Decomposition of Quadratic Variational Problems*

Florian Becker, Christoph Schnörr

Image and Pattern Analysis Group, Heidelberg Collaboratory for Image Processing, University of Heidelberg, Germany, {becker,schnoerr}@math.uni-heidelberg.de

Abstract. Variational problems have proved of value in many image processing and analysis applications. However increase of sensor resolution as occurred in medical imaging and experimental fluid dynamics demand for adapted solving strategies to handle the huge amount of data.
In this paper we address the decomposition of the general class of quadratic variational problems, which includes several important problems, such as motion estimation and image denoising. The basic strategy is to subdivide the originally intractable problem into a set of smaller convex quadratic problems. Particular care is taken to avoid ill-conditioned sub-problems. We demonstrate the approach by means of two relevant variational problems.

1 Introduction

Variational approaches to motion estimation are nowadays routinely used in many image processing applications. A key problem, however, concerns the ever increasing sizes of data sets to be processed, in particular for analysing 3D image sequences in medical imaging and experimental fluid dynamics. For example, the next generation of imaging sensors in experimental fluid will deliver data volumes taken with high-speed cameras, that cannot be handled within the working memory of a single PC. This necessitates to investigate coarse problem decompositions that can process in parallel large-scale problems with off-the-shelf hardware.

In this paper, we present a decomposition approach for a fairly general class of variational problems, including higher-order regularisation, for instance. We pay attention to obtain easy-to-solve subproblems that communicate in order to compute provably the unique global solution.

Related Work and Contribution. The use of variational domain decomposition for motion estimation has been introduced to the field of image processing by Kohlberger et al. [1]. This approach is not applicable however to variational

* This work was partially financed by the EC project FLUID (FP6-513663).

models involving higher-order regularisation, like [2] for instance, because sub-problems in inner domains become inherently singular, and because the corresponding handling of boundary conditions becomes involved.

In this paper, we therefore adopt the viewpoint of convex programming to tackle the decomposition problem [3]. We consider the general class of convex quadratic optimisation problems and propose a method to decompose any instance into a sum of still convex sub-functions. This allows to apply the dual decomposition approach. The initial problem can then be solved as several smaller independent convex problems and computationally cheap synchronisation steps without changing the overall objective. We describe an extension that allows to improve the numerical properties of the underlying problem, hence improving convergence rate. Two problems from motion estimation are decomposed exemplarily. Error measurements in comparison to single-domain solutions are presented.

Organisation. In section 2 we summarise the underlying idea of dual problem decomposition and specialise it to the considered class of optimisation problems. Section 3 describes an iterative method for solving the problem in its decomposed formulation. An extension is proposed to improve the numerical properties of the sub-problems. Experimental results for the application to two variational motion estimation problems are given in section 4.

2 Problem Decomposition

2.1 Dual Decomposition of Convex Optimisation Problems

Our approach is based on the idea known as dual decomposition [4]. The method requires that the objective function of a convex optimisation problem,

$$\min_u f(u)\,, \qquad f(u) : \mathbb{R}^n \mapsto \mathbb{R}\,, \tag{1}$$

can be decomposed into a sum of d convex sub-functions, so $f(u) = \sum_{l=1}^{d} f_l(u)$. For demonstrating the basic idea we restrict ourselves to $d = 2$. The variable vector u is split up into internal variables $u_l \in \mathbb{R}^{n_l}$ that are only involved in exactly one sub-function, and a vector of complicating variables $y \in \mathbb{R}^{n_C}$ that are common to f_1 and f_2. Next we introduce an own set of complicating variables y_l for each sub-function and enforce their identity by a consistency constraint, leading to an optimisation problem equivalent to (1):

$$\min_{\{u_l, y_l\}} f_1(u_1, y_1) + f_2(u_2, y_2) \qquad \text{s.t. } y_1 = y_2 \tag{2}$$

For briefness we denote a set of vectors $\{x_1, \ldots, x_d\}$ as $\{x_l\}$. With the Lagrangian function which is defined as $L(\{u_l, y_l\}, \lambda) := f_1(u_1, y_1) + f_2(u_2, y_2) + \lambda^\top (y_1 - y_2)$

we can make the constraints implicit and gain the primal (P) and dual (D) Lagrangian problem corresponding to (2),

$$\text{(P)} \qquad p^* := \min_{\{u_l, y_l\}} \sup_{\lambda} \; L(u, \lambda) \tag{3}$$

$$\text{(D)} \qquad d^* := \max_{\lambda} \inf_{\{u_l, y_l\}} \; L(u, \lambda) \tag{4}$$

