Proc. 1997 IEEE International Conference on Image Processing
(ICIP–97, Santa Barbara, Oct. 26–29, 1997), Vol. 3, pp. 396–399, 1997

Parallel Implementations of AOS Schemes: A Fast Way of Nonlinear Diffusion Filtering

J. Weickert,* K.J. Zuiderveld, B.M. ter Haar Romeny, W.J. Niessen
Image Sciences Institute, Utrecht University Hospital
E01.334, Heidelberglaan 100, NL–3584 CX Utrecht
E-mail: Joachim.Weickert@cv.ruu.nl

Abstract

In most cases nonlinear diffusion filtering is implemented by means of explicit finite difference schemes. These algorithms are not very efficient, since they are only stable for small time steps. We address this problem by presenting unconditionally stable semi-implicit schemes which are based on an additive operator splitting (AOS). They are very efficient since they can be implemented by recursive filtering, and their separability allows a straightforward implementation in any dimension. We analyse their behaviour on a parallel computer and demonstrate that parallel AOS schemes on a modern shared-memory multiprocessor system with 8 processors allow a speed-up of two orders of magnitude in comparison to the widely-used explicit scheme on a single processor.

1 Introduction

Although nonlinear diffusion filters are quite popular, their practical applicability suffers from the fact that the widely-used explicit (Euler-forward) discretization is only stable for very small time steps. This leads to poor efficiency. Until now, only a few attempts have been undertaken to overcome this problem [1, 2, 4, 9]. Among these are the so-called additive operator splitting (AOS) schemes, which have been proposed in a recent paper by the first author [9]. They are absolutely stable and can be separated into one-dimensional processes, for which a fast recursive algorithm with linear complexity is available. Under typical accuracy requirements they are at least ten times more efficient than explicit schemes.

The present paper describes how their efficiency can be increased by another order of magnitude by exploiting their intrinsic parallelism. We review the basic ideas behind AOS schemes and describe a parallel implementation of an AOS scheme. An example is presented where denoising of a medical data set from 3-D ultrasound is performed within less than one minute.

The paper is organized as follows: Section 2 describes a widely-used and mathematically well-founded prototype of a nonlinear diffusion filter: the model by Catté, Lions, Morel and Coll [3], a spatially regularized Perona–Malik filter [6]. In Section 3 we shall explain the basic ideas behind AOS schemes. The features of our parallel implementation are described in Section 4. It also gives experimental results for a 3-D ultrasound data set. We conclude with a summary in Section 5.

2 The Continuous Filter

In the m-dimensional case the filter of Catté, Lions, Morel and Coll [3] has the following structure:

Let $\Omega := (0,a_1)\times \dots\times(0,a_m)$ be our image domain and consider a (scalar) image $f(x)$ as a bounded mapping from Ω into the real numbers $\mathbb{R}$. Then a filtered image $u(x,t)$ of $f(x)$ is calculated by solving the diffusion equation with the original image as initial state, and reflecting boundary conditions:

$$\partial_t u \;=\; \sum_{l=1}^{m} \partial_{x_l}\Big(g(|\nabla u_\sigma|^2)\,\partial_{x_l} u\Big) \qquad (1)$$

$$u(x,0) \;=\; f(x), \qquad (2)$$

$$\partial_n u|_{\partial\Omega} \;=\; 0, \qquad (3)$$

*current address: DIKU, Universitetsparken 1, DK–2100 Copenhagen East where n denotes the normal to the image boundary $\partial\Omega$.

The "time" t is a scale parameter: larger values lead to simpler image representations. In order to reduce smoothing at edges, the diffusivity g is chosen as a decreasing function of the edge detector $|\nabla u_\sigma|$, where ∇u_σ is the gradient of a Gaussian-smoothed version of u:

$$\nabla u_\sigma \;:=\; \nabla(K_\sigma * u), \tag{4}$$

$$K_\sigma \;:=\; \frac{1}{(2\pi\sigma^2)^{m/2}} \exp\left(-\frac{|x|^2}{2\sigma^2}\right). \tag{5}$$

We use the diffusivity

$$g(s) \;:=\; \begin{cases} 1 & (s \le 0) \\ 1 - \exp\left(\frac{-3.315}{(s/\lambda)^4}\right) & (s > 0). \end{cases} \tag{6}$$

For such rapidly decreasing diffusivities smoothing on both sides of an edge is much stronger than smoothing across it. This selective smoothing process prefers intraregional smoothing to interregional blurring. The factor 3.315 ensures that the flux $\Phi(s) := sg(s)$ is increasing for $|s| \le \lambda$ and decreasing for $|s| > \lambda$. Thus, λ is a contrast parameter separating low-contrast regions with (smoothing) forward diffusion from high-contrast locations where backward diffusion may enhance edges [6]. After some time this filter creates segmentation-like results which are piecewise almost constant. For $t \to \infty$, however, the image becomes completely flat [7]. Well-posedness results for this filter can be found in [3, 7] and a scale-space interpretation in terms of an extremum principle and smoothing Lyapunov functionals is given in [7].

