Recasting Tomo-PIV reconstruction as constrained and L1-regularized nonlinear least squares problem

Sebastian Gesemann[1], Daniel Schanz[1], Andreas Schröder[1], Stefania Petra[2], Christoph Schnörr[2]

1: Dept. of Experimental Fluids, Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR)
{sebastian.gesemann, daniel.schanz, andreas.schroeder}@dlr.de
2: Dept. of Mathematics and Computer Science, University of Heidelberg, Germany
{petra,schnoerr}@math.uni-heidelberg.de

Abstract. Tomo-PIV is a volumetric flow measurement system. It is able to obtain instantaneous 3D-3C velocity fields using multiple cameras that observe the same volume of two subsequently illuminated particle distributions from different angles. Part of the measurement technique is the 3D reconstruction of the particle volume from 2D images. This work presents a promising alternative approach to the popular reconstruction algorithms MART and SMART. Simulations and experimental data of a turbulent free jet at Ma = 0.7 are used in order to assess and demonstrate the performance of the newly developed reconstruction algorithm based on constrained least squares strategies and L1-regularization.

1. Introduction

In Tomo-PIV the reconstruction problem is typically modeled as linear equation system where the volume is discretized (Elsinga et al, 2005). This system relates each unknown "voxel" intensity to a small subset of pixels under the assumption that light intensities accumulate along lines of sight and, typically, that the voxels' intensities are not negative. The relationship between voxels and pixels is completely determined by the system's matrix which encodes geometric as well as spreading function information about the measurement configuration and optical systems involved. In an ideal case this matrix is correctly determined by means of calibration. This linear model is appropriate for small particles since virtually no occlusion occurs. The size of the imaged particles is mainly due to diffraction and refraction effects of the optical system.
Unfortunately, these equation systems are severely underdetermined due to the low number of projections. Typically, only three or four cameras are used to record the illuminated particle volume from different perspectives. However, the nonnegativity constraint for voxel intensities seems to be one of the main reasons why reconstructions work reasonably well at least for low to moderate particle densities. The sheer size of these equation systems poses a problem in terms of computation time and working memory requirements. Iterative solvers are often applied with a low fixed number of iterations instead of iterating until a certain convergence criterion is satisfied. The performance of reconstruction algorithms during the first iterations is therefore an interesting study.
Recently, Petra et al (2008) drew attention to a generic optimization algorithm with interesting properties (superlinear convergence for a certain problem set) for the reconstruction problem called nonmonotone Spectral Projected Gradient Method – "SPG" (Birgin et al, 2000). The reconstruction problem is formulated as a linear least squares problem which accounts for errors in the pixel intensity measurements and constrains voxel intensities to be nonnegative. The least squares approach and superlinear convergence are attractive features of the algorithm as they could increase robustness and reduce reconstruction time.
In this work we investigated and compared the performance of SMART (Mishra et al, 1999) as one of the state-of-the-art reconstruction algorithms for Tomo-PIV (Atkinson, Soria, 2009) and the nonmonotone SPG algorithm through various simulations. We were able to improve the performance of reconstruction via SPG with a nonlinear substitution that warps the search space.

The algorithm that combines this substitution and the SPG solver is referred to as S-SPG throughout the rest of this paper. In addition, we tested the use of the L1-regularization as part of the cost function. The L1-regularization plays an important role in recovering sparse signals from undersampled data (Compressed Sensing paradigm, see Petra et al 2009). The combination of S-SPG with L1-regularization is referred to as L1S-SPG.
Finally, SMART and S-SPG are used to reconstruct 40 x 40 x 6 mm^3 volumes from a real free jet experiment with a magnification factor of approximately 30 voxels per millimeter in each dimension. The images show blurred as well as very sharp particles (aliasing) due to the trade-off between large depth of field and high light intensity. This typically calls for a slight lowpass filter as part of the image preprocessing but it gave us the chance to compare the performance between SMART and the SPG-based solvers under these difficult conditions.

2. Reconstruction algorithms

Given an *M* times *N* weighting matrix *A* with coefficients $a_{i,j}$ and a measurement vector *b* encoding the recorded pixel intensities we seek to compute the discrete volumetric intensity distribution *x* – typically a regular grid of voxels in the measurement volume – so that *Ax* approximates *b*, *x* is reasonably sparse and nonnegative.
The algorithms MART and SMART inherently enforce this nonnegativity constraint due to their multiplicative nature and the restriction of nonnegative weights in *A*. For reference, these are the SMART equations we used:

$$r_i^{(k)} = \frac{b_i}{\sum_{j=1}^{N} a_{i,j} x_j^{(k-1)}}$$

$$x_j^{(k)} = x_j^{(k-1)} \prod_{i, a_{i,j}>0} \left(r_i^{(k)} \right)^{a_{i,j}\mu}$$