Their optimal values are related as $p^* \geq d^*$ (weak duality). For convex functions without inequality constraints this relation holds strictly and allows to solve the primal problem by the way of its dual formulation,

$$\underbrace{\max_{\lambda} \underbrace{\inf_{u_1, y_1} \; f_1(u_1, y_1) + \lambda^\top y_1}_{\text{sub-problem 1}} + \underbrace{\inf_{u_2, y_2} \; f_2(u_2, y_2) - \lambda^\top y_2}_{\text{sub-problem 2}}}_{\text{master problem}} . \tag{5}$$

The problem decomposes into two convex sub-problems with independent variables embedded into a master problem that updates λ. This formulation allows to specify an iterative method,

`initialise` $\lambda \leftarrow 0$, $u_l \leftarrow 0$, $y_l \leftarrow 0$, $k \leftarrow 0$
`repeat`
  `for` $l = 1, \ldots, d$ `do:`
    $\left(u_l^{(k+1)}, y_l^{(k+1)}\right) \leftarrow$ `solution of sub-problem` l
  $\lambda^{(k+1)} \leftarrow \lambda^{(k)} + \alpha^{(k)} \nabla_\lambda L \left(\left\{ u_l^{(k+1)}, y_l^{(k+1)} \right\}, \lambda^{(k)} \right)$
  $k \leftarrow k + 1$
`until convergence` .

In each iteration the primal variables are refined by solving smaller sub-problems which are preferably of a simple structure. Due to their independence this task can be performed in parallel. Afterwards the dual variables are updated, e.g. by a sub-gradient method such as cutting plane, bundle or trust region which choose $\alpha^{(k)}$ – we refer to [3] for details. The update steps are iterated until a stopping criteria is met, for example the change in the primal and/or dual variables.

2.2 Considered Class of Problems

Here we consider the class of convex quadratic optimisation problems given in the form

$$\min_{u \in \mathbb{R}^n} \frac{1}{2} \left\| Du + c \right\|_2^2 \qquad \text{with } D \in \mathbb{R}^{m \times n},\ c \in \mathbb{R}^m \tag{6}$$

which for example includes the variational approach to optical flow estimation by Yuan et al. [2]:

$$\min_u \frac{1}{2}\int_\Omega \left\|\nabla g^\top u + g_t\right\|_2^2 + \beta_1\|\nabla \mathrm{div}\, u\|_2^2 + \beta_2\|\nabla \mathrm{curl}\, u\|_2^2 \,\mathrm{d}x + \frac{1}{2}\beta_3\int_{\partial\Omega}\|\partial_n u\|_2^2\,\mathrm{d}x \tag{7}$$

In section 4 we solve this problem in its decomposed form.

2.3 Decomposition of Convex Quadratic Problems

We assume that the problem description can be rearranged by permuting the order of variables in u and the row order in D and c in order to get a form that makes the separable structure explicit:

$$f(u) = \frac{1}{2}\|Du + c\|_2^2 = \frac{1}{2}\left\|\begin{pmatrix} D_{1,\mathrm{C}} & 0 & \cdots & 0 & D_{1,\mathrm{C}} \\ 0 & D_{2,\mathrm{I}} & \ddots & \vdots & D_{2,\mathrm{C}} \\ \vdots & \ddots & \ddots & 0 & \vdots \\ 0 & \cdots & 0 & D_{d,\mathrm{I}} & D_{d,\mathrm{C}} \end{pmatrix}\begin{pmatrix} u_1 \\ \vdots \\ u_d \\ y \end{pmatrix} + \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_d \end{pmatrix}\right\|_2^2 \tag{8}$$