3 AOS Schemes

3.1 Limitations of Conventional Schemes

Let us now consider finite difference approximations to the m-dimensional diffusion filter.

A discrete m-dimensional image can be regarded as a vector $f \in \mathbb{R}^N$, whose components f_i, $i \in \{1, ..., N\}$ display the grey values at the pixels. Pixel i represents the location x_i. Let h_l denote the grid size in the l direction. We consider discrete times $t_k := k\tau$, where $k \in \mathbb{N}_0$ and τ is the time step size. By u_i^k and g_i^k we denote approximations to $u(x_i, t_k)$ and $g(\nabla u_\sigma(x_i, t_k))$, respectively, where the gradient is replaced by central differences.

The simplest discretization of the diffusion equation with reflecting boundary conditions is given by

$$\frac{u_i^{k+1} - u_i^k}{\tau} \;=\; \sum_{l=1}^{m} \sum_{j \in \mathcal{N}_l(i)} \frac{g_j^k + g_i^k}{2h_l^2}\,(u_j^k - u_i^k). \tag{7}$$

where $\mathcal{N}_l(i)$ consists of the two neighbours of pixel i along the l direction (boundary pixels may have only one neighbour). In vector–matrix notation this becomes

$$\frac{u^{k+1} - u^k}{\tau} \;=\; \sum_{l=1}^{m} A_l(u^k)\, u^k. \tag{8}$$

A_l describes the diffusive interaction in l direction. One can calculate u^{k+1} directly (explicitly) from u^k without any matrix inversion:

$$u^{k+1} \;=\; \Big(I + \tau \sum_{l=1}^{m} A_l(u^k)\Big)\, u^k. \tag{9}$$

For this reason it is called *explicit scheme*. Each explicit iteration step can be performed very fast. Unfortunately, each step has to be very small as well: one can show [9] that in order to guarantee stability, the step size has to satisfy

$$\tau \;\le\; \frac{1}{\sum_{l=1}^{m} \frac{2}{h_l^2}}. \tag{10}$$

For most practical applications, this restriction requires to use a very high number of iterations, such that the explicit scheme is rather slow.

Thus, we consider a slightly more complicated discretization next, namely

$$\frac{u^{k+1} - u^k}{\tau} \;=\; \sum_{l=1}^{m} A_l(u^k)\, u^{k+1}. \tag{11}$$

This scheme does not give the solution u^{k+1} directly (explicitly): It requires to solve a linear system first. It is called a *linear-implicit (semi-implicit)* scheme. The solution u^{k+1} is given by

$$u^{k+1} \;=\; \Big(I - \tau \sum_{l=1}^{m} A_l(u^k)\Big)^{-1} u^k. \tag{12}$$

This scheme is unconditionally stable [9].

In the 1-D case the system matrix is tridiagonal and diagonally dominant. For such a system a Gaussian algorithm for tridiagonal systems (also called *Thomas algorithm*) is stable and solves the problem in linear complexity.

For dimensions $m \ge 2$, however, it is not possible to order the pixels in such a way that in the i-th row all nonvanishing elements of the system matrix can be found within the positions $[i, i-m]$ to $[i, i+m]$: Usually, the matrix reveals a much larger bandwidth. Applying direct algorithms such as Gaussian elimination would destroy the zeros within the band and would lead to an immense storage and computation effort. Typical iterative algorithms become slow for large τ, since this increases the condition number of the system matrix. Thus, in spite of its absolute stability, the semi-implicit scheme is often not so much faster than the explicit one.

3.2 AOS Schemes

In order to address this problem we consider a modification of (12), namely the *additive operator splitting (AOS) scheme*

$$u^{k+1} = \frac{1}{m} \sum_{l=1}^{m} \Big(I - m\tau A_l(u^k)\Big)^{-1} u^k. \qquad (13)$$

The operators $B_l(u^k) := I - m\tau A_l(u^k)$ describe one-dimensional diffusion processes along the x_l axes. Under a consecutive pixel numbering along the direction l they come down to strictly diagonally dominant tridiagonal matrices which can be efficiently inverted by the Thomas algorithm.