Here, the vector $r^{(k)}$ refers to the ratios between measured pixel intensities and current projections in the *k*-th iteration. The update of the volume is done by simultaneously applying a multiplicative corrective term which, for a choice of μ with $\mu \sum_i a_{i,j} = 1$, is a weighted geometric average of the previously computed intensity ratios. Smaller values for μ slow down convergence and higher values tend to produce oscillating sequences of intensity distributions which is why we chose it to satisfy the aforementioned equation.
The goal of the Spectral Projected Gradient method is to find a minimum of an objective function. The obvious choice for an objective function in this case is the sum of squared pixel intensity errors:

$$f(x) = \frac{1}{2} \| Ax - b \|_2^2$$

$$\nabla f(x) = A^T \left(Ax - b \right)$$

To enforce the nonnegativity constraint we can use a projection *P* as part of the algorithm that replaces negative coefficients with zero. The algorithm basically computes gradients of the objective function and derives a descent direction *d* according to the current gradient, a scale factor of $1/\sigma^{(k)}$ where $\sigma^{(k)} I_N$ is supposed to approximate the objective function's Hessian matrix at the current location, and the projection *P*:

$$d^{(k)} = P\left(x^{(k-1)} - \frac{1}{\sigma^{(k-1)}} \nabla f\left(x^{(k-1)} \right) \right) - x^{(k-1)}$$

$$x^{(k)} = x^{(k-1)} + \alpha^{(k)} d^{(k)}$$

$$\sigma^{(k)} = \max\left\{\sigma_{\min}, 0.1\cdot\sigma^{(k-1)}, \frac{< d^{(k)}, \nabla f\left(x^{(k)}\right) - \nabla f\left(x^{(k-1)}\right) >}{< d^{(k)}, \alpha d^{(k)} >}\right\}$$

For $\alpha^{(k)}$ with $0 < \alpha^{(k)} \leq 1$ the next $x^{(k)}$ is guaranteed to have nonnegative coefficients due to the projection *P*. In most iterations $\alpha^{(k)} = 1$ is used – even tolerating a nonmonotone development of the objective function's cost values. However, global convergence can be guaranteed if $\alpha^{(k)}$ is occasionally set to lower values so that the new cost stays below the maximum cost value seen during the last *T* iterations. In our case we used *T = 4*. For the next value of σ the difference between the new gradient and old gradient is taken into account as well as the last step αd. Here, $<\square,\square>$ refers to the scalar product. With a lower bound of $0 < \sigma_{\min}$ the step length is restricted and the algorithm can even navigate nonconvex regions of the objective function. In addition, we prevented new values for σ to go below one tenth of the previous value to reduce the chance of an unnecessarily large descent step *d* during the next iteration. For more details on the nonmonotone Spectral Projected Gradient Method see Birgin et al, 2000.

3. Evaluation criteria

Due to the low number of projections the quality of a reconstruction is not completely determined by the residual alone. Solutions to the reconstruction problem may differ greatly in quality while producing the same images when applied to the weighting matrix *A*. For simulations with known volumetric intensity distribution and correct matrix *A* – correct in the sense that synthetically generated images correspond to *Ax* except for optionally added noise – we not only can check the development of the residual of intermediate solutions but also compare the current iteration's reconstruction with the original volume. To measure the similarity between two intensity distributions we simply used the normalized cross correlation. This cross correlation is sometimes known as "quality factor". The use of the normalized cross correlation seems reasonable as it is expected to be indicative of the correlation peak heights during correlation-based displacement estimation.

Since all algorithms discussed here can be written in terms of matrix vector products – even SMART via *M* additional logarithms and *N* exponentiations per iteration – and these products are dominating the computation time, a fair comparison of computation time can be done by counting the number of matrix vector products each algorithm uses internally. A SMART iteration always uses two of these "costly operations" per iteration – one for the projection and one for the update step. The SPG algorithm requires slightly more of these operations at average due to the occasional step size reduction and recomputation of the residual with another step size.

4. Warping the search space as preconditioning for SPG

After initial tests we noticed the superiority of SMART over SPG when directly applied on the constrained linear least squares problem in two aspects: SMART managed to reduce the residual more quickly during the first iterations in many cases and the intermediate solutions $x^{(k)}$ produced by SMART had higher quality factors compared to the SPG reconstructions at *similar* residual levels. This suggests that in these instances there was no unique solution to find and that the natures of the respective algorithms affected the search paths differently. These results are included in section 5.

An attempt to combine the good features of both algorithms led to a nonlinear substitution of the search space. Instead of applying the SPG directly on the linear least squares problem, we substitute x for z^2 (componentwise, $\forall j, 1 \le j \le N : x_j =: z_j^2$) and let the SPG find the optimal z. This substitution affects the gradients in a way that higher voxel intensities change more quickly than lower voxel intensities just like it is the case with multiplicative reconstruction approaches. However, the objective function it not convex quadratic in z anymore. Turning a constrained linear least squares problem into a constrained nonlinear least squares problem seems counterproductive at first, but the SPG algorithm is still applicable and shows a better convergence behavior as can be seen in the next section. The substitution does not affect the complexity of the implementation in any significant way.

5. Simulations

We simulated discrete particle volumes of 256 x 256 x 128 voxels with varying particle densities and image noise levels. Four virtual cameras observed the volume under parallel projection with viewing directions $(\pm 1;\ \pm 1;\ 2)^T$ and a 1:1 voxel-pixel size ratio. The particle positions were determined by a pseudo random number generator using uniform distributions in the volume including nonzero sub-voxel positions. For each generated particle up to 4^3 voxels have been altered by sampling and adding a three-dimensional cubic B-Spline to the discrete volume. The projection of each voxel with nonzero intensity is done similarly. The three-dimensional voxel coordinate is mapped to a pixel coordinate and a two-dimensional cubic B-Spline is sampled, thus, affecting up to 4^2 pixels per camera. This results in a projection matrix A where each of the 2^{23} columns contains up to 4^3 nonzero entries which always sum up exactly to the number of cameras due to the partition of unity property of B-Splines. We chose the SMART parameter $\mu = 1/4$ accordingly.

