Parallel implementation of Total-FETI DDM with application to medical image registration

Michal Merta[1], Alena Vašatová[1], Václav Hapla[1], and David Horák[1]

Key words: Total-FETI, image registration, parallel computing

1 Introduction

The main task of image registration is to determine an optimal spatial transformation such that two (or more) images become, in a certain sense, similar. Therefore, it plays a crucial role in image processing if there is a need to integrate information from two (or more) source images. These images usually show the same scene, but taken at different times, from different viewpoints or by different sensors.

Image registration is used in various areas. In medical applications it serves to obtain more complete information about the patient (e.g., to monitor a progression or regression of a disease, to align pre- and post- contrast images, or to compare patient's data with anatomical atlases), to compensate a motion of a subject during medical scanning, to correct calibration differences across scanners etc. [10, 12]. For more examples of usage of medical image registration see [8].

The first attempts at medical image registration focused mainly on the processing of brain images. Hence, a rigid body approximation was sufficient, because of a relatively small possibilities for deformation inside the skull. Later, it was extended to the affine registration. However, rigid or even affine approximations are usually not sufficient for a registration of a human body. Therefore, the research in medical image processing is now focused on the development of non-rigid registration methods. One of them is the elastic registration introduced by Broit [1]. In this method, images are considered to be 2D elastic bodies. Volume forces defined from 'differences' of the two images then deform one image so that it becomes similar to the other. The disadvantage of this linear model is that it assumes small deformations. For large deformations it can be replaced by the viscous fluid model [2].

With the increasing amount of data provided by medical instruments like CT or MRI, a parallel implementation of image registration seems to be necessary. In this work we combine the method of elastic registration together with the Total-FETI method [3] to obtain scalable algorithm for registration of medical images.

[1] Centre of Excellence IT4Innovations, VŠB-Technical University of Ostrava, tř. 17. listopadu 15/2172, Ostrava, 708 33 Czech Republic, e-mail: {michal.merta}{alena.vasatova}{vaclav.hapla}{david.horak}@vsb.cz

2 Elastic registration

Image registration usually consists of three parts: choosing an appropriate transformation model, choosing a distance (similarity) measure, and optimization process. Let us use the notation from [10] and briefly describe the process.

In order to find a transformation of the template image T, such that after its application it becomes, in a certain sense, similar to the reference image R, we define a suitable distance measure $\mathscr{D}$ and minimize the distance between R and T with respect to searched transformation φ:

$$\min_{\varphi:\mathbb{R}^2\to\mathbb{R}^2} \mathscr{D}\left[R,T;\varphi\right], \tag{1}$$

where $\mathscr{D}\left[R,T;\varphi\right] := \mathscr{D}\left[R,T_\varphi\right]$.

However, this approach has its drawbacks: a solution is not necessarily unique and it actually may not exist. Thus, the problem (1) is ill-posed. Moreover, additional implicit constraints can emerge, e.g., in medical images no additional cracks or folding of the tissue are allowed (the transformation should be diffeomorphic). Both these situations can be solved by adding a regularizer [10].

Transformation model of elastic registration is based on a physical motivation that the images are two different observations of an elastic body, one before and one after a deformation. The transformation $\varphi : \mathbb{R}^2 \to \mathbb{R}^2$ is split into the identity part and the displacement $u : \mathbb{R}^2 \to \mathbb{R}^2$:

$$\varphi(x) := x - u(x). \tag{2}$$

As the regularizer we use the linearized elastic potential

$$\mathscr{P}\left[u\right] := \int_\Omega \frac{\mu}{4}\sum_{j=1}^{2}\sum_{k=1}^{2}\left(\partial_{x_j}u_k + \partial_{x_k}u_j\right)^2 + \frac{\lambda}{2}(\operatorname{div} u)^2\,\mathrm{d}V, \tag{3}$$

where λ and μ are the Lamé parameters. The regularizer has the meaning of volume forces, which implicitly constrain the displacement to fulfill a smoothness criteria. We obtain the following regularized problem which is more suitable for a numerical realization:

$$\mathscr{J}\left[u\right] = \min_{v:\mathbb{R}^2\to\mathbb{R}^2} \mathscr{J}\left[v\right], \quad \text{where} \quad \mathscr{J}\left[v\right] := \mathscr{D}\left[R,T;v\right] + \alpha\mathscr{P}\left[v\right]. \tag{4}$$

Here, the parameter $\alpha \in \mathbb{R}^+$ controls the strength of the smoothness of the displacement versus the similiarity of the images. In the case of the elastic registration it is usually omitted, since it can be included in the Lamé parameters. Therefore, let us assume $\alpha = 1$ in what follows.