Moreover, (13) has the same first-order Taylor expansion in τ as the explicit and semi-implicit scheme: all methods are $O(\tau + h_1^2 + ... + h_m^2)$ approximations to the continuous equation.

Since it is an *additive* splitting, all coordinate axes are treated in exactly the same manner. This is in contrast to conventional splitting techniques from the literature, which are *multiplicative* [5]. They may produce different results if the image is rotated by 90 degrees.

Recently a general framework for discrete nonlinear diffusion scale-spaces has been discovered, which guarantees that the discretization reveals the same scale-space properties as its continuous counterpart [8]. One can verify [9] that the AOS scheme creates such a discrete nonlinear diffusion scale-space for every (!) step size τ. As a consequence, it preserves the average grey level μ, satisfies a causality property in terms of a maximum–minimum principle, and converges to a constant steady state. Moreover, the process is a simplifying, information-reducing transform with respect to many aspects: The p-norms

$$\|u^k\|_p := (\sum_{i=1}^{N} |u_i^k|^p)^{1/p} \qquad (14)$$

and all even central moments

$$M_{2n}[u^k] := \frac{1}{N} \sum_{j=1}^{N} (u_j^k - \mu)^{2n} \qquad (15)$$

Table 1: Overview of the different schemes.

scheme	eqn.	stability	costs/iter.	efficiency
explicit	(9)	$\tau < \frac{h^2}{2m}$	very low	low
semi-implicit	(12)	absolute	high	medium
AOS	(13)	absolute	low	high

are decreasing in k, and the entropy

$$S[u^k] := -\sum_{j=1}^{N} u_j^k \, \ln u_j^k, \qquad (16)$$

a measure of uncertainty and missing information, is increasing in k (if f_j is positive for all j).

Table 1 summarizes the features of the discussed schemes. Full algorithmic details of AOS schemes can be found in [10].

4 Parallel Implementation

Since AOS schemes split the whole process in a sequence of 1-D processes, they are intrinsically parallel. Let N_j denote the number of pixels in the j direction. Then $(I - m\tau A_l(u^k))^{-1}$ creates $\prod_{j \neq l} N_j$ one-dimensional diffusion processes along the l direction. These processes are completely independent from each other, so that they can be distributed to the different processors of a parallel computer.

As a demonstrator we have implemented a three-dimensional AOS scheme on an SGI Power Challenge XL (eight 195 MHz R10000 processors with 1 MByte secondary cache). The test image is a 3-D ultrasound data set of size $138 \times 208 \times 138$. It depicts a 10-week old human foetus.

The filter parameters of the AOS scheme were determined as follows: The contrast parameter λ was set to the 40% quantile of the cumulative histogramme for $\nabla(K_\sigma * f)$, and the standard deviation σ of the Gaussian was chosen as the unit length of the cubic voxel. From Section 3 we know that the variance $M_2[u^k]$ is deceasing in k. We may use this Lyapunov functional as a measure for the simplicity of u^k, and stop the process when the relative variance $M_2[u^k]/M_2[f]$ is below a specified number. For our test image 8 iterations with $\tau = 10$ were sufficient to decrease the variance by 5 %. This already gives good denoising, as is seen in Fig. 1. Note the significantly improved visibility of the skull and the hands.

Implementing the parallel AOS scheme was done in two steps: First we have further optimized our existing serial AOS code by using the C++-based dedicated TULIP library. TULIP has been developed by

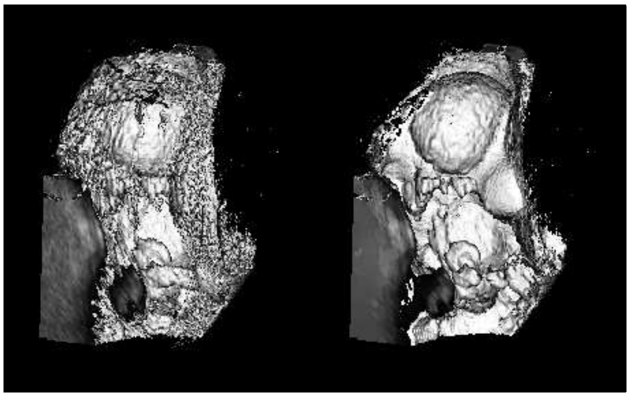

Figure 1: Rendering of a 3-D ultrasound image of a 10-week old foetus. LEFT: Original. RIGHT: Filtered.