As an initial guess for all algorithms we computed $x^{(0)}$ to be

$$x_j^{(0)} := s\left(\min_c \left\{ b^T A_{\cdot,j}^{(c)} \right\} + \varepsilon\right),$$

where the scale factor s minimizes $\left\| Ax^{(0)} - b \right\|_2^2$,

and $A = \begin{bmatrix} A^{(1)} \\ \vdots \\ A^{(4)} \end{bmatrix}$.

$A^{(1)}$, $A^{(2)}$,…, $A^{(4)}$ are the respective submatrices for all four cameras and ε is a small positive number to prevent initializing a voxel intensity with zero.

5.1 Simulation of various particle densities

Four cases have been simulated: 3 000, 7 000, 10 000 and 15 000 random particles have been rendered into a volume and projected to four image planes. Approximately 79 000 pixels are potentially affected by particles in each camera. This corresponds to particle per pixel (ppp) values of approximately 0.038, 0.089, 0.127, and 0.190, see Fig. 1 for a combination of all four cases in one image of the first virtual camera. The tests have been repeated with different initialization parameters for the random number generator to compute representative residual and quality factor curves against the number of costly operations.

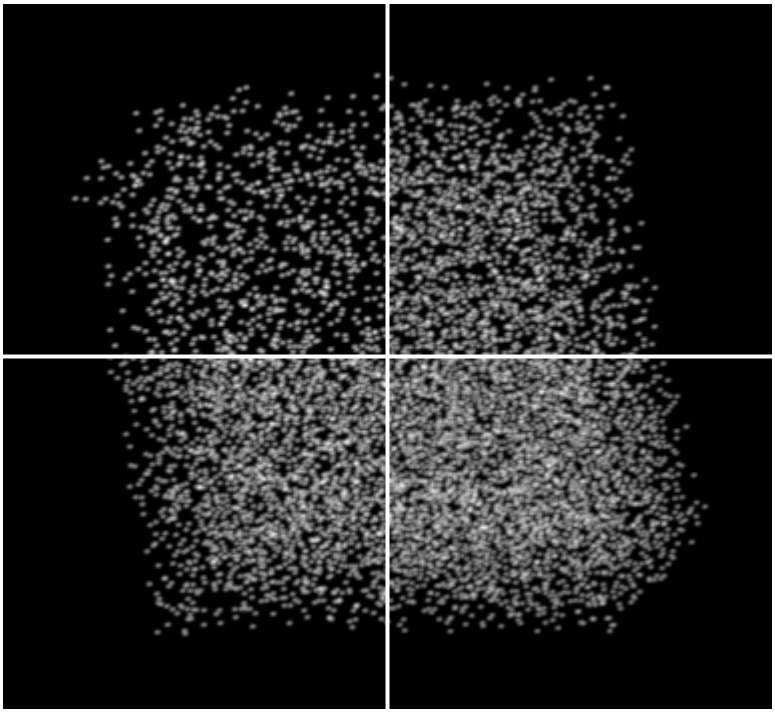

Fig. 1 Simulated first camera showing ppp = 0.038 (top left), ppp = 0.089 (top right), ppp = 0.127 (bottom left), ppp = 0.190 (bottom right)

Fig. 2 through Fig. 5 show the average results of the simulations. In Fig. 2 and Fig. 3 we can clearly see that SMART tends to outperform both of the alternative reconstruction approaches during the first couple of iterations. After about 10 costly operations in the first case with few particles, S-SPG and SMART share the same slope of the residual curve. SPG applied on the original least squares problem (red curve) has a much weaker residual curve slope. The difference to S-SPG is apparent. The residual level of 10^2 is reached by SMART after 10 costly operations, by S-SPG after 16 costly operations while SPG reaches this level after 65 costly operations. But also the quality factors of the reconstructed intermediate volumes at this residual level differ greatly. For a residual of 10^2 SPG reconstructs an intensity distribution with a quality factor below 0.7 while the quality factors of the reconstructed volumes for both of the other algorithms are above 0.9. SMART and S-SPG inherently picked solutions that are closer to the original than the reconstructions of SPG at the same residual level.