Then $D_{l,\mathrm{I}} \in \mathbb{R}^{m_l \times n_l}$ and $D_{l,\mathrm{C}} \in \mathbb{R}^{m_l \times n_\mathrm{C}}$ represent the coefficients of the internal respectively complicating variables and $c_l \in \mathbb{R}^{m_l}$ the constant parts. The objective function decomposes into

$$\begin{aligned} &\sum_{l=1}^{d}\left(\frac{1}{2}\begin{pmatrix} u_l \\ y \end{pmatrix}^\top \underbrace{\begin{pmatrix} D_{l,\mathrm{I}} \\ D_{l,\mathrm{C}} \end{pmatrix}^\top \begin{pmatrix} D_{l,\mathrm{I}} \\ D_{l,\mathrm{C}} \end{pmatrix}}_{A_l}\begin{pmatrix} u_l \\ y \end{pmatrix} + \begin{pmatrix} u_l \\ y \end{pmatrix}^\top \underbrace{\begin{pmatrix} D_{l,\mathrm{I}}^\top \\ D_{l,\mathrm{C}}^\top \end{pmatrix} c_l}_{-b_l :=} + \underbrace{\frac{1}{2}c_l^\top c_l}_{\text{const}}\right) \\ =&\sum_{l=1}^{d}\underbrace{\frac{1}{2}\begin{pmatrix} u_l \\ y \end{pmatrix}^\top A_l \begin{pmatrix} u_l \\ y \end{pmatrix} - \begin{pmatrix} u_l \\ y \end{pmatrix}^\top b_l}_{f_l(u_l,y):=} + \text{const} = \sum_{l=1}^{d} f_l(u_l, y) + \text{const}\,. \end{aligned} \tag{9}$$

Note that A_l are symmetric and due to $x^\top A_l x = \left\|\begin{pmatrix} D_{l,\mathrm{I}} & D_{l,\mathrm{C}} \end{pmatrix} x\right\|_2^2 \geq 0\ \forall x$ positive semidefinite matrices. Hence each sub-function is convex and quadratic by construction. We further assume that the decomposition is chosen such that the matrices A_l are also non-singular.

3 Optimisation of Decomposed Quadratic Problems

In the previous section we showed that the objective function (8) can be decomposed into a sum of convex sub-functions and hence we can apply the dual decomposition method for solving the problem (6).

First of all we combine the internal and complicating variables of each sub-function into $x_l := \left(u_l^\top \; y_l^\top\right)^\top$ and define a set of linear operators $\{C_l\}$ to represent the consistency constraints, e.g. as

$$\begin{pmatrix} y_1 - y_2 \\ y_2 - y_3 \\ \vdots \\ y_{d-1} - y_d \end{pmatrix} = \sum_{l=1}^{d} C_l x_l = 0 \; .$$

Note that although x_l contains local variables the constraints only involve complicating ones. Then the following decomposed problem is equivalent to (6).

$$\min_{\{x_l\}} \sum_{l=1}^{d} \frac{1}{2} x_l^\top A_l x_l - x_l^\top b_l \qquad \text{s.t. } \sum_{l=1}^{d} C_l x_l = 0 \tag{10}$$

With the definition of the Lagrange function,

$$L\left(\{x_l\}, \lambda\right) = \sum_{l=1}^{d} \left(\frac{1}{2} x_l^\top A_l x_l - x_l^\top b_l + \lambda^\top C_l x_l \right) \; ,$$

the problem in its decomposed *dual* form reads

$$\max_{\lambda} \sum_{l=1}^{d} \left(\inf_{x_l} \; \frac{1}{2} x_l^\top A_l x_l - x_l^\top \left(b_l - C_l^\top \lambda \right) \right) \; . \tag{11}$$

For the considered class of optimisation problems, each sub-problem is an unconstrained convex qua et aldratic optimisation problem. Due to this beneficial properties, any first-order optimal solution x_l with $\nabla f_l(x_l) = 0$ is also a global minimum for f_l. This can efficiently be found by solving the linear system $A_l x_l = b_l - C_l^\top \lambda$, e.g. by a conjugate gradient method [5].

The variables λ of the enveloping master problem are updated by moving along the gradient of the Lagrange function, $\nabla_\lambda L(\{x_l\}, \lambda) = \sum_{l=1}^{d} C_l x_l$. The update steps for the primal and dual variables then read

$$x_l^{(k+1)} \leftarrow A_l^{-1} \left(b_l - C_l^\top \lambda^{(k)} \right), \; l = 1, \ldots, d \tag{12}$$

$$\lambda^{(k+1)} \leftarrow \lambda^{(k)} + \alpha^{(k)} \sum_{l=1}^{d} C_l x_l^{(k+1)} \; . \tag{13}$$

The step scaling $\alpha^{(k)} > 0$ may be chosen constant or – as proposed in [3] – as a sequence with $\lim_{k\to\infty} \alpha^{(k)} \to 0$ and $\sum_{k=1}^{\infty} \alpha^{(k)} \to \infty$.