A distance measure is a cost function which determines a similarity of two images. We choose the so-called sum of squared differences (SSD):

$$\mathscr{D}[R,T;u] := \frac{1}{2}\|T_u - R\|^2_{L_2(\Omega)}, \tag{5}$$

where $T_u(x) := T(x - u(x))$. The volume forces

$$f(x,u(x)) := (R(x) - T_u(x))\nabla T_u(x), \tag{6}$$

$f : \mathbb{R}^2 \to \mathbb{R}^2$, derived from its Gâteaux derivative, push a transformed image into the direction of a reference.

Images are represented by the compactly supported mappings $R,T : \Omega \to \mathbb{R}$, where $\Omega := (0,1)^2$. $T(x)$ and $R(x)$ denote the intensities of images at the spatial position x; we set $R(x) := 0$ and $T(x) := 0$ for all $x \notin \Omega$.

By applying the Gâteaux derivative to the elastic potential (3) we obtain the Navier-Lamé operator of classical elasticity. The displacement of the elastic body and therefore the transformation of the image T is then obtained as the solution of the partial differential equation with zero Dirichlet boundary condition:

$$\begin{cases} \mu\Delta u(x) + (\lambda+\mu)\nabla \operatorname{div} u(x) = -f(x,u(x)) & \text{in } \Omega, \\ u(x) = 0 & \text{on } \partial\Omega. \end{cases} \tag{7}$$

There are several possibilities how to overcome the non-linearity of the previous equation. In the simplest case, when the difference between the reference and the template image is small enough, we set

$$f(x,u(x)) := f(x,0) = (R(x) - T(x))\nabla T(x), \tag{8}$$

and obtain a linearized problem. Otherwise, we solve the problem iteratively using the Algorithm 1. The similar algorithm is presented in [10], where the finite differ-

Algorithm 1 Fixed-point iteration for the solution of Equation (7)

$T_0(x) := T(x)$
$f_0(x) := (R(x) - T_0(x))\nabla T_0(x)$
for $k = 1$ **to** K **do**
  solve (7) for u_k with $f(x,u(x)) := f_{k-1}$
  $T_k(x) := T_{k-1}(x - u_k)$
  $f_k(x) := (R(x) - T_k(x))\nabla T_k(x)$
end for

ence method is used for the solution of the linearized problem.

We discretize the linearized problem using a finite element method with piecewise affine basis functions on triangular elements. To approximate the gradient of T_u, which is necessary for the evaluation of forces f, we use a convolution with an appropriate kernel of the Sobel operator (see, e.g., [11]). The solution can be easily parallelized by the Total-FETI method described in the following part.

3 Parallelization using Total-FETI method

The numerical solution of the linearized version of the problem (7) can be effectively parallelized by the Total-FETI (TFETI) method which is a variant of the FETI method originally proposed by Farhat et. al. [6]. The method is based on the decomposition of the spatial domain into non-overlapping subdomains. The continuity of the solution among subdomains is enforced by Lagrange multipliers. Total-FETI by Dostál et al. [3] simplifies the inversion of stiffness matrices of subdomains by using Lagrange multipliers also to enforce the Dirichlet boundary condition. Using this approach, all subdomains are floating and their stiffness matrices have the same kernels formed by the vectors of the rigid body modes.

To apply the FETI based domain decomposition, we partition the rectangular domain Ω, representing the processed image, into N geometrically identical rectangular subdomains Ω_s. We denote K_s, f_s, u_s, and B_s the subdomain stiffness matrix, the subdomain load vector, the subdomain displacement vector, and the subdomain constraint matrix, respectively. Let us also denote R_s as the matrix with columns forming the basis of the kernel of K_s. Notice, that because of this regular decomposition, the matrices K_s, as well as R_s, are the same for all subdomains. Therefore, they are computed only once and then redistributed among processors. Eventually, they can be stored in a shared memory.