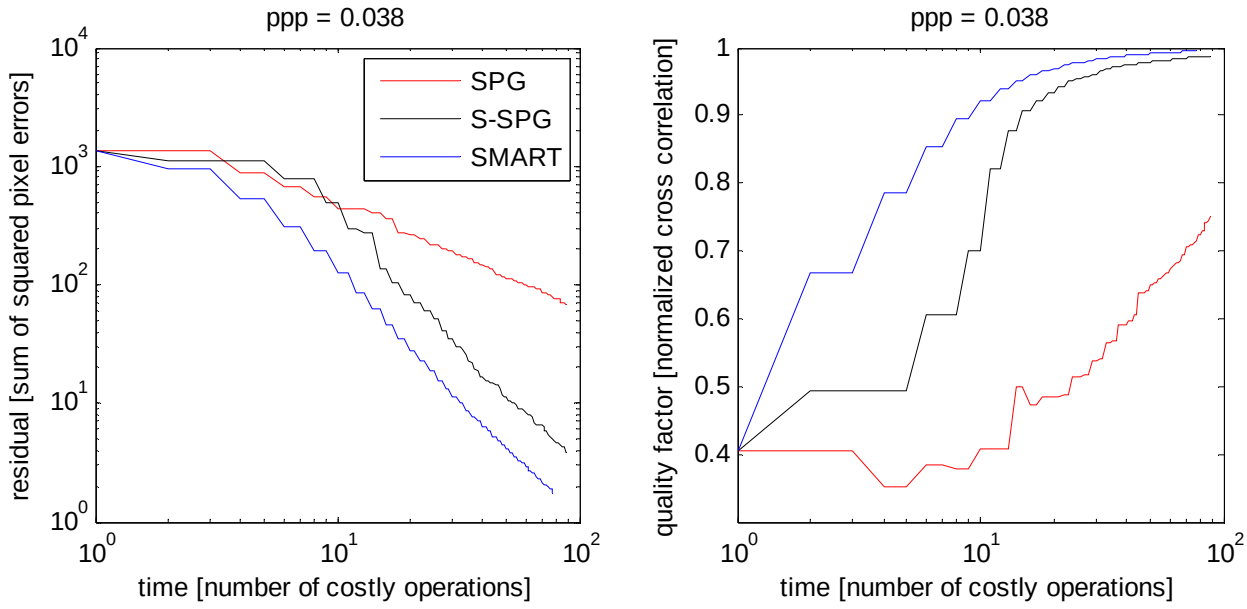

Fig. 2 Results for the first case with 3 000 particles and 40 iterations, residual (left) and quality factor (right)

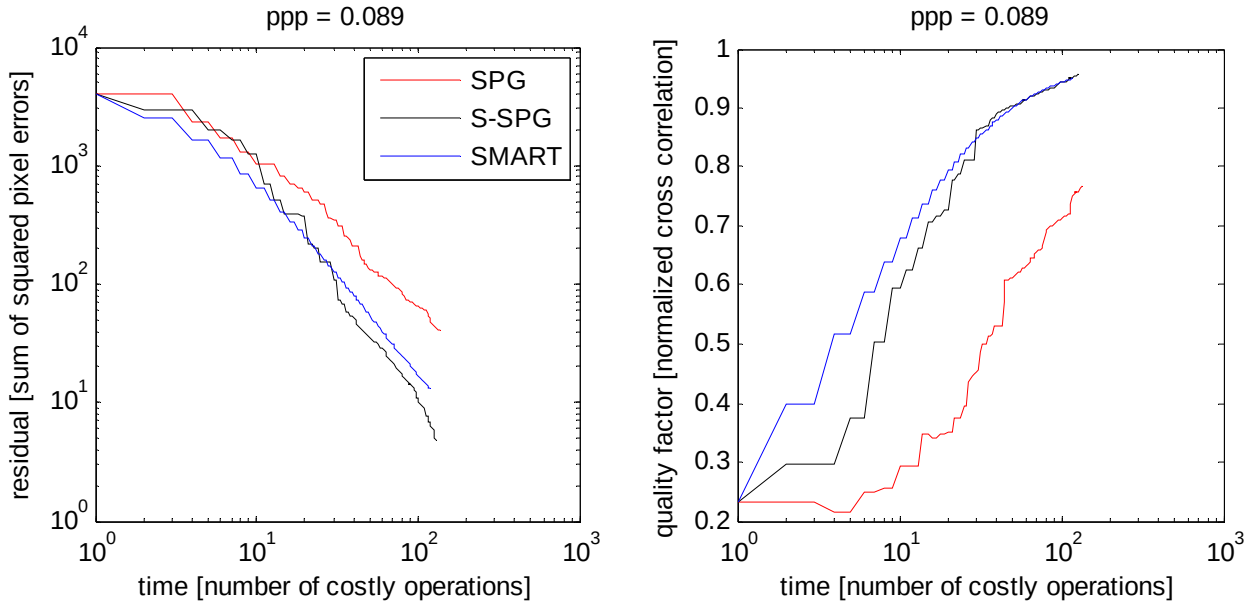

Fig. 3 Results for the second case with 7 000 particles and 60 iterations, residual (left) and quality factor (right)

In the second case with 7000 particles (ppp = 0.106) S-SPG starts slow but can catch up with SMART after about 20 costly operations. After 30 costly operations (15 SMART iterations), S-SPG manages to outperform SMART by approximately 50% in terms of residual minimization. The quality factors for volumes derived by SMART and S-SPG are comparable after that time. SPG is still behind in terms of the residual as well as the quality factor.