3.1 Extension for Badly Conditioned Sub-Problems

The matrices A_l which arise from the decomposition (9) are positive definite but may be badly conditioned, depending on the problem and actual decomposition.

We propose a framework that allows to improve the numerical properties of the underlying optimisation problems without altering the overall objective.

Instead of solving $\min_{x_l} f_l(x_l)$ in each iteration and for each sub-problem we modify the objective functions by adding a regularisation term which involves the current value of $x_l^{(k)}$:

$$f_l(x) + \frac{1}{2}\left\| B_l^{\frac{1}{2}}\left(x - x_l^{(k)}\right)\right\|_2^2 = \frac{1}{2}x^\top (A_l + B_l)\, x - x^\top\left(b_l + B_l x_l^{(k)}\right) + \frac{1}{2}\left\| B_l^{\frac{1}{2}} x_l^{(k)}\right\|_2^2$$

with an arbitrary symmetric positive semidefinite matrix B_l. Then the new iteration steps (replacing (12)) for the primal variables are

$$x_l^{(k+1)} \leftarrow (A_l + B_l)^{-1}\left(b_l + B_l x_l^{(k)} - C_l^\top \lambda^{(k)}\right) . \tag{14}$$

In this representation it becomes apparent that B_l allows to directly modify the linear system that has to be solved for each sub-problem. For $B_l = 0\ \forall l$ we obtain the original update step (12).

In order to show that the altered method still solves the original problem, we rewrite the steps (14) together with the unmodified update for the dual variables (13) as a single system of the form

$$\underbrace{\begin{pmatrix} A_1 + B_1 & & 0 & 0 \\ & \ddots & & \vdots \\ 0 & & A_d + B_d & 0 \\ C_1 & \cdots & C_d & -\frac{1}{\alpha} I \end{pmatrix}}_{M:=} \underbrace{\begin{pmatrix} x_1^{(k+1)} \\ \vdots \\ x_d^{(k+1)} \\ \lambda^{(k+1)} \end{pmatrix}}_{z^{(k+1)}:=} = \underbrace{\begin{pmatrix} B_1 & & 0 & -C_1^\top \\ & \ddots & & \vdots \\ 0 & & B_d & -C_d^\top \\ 0 & \cdots & 0 & -\frac{1}{\alpha} I \end{pmatrix}}_{N:=} \underbrace{\begin{pmatrix} x_1^{(k)} \\ \vdots \\ x_d^{(k)} \\ \lambda^{(k)} \end{pmatrix}}_{z^{(k)}:=} + \underbrace{\begin{pmatrix} b_1 \\ \vdots \\ b_d \\ 0 \end{pmatrix}}_{b:=}$$

which can be interpreted as a splitting method, a class of iterative methods for solving linear problems – we refer to [5] for details.

Theorem 1. *[5, Th. 10.1.1] If $M - N$ and M are nonsingular and the singular values of $G := M^{-1}N$ lie within the unit circle, i.e. $\rho(G) < 1$, then the iteration $Mz^{(k+1)} = Nz^{(k)} + b$ converges to a solution of $(M - N)z = b$ for any initial $z^{(0)}$.*

Corollary 2. *The iteration* (13)-(14) *solves* (11) *if it converges.*

Proof. Due to Theorem 1 the fix-point of the proposed iteration method reads

$$(M - N)z = \begin{pmatrix} A_1 & & 0 & C_1^\top \\ & \ddots & & \vdots \\ 0 & & A_d & C_d^\top \\ C_1 & \cdots & C_d & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_d \\ \lambda \end{pmatrix} = \begin{pmatrix} b_1 \\ \vdots \\ b_d \\ 0 \end{pmatrix} = b\ .$$

These equations exactly represent the first order Karush-Kuhn-Tucker conditions of the original problem,

$$\nabla_{x_l} L(\{x_l\}, \lambda) = 0\ \forall l \quad \text{and} \quad \nabla_\lambda L(\{x_l\}, \lambda) = 0\ . \tag{15}$$

Hence the method solves the original problem for *any* B_l if it converges. □

4 Experiments and Discussion

In this section we apply the proposed decomposition method to two variational motion estimation approaches. In order to measure the exactness we compare the results to the non-decomposed solution of the problem. We examine the evolution of the error measurements over iterations of the master problem to investigate the time complexity of the method.