After the decomposition we obtain the quadratic minimization problem with equality constraints

$$\min \frac{1}{2} u^T K u - u^T f \quad \text{s. t.} \quad Bu = c, \tag{9}$$

where

$$K := \begin{bmatrix} K_1 & & \\ & \ddots & \\ & & K_N \end{bmatrix}, \quad f := \begin{bmatrix} f_1 \\ \vdots \\ f_N \end{bmatrix}, \quad u := \begin{bmatrix} u_1 \\ \vdots \\ u_N \end{bmatrix}, \quad B := [B_1, \dots, B_N]. \tag{10}$$

Applying the duality theory to the equivalent saddle-point problem and establishing the notation

$$F := BK^\dagger B^T, G := R^T B^T, d := BK^\dagger f, e := R^T f,$$

where $K^\dagger$ denotes a generalised inverse matrix satisfying $KK^\dagger K = K$ (see, e.g., [4]), and R denotes the block-diagonal matrix with blocks R_s, we obtain the following minimization problem:

$$\min \frac{1}{2} \lambda^T F \lambda - \lambda^T d \quad \text{s.t.} \quad G\lambda = e. \tag{11}$$

We can further homogenize the equality constraints $G\lambda = e$ to $G\mu = 0$ by decomposing λ into $\mu \in \operatorname{Ker} G$ and $\tilde{\lambda} \in \operatorname{Im} G^T$ as

$$\lambda := \mu + \tilde{\lambda}. \tag{12}$$

We get $\tilde{\lambda}$ easily by $\tilde{\lambda} = G^T(GG^T)^{-1}e$. To enforce the condition $G\mu = 0$ we introduce the projector $P := I - Q$ to the null space of G. Here $Q := G^T(GG^T)^{-1}G$ is the projector onto the image space of G^T. The final problem for μ reads (note that $P\mu = \mu$):

$$PF\mu = P(d - F\tilde{\lambda}). \tag{13}$$

This problem can be effectively solved by the conjugate gradient method.

One of the advantages of the approach based on the Lagrange multipliers is the possibility to include other constraints to the matrix B than 'gluing' and Dirichlet conditions. One possibility is to use it to enforce the rigidity of certain parts of the processed image. These rigid parts can represent, e.g., bones. As mentioned in Section 2, the new coordinates $\varphi(x)$ of any point x after transformation are

$$\varphi(x) := x - u(x). \tag{14}$$

Using rigid body motions with a linearized rotation, this transformation can also be described by

$$x - u(x) = R_x a, \tag{15}$$

where

$$R_x := \begin{bmatrix} -x_2 & 1 & 0 \\ x_1 & 0 & 1 \end{bmatrix}, \tag{16}$$

and a is the vector of motion parameters (shifts and rotation). Conditions necessary to enforce a rigidity of a motion of two point $\tilde{x}, \tilde{y}$ can be derived from the following system of equations

$$\begin{cases} \tilde{x} - u(\tilde{x}) = R_{\tilde{x}} a, \\ \tilde{y} - u(\tilde{y}) = R_{\tilde{y}} a. \end{cases} \tag{17}$$

We eliminate a and obtain

$$-ou_1(\tilde{x}) - pu_2(\tilde{x}) + ou_1(\tilde{y}) + pu_2(\tilde{y}) = p^2 + o^2, \tag{18}$$

where $p := \tilde{y}_2 - \tilde{x}_2$, $o := \tilde{y}_1 - \tilde{x}_1$, and $u(x) := (u_1(x), u_2(x))$. These conditions are added to appropriate positions in the matrix B. To reduce the number of additional constraints, one can enforce the rigidity only on the boundaries of given areas.

