is a 12 x 1 vector. There are, however, only 30 of the 144 elements in the A-matrix that need to be calculated, since the rest is zero or can be obtained from the fact that the A-matrix is symmetric, i.e. A(i, j) = A(j, i). To setup the equation system in the CPU implementation we simply sum over all filters, voxels and non zero parameters. In the GPU implementation we take advantage of the high number of simultaneous threads by summing at many positions at the same time. The first CUDA kernel calculates the A-matrix and h-vector for each parameter and at the same time sums over x, for each position (y,z). The second CUDA kernel sums over z and creates the final A-matrix and h-vector.

D. Solving the equation system

Since the equation system is very small, we send the Amatrix and the h-vector to the host and solve the equation system there. Only 12 floats need to be sent for the hvector and 30 floats for the A-matrix.

E. Motion vectors and interpolation

A big advantage with graphic cards, compared to CPU's, is that they have hardware support for interpolation, at least for linear interpolation, but cubic interpolation is also very fast. We use a 3D texture for fast interpolation and use a CUDA kernel that first calculates the motion vector, from the voxel's position (x,y,z) and the parameter vector p, and then simply reads the value from the target volume at that position using a texture call.

IV. RESULTS

The CPU implementation ran on an Intel Core 2 Quad with 4 cores at 2.83 GHz and 4 GB of memory. Our implementation ran on one of the four processor cores but all the execution times for the CPU implementation have been divided by 4 to make the comparison fair. The graphics card used was a Nvidia GTX 285 with 240 stream processors and 1 GB of memory. Neither the processor or the graphics card was overclocked. As test volumes we used one synthetic test volume with a 3D cross with the resolution 128 x 128 x 128 voxels, and a MRI volume with the resolution 240 x 240 x 140 voxels, to which we applied translations and rotations. In the following sections we will present the computational times for the different parts in the algorithm, for the CPU implementation and for the GPU implementation, and the resulting speedups. We will also present the total time for 10 iterations of the algorithm. 10 iterations, on the original scale, is sufficient to for example compensate for 1-5 pixels of translation and 1-5 degrees of rotation in each dimension.

A. Filtering

Volume size	CPU	GPU	Speedup
128 x 128 x 128	0.079 s	6.8 ms	11.6
240 x 240 x 140	0.31 s	187.7 ms	1.7

B. Phase differences, certainties and phase gradients

Volume size	CPU	GPU	Speedup
128 x 128 x 128	0.60 s	7.8 ms	76.9
240 x 240 x 140	1.61 s	25.7 ms	62.6

C. Setting up the equation system

Volume size	CPU	GPU	Speedup
128 x 128 x 128	0.1 s	10 ms	10
240 x 240 x 140	0.33 s	50.2 ms	6.6

D. Solving the equation system

On average it takes 1 ms to send the data to the host and to solve the equation system there.

E. Motion vectors and interpolation

Volume size	CPU	GPU	Speedup
128 x 128 x 128	0.17 s	0.40 ms	425
240 x 240 x 140	0.30 s	1.44 ms	208.3

F. The whole algorithm, 10 iterations, including transfers to and from graphics card

Volume size	CPU	GPU	Speedup
128 x 128 x 128	9.48 s	0.31 s	30.6
240 x 240 x 140	29.80 s	2.97 s	10

G. fMRI-volumes

We also tested the implementation on a set of fMRIvolumes. In fMRI, functional magnetic resonance imaging, for example one volume per second is aquired to estimate a level of brain activity in each voxel. Since the subject can move the head during the acquisition, image registration is important. Each volume in the dataset had a resolution of $80 \times 80 \times 20$ voxels and the dataset contained 324 volumes. 3 iterations per volume was sufficient and the registration took 5.1 s, or 0.0157 s per volume, compared to 0.1 s for a volume of size 64 x 64 x 30, for the fastest algorithm in the comparison by Oakes et al. [8]. This algorithm does not use any filters at all.

H. Quantitative MRI

Quantitative magnetic resonance imaging (QMRI) has the major advantage that it handles absolute measurements of physical parameters. Parameters such as relaxation rates, R_1 and R_2 , and proton density (PD) are independent of MR scanner settings and imperfections and are hence directly representative of the underlying tissue characteristics.