A synthetic image pair of size 500 by 500 pixels was used as input data for the experiments. The ground truth vector field (see Fig. 1(d)) is affine in the coordinates and has a maximum magnitude of 1 and 0.46 pixels in average.

In all experiments we chose a geometrical based method to determine the decomposition into sub-functions. The sub-problems were solved as linear programs with a conjugate gradient method. For the master problem we defined

$$\alpha^{(k)} = \frac{1}{a+bk} \frac{1}{\|\nabla_\lambda L(\{x_l\}, \lambda)\|_2} \,. \tag{16}$$

Both the primal and dual variables were initialised with all-zero vectors.

For every grid position $x \in \Omega$ we determined the Euclidean distance $e(x)$ between the solutions of the decomposed and non-decomposed problems. The following overall error measurements were evaluated:

$$\mu_e := \frac{1}{|\Omega|} \sum_{x \in \Omega} e(x) \,, \quad \sigma_e := \sqrt{\sum_{x \in \Omega} \frac{(e(x) - \mu_e)^2}{|\Omega| - 1}} \text{ and } \quad \max\, e := \max_{x \in \Omega}\, e(x)$$

The results and error measurements were recorded for each iteration. Reaching the maximum number of iterations - ten - was the only stopping criteria used.

4.1 Horn and Schunck

In [6], Horn and Schunck proposed the following variational approach to globally estimate the optical flow field:

$$\min_u \int_\Omega \frac{1}{2} \left\|\nabla g^\top u + g_t\right\|_2^2 + \frac{1}{2}\beta \left\|\nabla u_1\right\|_2^2 + \frac{1}{2}\beta \left\|\nabla u_2\right\|_2^2 \mathrm{d}x$$

where $u(x) := \begin{pmatrix} u_1(x) & u_2(x) \end{pmatrix}^\top$. A finite difference scheme was used for discretisation, allowing to write the discrete version of the objective function as

$$f(u) := \frac{1}{2} \left\|D_{\mathrm{OFC}} u + c_{\mathrm{OFC}}\right\|_2^2 + \frac{1}{2}\beta \left\|D_{\partial_x} u\right\|_2^2 + \frac{1}{2}\beta \left\|D_{\partial_y} u\right\|_2^2$$

with the linear operators D_{OFC}, D_{∂_x} and D_{∂_y} and constant vector c_{OFC}. To make the proposed decomposition method applicable, we rewrite the objective function as

$$f(u) = \frac{1}{2} \left\| \begin{pmatrix} D_{\mathrm{OFC}} \\ \sqrt{\beta} D_{\partial_x} \\ \sqrt{\beta} D_{\partial_y} \end{pmatrix} u + \begin{pmatrix} c_{\mathrm{OFC}} \\ 0 \\ 0 \end{pmatrix} \right\|_2^2 .$$

For the decomposition we divided the grid into four areas which overlap by one grid unit, see Fig. 1(a).

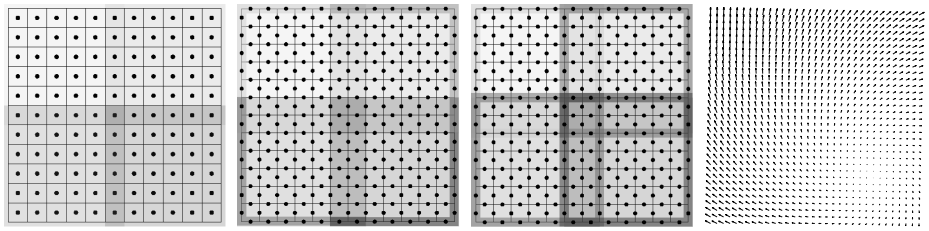

(a) (b) (c) (d)

Fig. 1. Discretisation grid and decomposition, dots represent variables, shadings represent the four (overlapping) subdivisions: (a) approach by Horn&Schunck, (b) approach by Yuan et al. with original boundary terms and (c) with additional terms on subdivision boundaries, (d) ground truth vector field (sub-sampled).

Results. For the experiments we chose the regularisation parameter as $\beta = 0.1$. The parameters for the dual variable update rule (16) were set to $a = -300$ and $b = 400$.