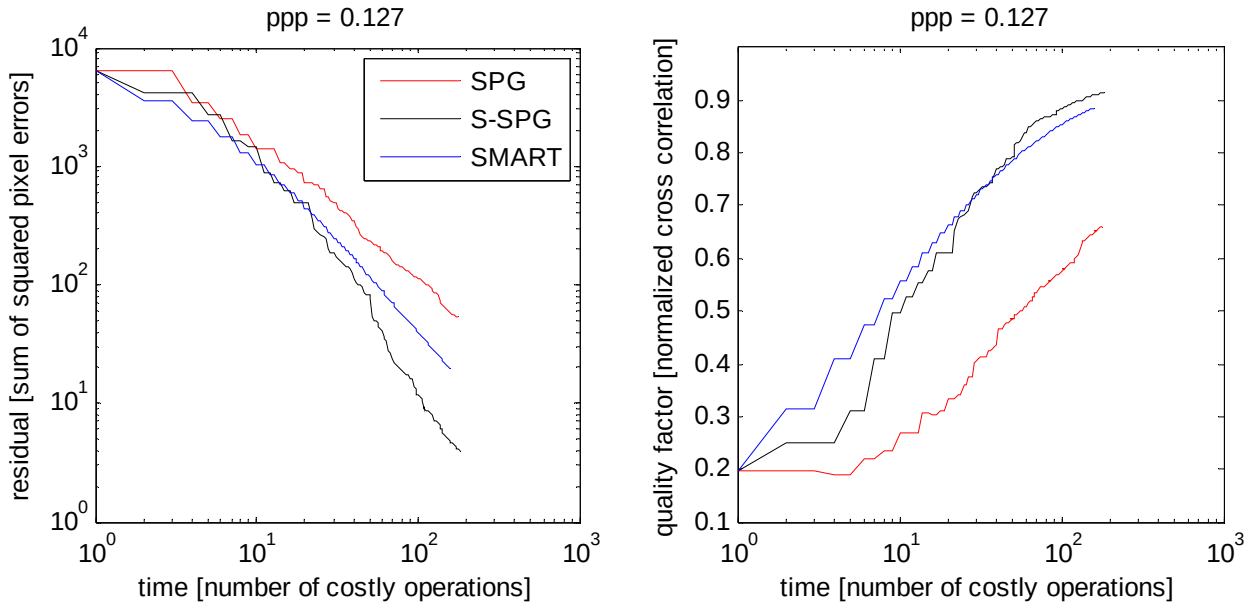

Fig. 4 Results for the third case with 10 000 particles and 80 iterations, residual (left) and quality factor (right)

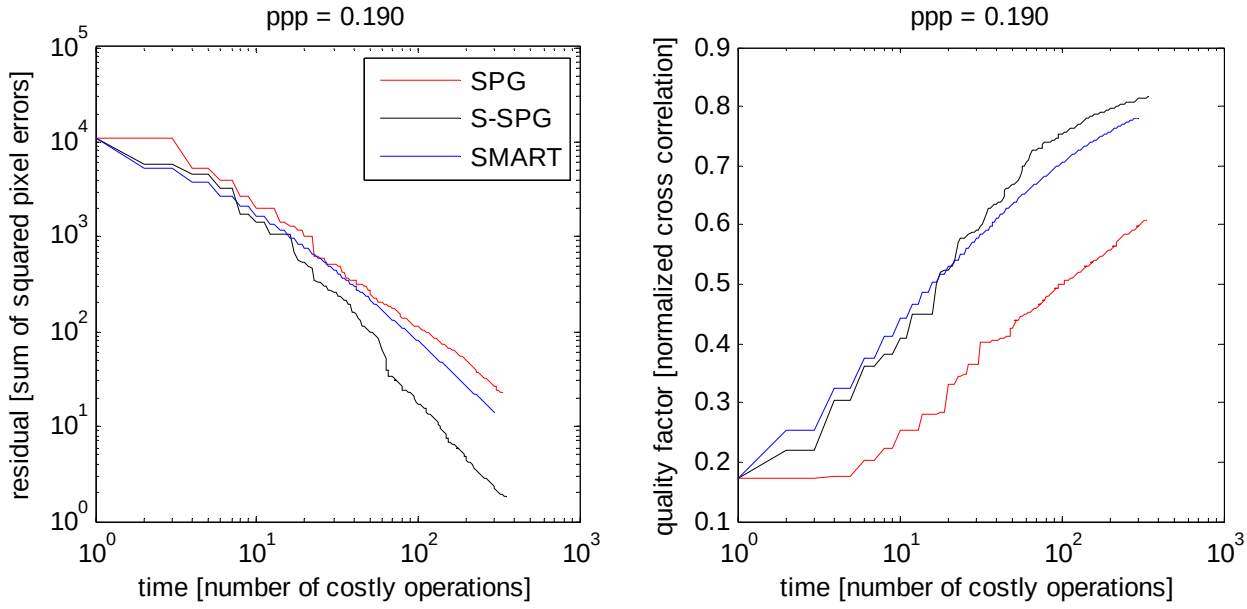

Fig. 5 Results for the fourth case with 15 000 particles and 150 iterations, residual (left) and qualify factor (right)

With increased particle density S-SPG extends its lead in terms of residual minimization and quality factors of reconstructed volumes. Even though SPG (applied on the original least squares problem) catches up in terms of residual, it computes volumes of lower quality. At such high particle densities there probably is no unique solution for the minimization problem. Nevertheless SMART and S-SPG manage to reconstruct volumes that are remarkably close to the original.