4 Data parallelization and implementation using Trilinos framework

Parallelization of FETI/TFETI can be implemented using SPMD technique – distributing matrix portions among the processing units. The distribution of primal data is straightforward because of the block-diagonal structure of the system stiffness

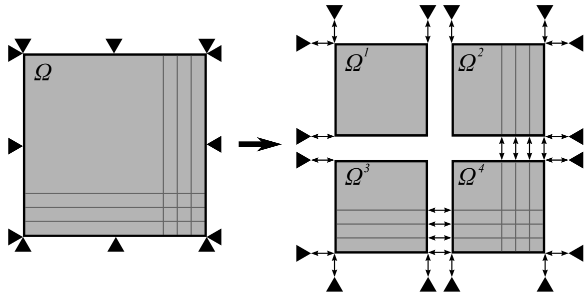

Fig. 1 Total-FETI domain decomposition of the 2D rectangular area. Dirichlet boundary conditions are enforced by Lagrange multipliers.

matrix. Each processor is assigned one rectangular part of the images R and T, and the corresponding primal data – one block of the global stiffness matrix K, one block of the kernel matrix R, and corresponding parts of the constraint matrix B, solution vector u, and right-hand side vector f. On the other hand, if we want to accelerate also the dual actions we have to distribute the dual objects as well. We distribute the matrix G into vertical blocks. All dual vectors are distributed accordingly to this (for more details see [9]).

For the parallel implementation we use the Trilinos framework [7] which is a collection of relatively independent packages developed by Sandia National Laboratories. It provides a tool kit for basic linear algebra operations (both serial and parallel), direct and iterative solvers to linear systems, PDE discretization utilities, etc. Its main advantages are object oriented design, high modularity and use of modern features of C++ language such as templating. It is currently in version 11.

In our codes we use the Epetra package as a base for linear algebra operations. It provides users with distributed dense vectors and matrices, as well as sparse matrices in compressed row format (`Epetra_CrsMatrix`), linear operators, distributed graphs, etc. As the object-oriented wrapper to direct linear system solver SuperLU, which is used for the solution of the coarse problem (application of $(GG^T)^{-1}$) and the application of the pseudoinverse $K^\dagger$, we use the Amesos package.

5 Numerical experiments

The numerical experiments were performed on the cluster consisting of 16 SMP nodes, each of the nodes is equipped with two Intel Xeon QuadCore 2.5 GHz CPUs and 18 GB of RAM. Table 1 shows the results of the scalability tests for the data obtained from Department of Oncology of University Hospital of Ostrava. We perfomed two experiments – one with no additional constraints, and the second on the same data but with a rigidity of the bones enforced by additional Lagrange multipliers. The processed data are depicted in Figure 2.

The problem is linearized using the approach (8). For the first experiment, the number of CG iterations is relatively low. For these numbers of dual variables the coarse problem (which is usually the main bottleneck of the FETI methods) is not big enough to affect the scalability and the increasing time per iteration is caused mainly by the communication and vector redistribution routines within the Trilinos framework. The second experiment shows that the additional constraints lead to the increase of the number of CG iterations. To reduce this number we can use the cheap lumped preconditioner $\overline{F^{-1}} = BKB^T$ (see [5]).

Table 1 Performance of the TFETI implementation for varying decomposition and discretization

Number of subdom.	1	4	16
Primal dimension	20,402	81,608	326,432
Dual dimension	808	2,424	8,080
CG time [s]	0.50	1.53	4.35
CG iterations	25	39	47
Time per iteration [s]	0.02	0.04	0.09
Example 1: Without rigid body parts			

Number of subdom.	1	4	16
Primal dimension	20,402	81,608	326,432
Dual dimension	903	2,641	8,254
CG time [s]	41.01	34.54	57.44
CG iterations	2467	990	665
Time per iteration [s]	0.01	0.03	0.08
Example 2: With rigid body parts			