Table 2: Execution time for eight AOS iterations on an SGI Power Challenge for an image of size $138 \times 208 \times 138$ using a varying number of processors.

# proc.	exec. time	speed-up	efficiency
1	158.22 s	1.00	100.0 %
2	93.11 s	1.70	85.0 %
3	71.48 s	2.21	73.7 %
4	59.75 s	2.64	66.0 %
5	51.96 s	3.05	61.0 %

Karel Zuiderveld and Fred Appelman (Image Sciences Institute, Utrecht) using cache memory addressing in a sophisticated way in order to speed up the way of accessing neighbouring voxels. This feature has led to a speed-up by a factor three. In a second step the optimized code has been split up into functions acting only within 2-D slices of the 3-D image. These slices were then distributed to the different processors.

Table 2 shows the measured execution times as a function of the number of processors. The execution time of an algorithm on a parallel systems is given by

$$T = T_s + \frac{T_p}{p} + T_{sc}(p), \tag{17}$$

where T_s and T_p are the serial and parallel components, respectively, and $T_{sc}(p)$ is the synchronization time for the parallel processes.

If we neglect $T_{sc}(p)$, we find from Table 2 that $T_s \approx 0.16\,T$, that means about 84 % of the code is parallel. The serial part is mainly caused by I/O from and to the harddisk, and from necessary copies of the volume data, which are hardly parallelizable. For more iterations we may expect further improvement with respect to the scaling behaviour. The measured data indicate that due to parallel implementations, an efficiency increase by one order of magnitude is realistic.

5 Summary

We have illustrated the essential ideas behind AOS schemes by applying them to a regularized Perona–Malik diffusion filter. A parallel implementation of AOS schemes on a few processors has been demonstrated to be sufficient for filtering 3-D medical data sets within their typical acquisition time. This can be considered as an important step towards many challenging real-life applications of these theoretically well-founded methods.

References

[1] S.T. Acton, A.C. Bovik, M.M. Crawford, *Anisotropic diffusion pyramids for image segmentation*, Proc. IEEE Int. Conf. Image Processing (ICIP–94, Austin, Nov. 13–16, 1994), Vol. 3, IEEE Computer Society Press, Los Alamitos, 478–482, 1994.

[2] E. Bänsch, K. Mikula, *A coarsening finite element strategy in image selective smoothing*, Preprint No. 18/1996, Faculty of Mathematics, University of Freiburg, Hebelstr. 27, 79104 Freiburg, Germany, 1996; to appear in Computation and Visualization in Sciences.

[3] F. Catté, P.-L. Lions, J.-M. Morel, T. Coll, *Image selective smoothing and edge detection by nonlinear diffusion*, SIAM J. Numer. Anal., Vol. 29, 182–193, 1992.

[4] T. Gijbels, P. Six, L. Van Gool, F. Catthoor, H. De Man, A. Oosterlinck, *A VLSI-architecture for parallel nonlinear diffusion with applications in vision*, Proc. IEEE Workshop on VLSI Signal Processing, Vol. 7, 398–407, 1994.

[5] G.I. Marchuk, *Splitting and alternating direction methods*, P.G. Ciarlet, J.-L. Lions (Eds.), Handbook of numerical analysis, Vol. I, 197–462, 1990.

[6] P. Perona, J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 629–639, 1990.

[7] J. Weickert, *Theoretical foundations of anisotropic diffusion in image processing*, Computing, Suppl. 11, 221–236, 1996.

[8] J. Weickert, *Nonlinear diffusion scale-spaces: From the continuous to the discrete setting*, M.-O. Berger, R. Deriche, I. Herlin, J. Jaffré, J.-M. Morel (Eds.), ICAOS '96: Images, wavelets and PDEs, Lecture Notes in Control and Information Sciences, Vol. 219, Springer, London, 111–118, 1996.

[9] J. Weickert, *Recursive separable schemes for nonlinear diffusion filters*, B. ter Haar Romeny, L. Florack, J. Koenderink, M. Viergever (Eds.), Scale-space theory in computer vision, Lecture Notes in Comp. Science, Vol. 1252, Springer, Berlin, 260–271, 1997.

[10] J. Weickert, B.M. ter Haar Romeny, M.A. Viergever, *Efficient and reliable schemes for nonlinear diffusion filtering*, IEEE Trans. Image Proc., 1998, to appear.