5.2 Simulation of noisy images

An additional L1-regularization of the objective function *f* did not seem to help much in the simulated cases with noise-free images. But we noticed a slight quality improvement of the reconstructed volumes when it was applied on problem instances with noisy images. L1-regularization introduces a new parameter λ which controls the effect of the sum of voxel intensities on the cost value of the objective function. In the following formulas 1_N refers to a column vector containing *N* entries of all ones.

$$f(x) = \frac{1}{2}\left\|Ax - b\right\|_2^2 + \lambda 1_N^T x$$
$$\nabla f(x) = A^T\left(Ax - b\right) + \lambda 1_N$$

A simulation with 15 000 particles in a 256 x 256 x 128 volume has been repeated six times with different initializations of the pseudo random number generator. Fig. 6 shows one of the images for the first camera. Each projected particle contributes exactly an intensity of one to the sum of pixel intensities for one camera. The generated noise was Gaussian with a standard deviation of 0.05. Pixels with negative intensity have been set to zero prior to reconstruction.

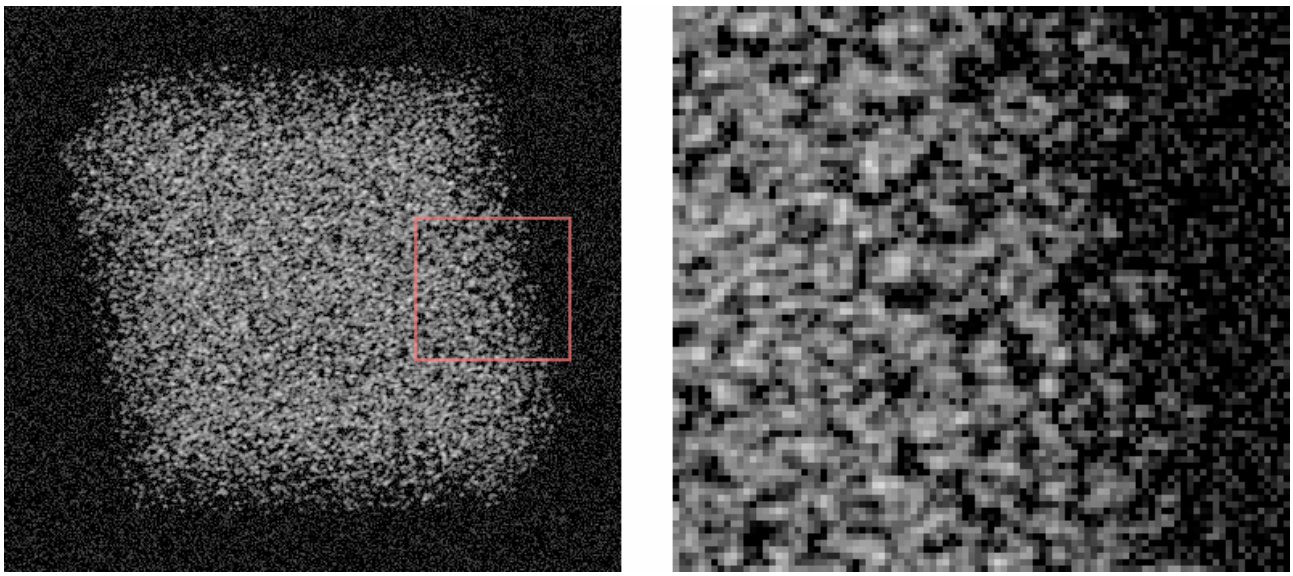

Fig. 6 Simulation with 15 000 particles (ppp = 0.190), images with Gaussian noise (standard deviation of 0.05)

Fig. 7 shows again the development of the residuals and quality factors over time for three algorithms. This time, SPG has been replaced with L1S-SPG. Due to the penalization of high voxel intensity sums via L1-regularization L1S-SPG allows the residual to be higher if the L1-norm can be reduced to a certain extent. This can be seen on the left side of the figure. S-SPG reconstructed volumes after 50 costly operations have a lower residual than corresponding volumes reconstructed by L1S-SPG. But we can see on the right side that this L1-regularization actually improves the quality factor of the reconstructed volumes slightly. The regularization parameter λ was chosen to be 0.1.