When performing QMR a number of volumes are collected from the MR scanner by using different settings. From these volumes it is possible to estimate the necessary physical parameters. To get a good estimate of the parameters, this however requires that the volumes are perfectly aligned. The collected volumes, shown in Figure 1, differ significantly in intensity and are therefore well suited for phase based registration. The four volumes have a resolution of $256 \times 256 \times 30$ voxels. QMR is further described by Warntjes et al. [9].

We compared our registration algorithm to the statistical parametric mapping (SPM) software by Friston et al. [10]. SPM needed about 2 minutes to register the volumes and almost failed to register volume 4 to volume 1, due to the large rotation and the large intensity difference. Our algorithm did the registration in 8 seconds.

V. DISCUSSION

We have presented a method for fast phase based volume registration. Our implementation shows a speedup between 10 - 30, compared to our CPU implementation, for 10 iterations of the whole algorithm.

The filtering is significantly faster on the GPU. The problem with filtering in the frequency domain is however the limited memory size on the graphics card, since the filters require to be stored as volumes that are as big as the reference volume and target volume.

Calculation of the phase differences, phase gradients and certainties is very fast since it suits perfectly for parallelization. The setup of the equation system is significantly faster on the GPU since we can sum at many positions at the same time. As expected, there is a huge performance gain by using the hardware support for trilinear interpolation when updating the compensated volume in each iteration.



Fig. 1. A slice from each of the four volumes that has to be registered to each other in order for the brain tissue quantification to work properly. This is a challenging registration problem since the volumes differ significantly in intensity. The most common approach is to maximize the mutual information between the volumes, with the help of some optimization algorithm. This may however require hundreds of iterations before a good match is found. We instead use the local phase from quadrature filters, since it is invariant to the intensity, and calculate what the translation and rotation is, instead of searching for the best parameters.

ACKNOWLEDGMENT

This work was supported by the Strategic Research Center MOVIII, funded by the Swedish Foundation for Strategic Research, SSF.

REFERENCES

- M. Hemmendorff, M. Andersson, T. Kronander, and H. Knutsson, "Phase-based multidimensional volume registration," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1536–1543, 2002.
- [2] M. Mellor and M. Brady, "Non-rigid multimodal image registration using local phase," in *Lecture notes in computer science*, *Medical Image Computing and Computer Assisted Intervention* (*MICCAI*), vol. 3216, 2004, pp. 789–796.
- [3] P. Bui and J. Brockman, "Performance analysis of accelerated image registration using GPGPU," in *Proceedings of 2nd workshop* on general purpose processing on graphics processing units, 2009, pp. 38–45.
- [4] P. Muyan-Ozcelik, J. D. Owens, J. Xia, and S. S. Samant, "Fast deformable registration on the GPU: A CUDA implementation of demons." in *Proceedings of the 2008 International Conference on Computational Science and Its Applications (ICCSA)*, 2008.
- [5] A. Wong and P. Fieguth, "Fast phase-based registration of multimodal data," *Signal Processing*, vol. 89, pp. 724 – 737, 2008.
- [6] K. Pauwels and M. M. Van Hulle, "Realtime phase-based optical flow on the GPU," in *Computer Vision and Pattern Recognition Workshops*, 2008, pp. 1–8.
- [7] H. Knutsson, M. Andersson, and J. Wiklund, "Advanced filter design," in *Scandinavian conference on image analysis (SCIA)*, 1999.
- [8] T. Oakes, T. Johnstone, K. Ores Walsh, L. Greischar, A. Alexander, A. Fox, and R. Davidson, "Comparison of fMRI motion correction software tools," *Neuroimage*, vol. 28, pp. 529–543, 2005.
- [9] J. Warntjes, L. Dahlqvist, J. West, and P. Lundberg, "Optimization for clinical usage of rapid magnetic resonance quantification on the brain," *Magnetic Resonance in Medicine*, vol. 60, pp. 320– 329, 2008.
- [10] K. Friston, J. Ashburner, and J. Poline, "Spatial realignment and normalization of images," *Human Brain Mapping*, vol. 2, pp. 165– 189, 1995.