The error plot over the iterations one to ten in Fig. 2(b) shows a steep drop of all error measurements within few iterations. After ten iterations they reach $\mu_e = 3.57 \cdot 10^{-6}$, $\sigma_e = 4.04 \cdot 10^{-5}$ and $\max e = 0.0028$. According to Fig. 2(a) the errors of the order of $\max e$ are located at only few positions at the artificial borders especially where all four areas meet and quickly reach 10^{-15} in the remaining parts of the subdivisions.

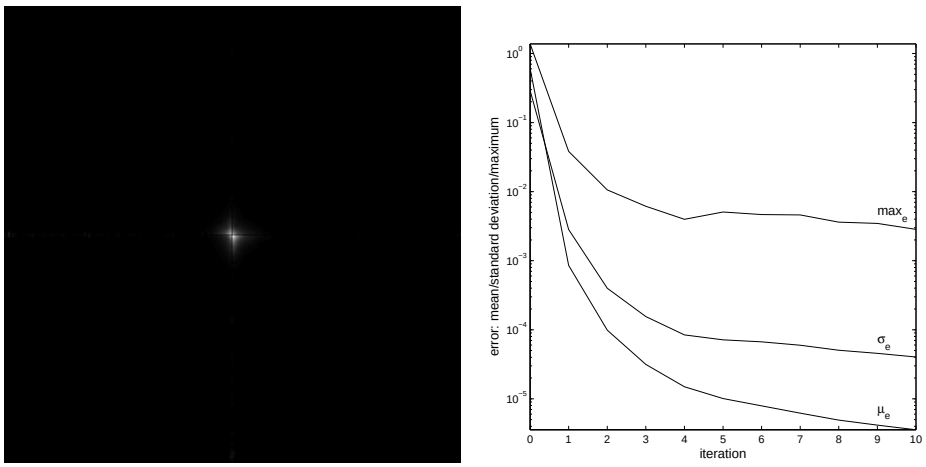

(a) error positions, $\max e = 0.0028$ (b) error history (logarithmic scale)

Fig. 2. Horn&Schunck: error of decomposed solution compared to single-domain result, (a) distribution over coordinates after ten iterations, (b) evolution over ten iterations.

4.2 Optical Flow Estimation with Higher Order Regularisation

The second approach we consider for decomposition uses higher order regularisation of the vector field (7). Mimetic differences as described in [7] were used for discretisation. In Fig. 1(b) the underlying variable grid is depicted. Using the notation introduced in [2] the discrete version of the objective function reads

$$\frac{1}{2}\left\|Gu+\partial_t g\right\|_{H_V}^2+\frac{1}{2}\beta_1\left\|\overline{\mathbb{G}}\mathbb{D}iv\ u\right\|_{H_S}^2+\frac{1}{2}\beta_2\left\|\mathbb{G}\overline{\mathbb{C}url}\ u\right\|_{H_E}^2+\frac{1}{2}\beta_3\left\|\mathbb{P}u\right\|_{bc}^2$$

where G, $\overline{\mathbb{G}}$, $\mathbb{D}iv$, ∂_t, $\overline{\mathbb{C}url}$ and $\mathbb{P}$ are linear operators and $g=\frac{1}{2}(g_1+g_2)$ represents the image information in a vector representation, i.e. the columns stacked together. Due to the fact that the definitions of the H_V-, H_S-, H_E- and bc-norms are equivalent to the Euclidean norm we are able to rewrite the objective function as

$$f(u)=\frac{1}{2}\left\|\begin{pmatrix} G \\ \sqrt{\beta_1}\,\overline{\mathbb{G}}\,\mathbb{D}iv \\ \sqrt{\beta_2}\,\mathbb{G}\,\overline{\mathbb{C}url} \\ \sqrt{\beta_3}\,\mathbb{P}\end{pmatrix}u+\begin{pmatrix}\partial_t g\\0\\0\\0\end{pmatrix}\right\|_2^2.$$

This is exactly the form presumed by the proposed decomposition method.