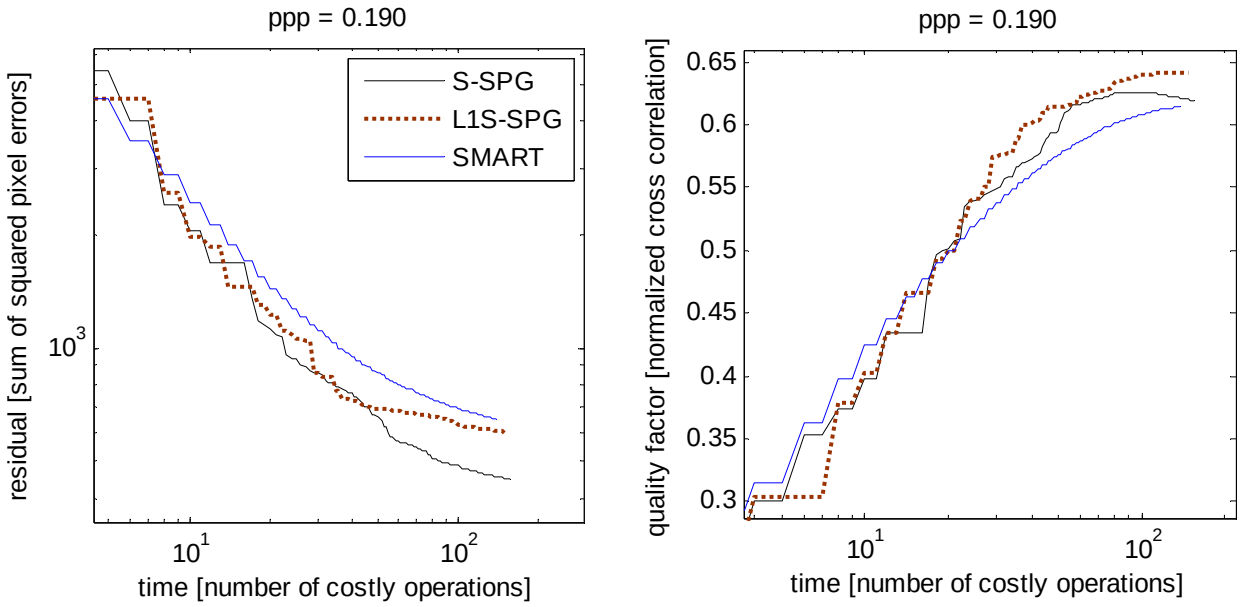

Fig. 7 Results for a simulation with noisy images and 70 iterations, residual (left) and quality factor (right)

6. Application on real experimental data

To verify the suitability of the new algorithm in practice we prepared a Tomo-PIV experiment measuring a free turbulent jet at Ma = 0.7. Four high resolution *PCO 4000* cameras equipped with f = 100 mm *Zeiss* lenses arranged in a pyramidal viewing set-up observing an investigation area of 120 x 50 x 5 mm^3 that is located in a region between x/d = 18 to 22 downstream a round nozzle of 15 mm diameter. As tracer particles 1µm diameter DEHS droplets were used which were illuminated by two combined *BigSky CFR200* double cavity lasers with 400 mJ pulse energy each in an extended and collimated laser beam, which was back-reflected in itself using a mirror. The illuminated volume was confined using two 50 x 5 mm^2 knife-edges, see Fig. 8. Several typical imaging problems especially when using Tomo-PIV in air flows have been modeled experimentally. Due to the change of the local light scattering angle a relatively strong intensity variation of the particle images can be observed along the field of views. Several aperture numbers between $f_\# = 2.8$ and 8 have been tested in order to change the dynamic range of particle image intensities and densities. At the same time this $f_\#$ range produces a wide variation of particle image diameters and introduces background noise due to limited depth-of-focus. Reconstructions with self-calibration will be compared in order to assess the robustness of the individual algorithms.
Fig. 9 shows one of the images recorded by the first camera where parts not affected by the discretized volume have been masked black. We can clearly see the shape of the volume and some slight perspective distortion. The volume that is selected here is 40 x 40 x 6 mm^3 which is about a millimeter thicker than the light sheet. In the close-up on the right side we can see some very sharp particles (aliasing) as well as blurred particle images.

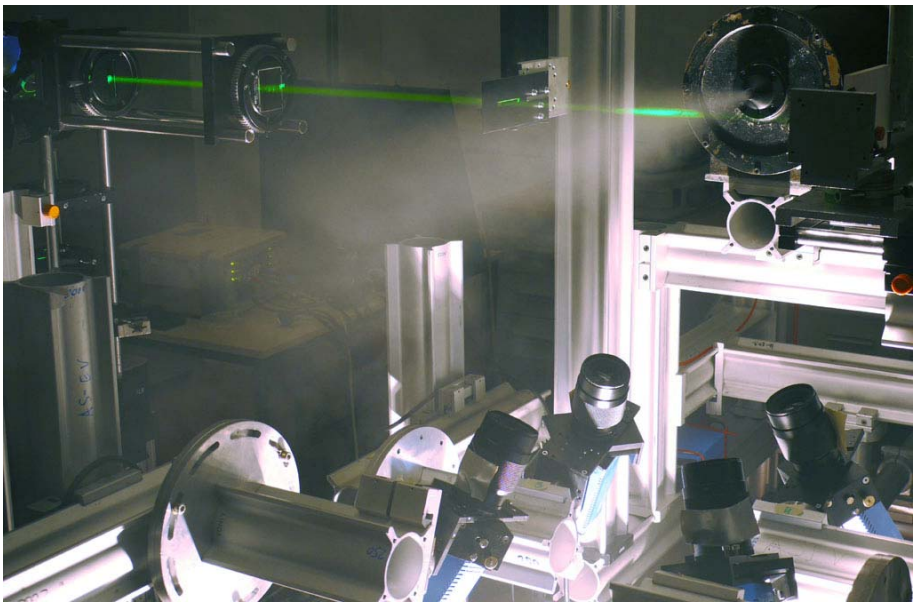

Fig. 8 Experimental setup: light sheet (top), nozzle (top right), four cameras in a pyramidal setup (bottom right)