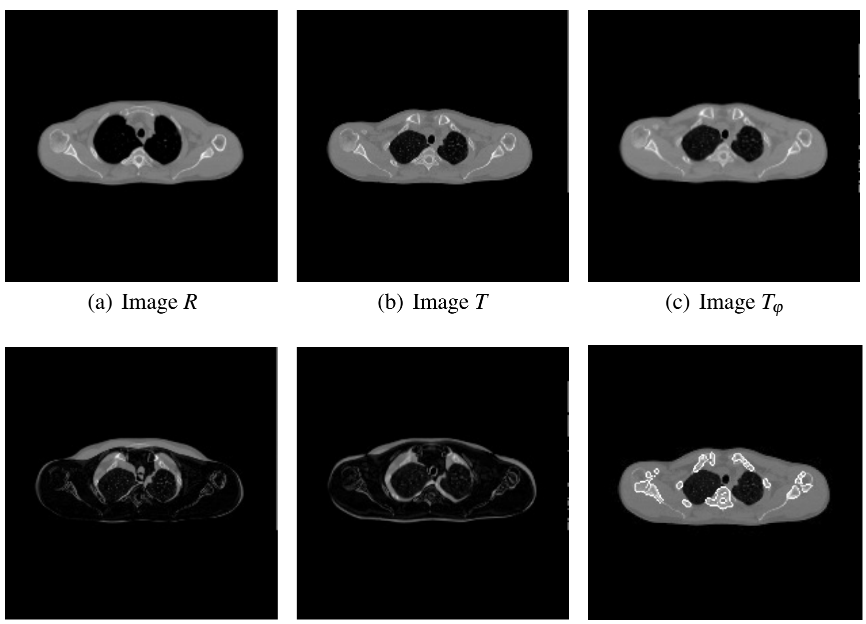

(a) Image R (b) Image T (c) Image T_φ

(d) Difference between R and T (e) Difference between R and T_φ (f) Boundaries of rigid parts (bones)

Fig. 2 Processed data - computer tomography of patient's chest. We search for a transformation φ of the image T (in exhalation) so it becomes similar to the image R (in inflation). For this experiment, we set $\mu = 5 \times 10^5$ and $\lambda = 0$.

6 Conclusion

We have demonstrated the applicability of the Total-FETI method to a parallelization of a process of image registration. Our implementation was tested on 2D computer tomography data obtained from University Hospital of Ostrava. Because of relatively low resolution of the images the total number of unknowns in the resulting systems did not exceed hundreds of thousands. However, these results enable us to focus on the development of domain decomposition-based methods for the image registration of 3D data, where the number of unknowns can easily reach tens or hundreds of millions.

Acknowledgements This work was supported by the European Regional Development Fund in the IT4Innovations Centre of Excellence project (CZ.1.05/1.1.00/02.0070) and by the project SPOMECH - Creating a multidisciplinary R&D team for reliable solution of mechanical problems, reg. no. CZ.1.07/2.3.00/20.0070 within Operational Programme 'Education for competitiveness' funded by Structural Funds of the European Union and state budget of the Czech Republic.

The research has also been supported by the grants: HPC-Europa2 project funded by the European Commission - DG Research in the Seventh Framework Programme under grant agreement No. 228398, PRACE 2IP project receiving funding from the EU's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. RI-283493.

References

1. Broit, C.: Optimal registration of deformed images. Ph.D. thesis, University of Pensylvania (1998)
2. Christenssen, G.: Deformable shape models for anatomy. Ph.D. thesis, Washington University (1994)
3. Dostál, Z., Horák, D., Kučera, R.: Total feti - an easier implementable variant of the feti method for numerical solution of elliptic pde. Commun. in Numerical Methods in Engineering **22**, 1155–1162 (2006)
4. Dostál, Z., Kozubek, T., Markopoulos, A., Menšík, M.: Cholesky decomposition and a generalized inverse of the stiffness matrix of a floating structure with known null space. Applied Mathematics and Computation **217**, 6067–6077 (2011)
5. Farhat, C., Mandel, J., Roux, F.: Optimal convergence properties of the feti domain decomposition method. Comput. Methods Appl. Mech. Eng. **115**, 365–385 (1994)
6. Farhat, C., Roux, F.: An unconventional domain decomposition method for an efficient parallel solution of large-scale finite element systems. SIAM Journal on Scientific Computing **13**, 379–396 (1992)
7. Heroux, M.: Trilinos web page. Available from http:://trilinos.sandia.gov (2012)
8. Hill, D., Holden, M., Hawkes, D.: Medical image registration. Physics in medicine and biology **46**, 1–45 (2001)
9. Horák, D., Hapla, V.: Tfeti coarse space projectors parallelization strategies. Lecture Notes in Computer Science. **7203**, 152–162 (2012)
10. Modersitzki, J.: Numerical Methods for Image Registration. Oxford University Press, Oxford (2004)
11. Pratt, W.K.: Digital Image Processing: PIKS Inside. John Wiley & Sons, Inc., New York (2001)
12. Zitova, B., Flusser, J.: Image registration methods: a survey. Image and Vision Computing **21**, 977–1000 (2003)