Results. First experiments showed that the decomposition by the proposed method leads to convex sub-functions but with badly conditioned matrices A_l. This causes a divergent behaviour of the algorithm. In order to improve the problem properties we imposed additional terms $\frac{1}{2}\left\|\partial_n u\right\|_2^2$ on the artificial boundaries. Thereby each modified sub-problems has regularisation terms at *all* its boundaries just as if it was a smaller instance of the original one. In Fig. 1(c) the location of the boundary terms are represented as shaded areas. In contrast Fig. 1(b) shows the unmodified grid decomposition where each sub-problem is only regularised by boundary terms at the borders of the original problem. The discrete forms of the new regularisation terms were concentrated into the matrices B_l and applied as described in section 3.1. As a result of this modification the condition number dropped from $\rho(A_l)\approx 10^8$ to $\rho(A_l+B_l)\approx 10^5$ and the algorithm converged quickly. The problem parameters were set to $\beta_1=\beta_2=0.4$ and $\beta_3=0.2$. The algorithm parameters were chosen by hand, $a=900$ and $b=100$. The four areas overlap by two grid cells.

Figure 3(b) shows that the error measurements reduce drastically within the first iteration and reach the order of their final values, $\mu_e=1.22\cdot 10^{-4}$, $\sigma_e=6.39\cdot 10^{-4}$ and $\max e=0.02$. The errors are mainly located around the artificial boundaries (see Fig. 3(a)) and tend to 10^{-15} in the remaining areas.

5 Conclusion and Further Work

We presented a convex programming approach to the decomposition of a fairly general class of variational approaches to image processing. The overall algorithm

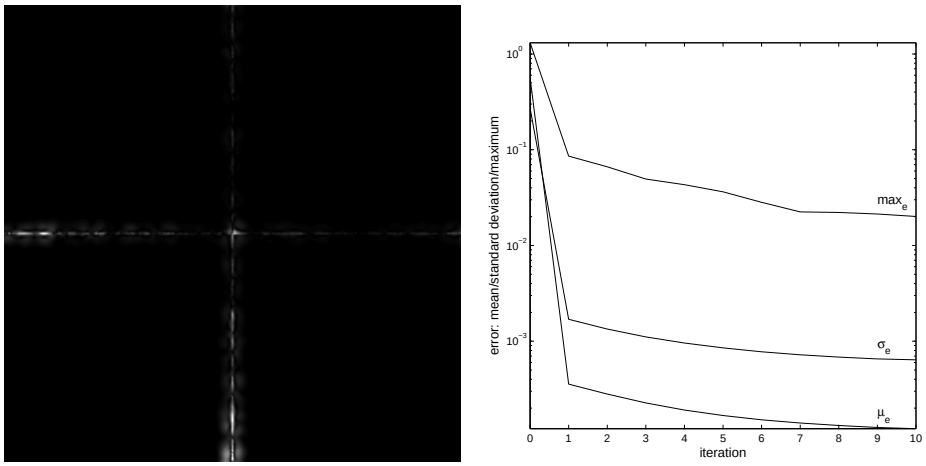

(a) error positions, $\max e = 0.02$ (b) error history (logarithmic scale)

Fig. 3. Yuan et al.: error of decomposed solution compared to single-domain result, (a) distribution over coordinates after ten iterations, (b) evolution over ten iterations.

converges to the optimality conditions and can be carried out by solving well-conditioned subproblems in parallel.

Our further work will focus on estimates of the convergence rate, and on connections to the domain decomposition literature [8] concerning preconditioning.

References

1. Kohlberger, T., Schnörr, C., Bruhn, A., Weickert, J.: Domain Decomposition for Variational Optical Flow Computation. IEEE Trans. Image Process. **14**(8) (2005) 1125–1137
2. Yuan, J., Schnörr, C., Mémin, E.: Discrete Orthogonal Decomposition and Variational Fluid Flow Estimation. J. Math. Imaging Vis. **28**(1) (2007) 67–80
3. Conejo, A.J., Castillo, E., Minguez, R., García-Bertrand, R.: Decomposition Techniques in Mathematical Programming. Springer (2006)
4. Boyd, S., Xiao, L., Mutapcic, A.: Notes on Decomposition Methods. Technical report, Stanford University (2003)
5. Golub, G.H., van Loan, C.F.: Matrix Computations. The John Hopkins University Press (1996)
6. Horn, B., Schunck, B.: Determining Optical Flow. Artif. Intell. **17** (1981) 185–203
7. Hyman, J.M., Shashkov, M.J.: Natural Discretizations for the Divergence, Gradient, and Curl on Logically Rectangular Grids. Comput. Math. Appl. **33**(4) (1997) 81–104
8. Toselli, A., Widlund, O.: Domain Decomposition Methods - Algorithms and Theory. Springer, Berlin, Germany (2004)