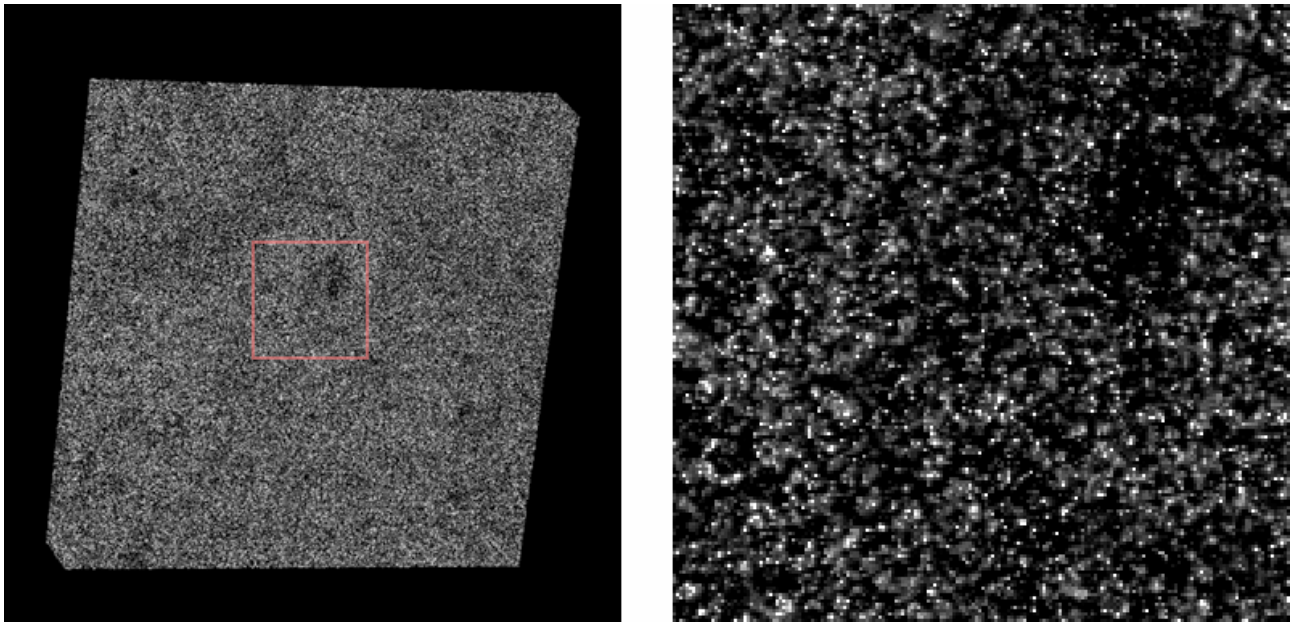

Fig. 9 Masked image of the first camera, overview (left) and close-up (right)

With no further image processing such as smoothing SMART tends to remove such sharp particles since the black surroundings of these isolated bright pixels lead to multiplicative corrective terms close to zero and this in turn sets many voxel intensities near the respective lines of sight irrevocably to zero. The least squares based approach on the other hand computes an intensity distribution which, when projected back to the image plane according to the projection matrix *A*, approximates the original image within the limits of *A*. This can be seen in Fig. 10. While the projection of the reconstructed volume using SMART is slightly less bright and missing some particles, the projection of the reconstructed volume using S-SPG preserves the overall brightness and particles.

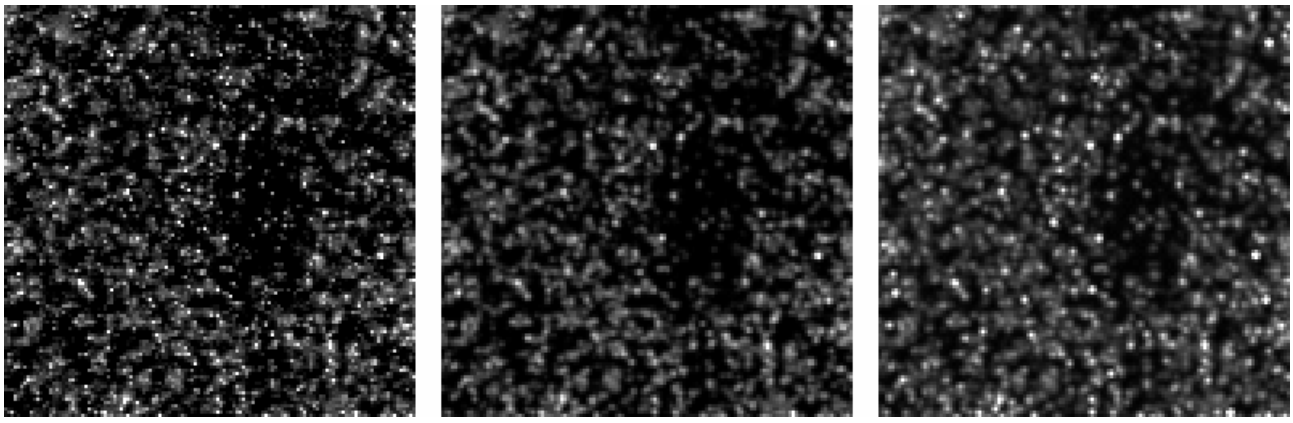

Fig. 10 Original recording (left), projection of SMART reconstruction (center), projection of S-SPG reconstruction (right), in both cases after 20 iterations

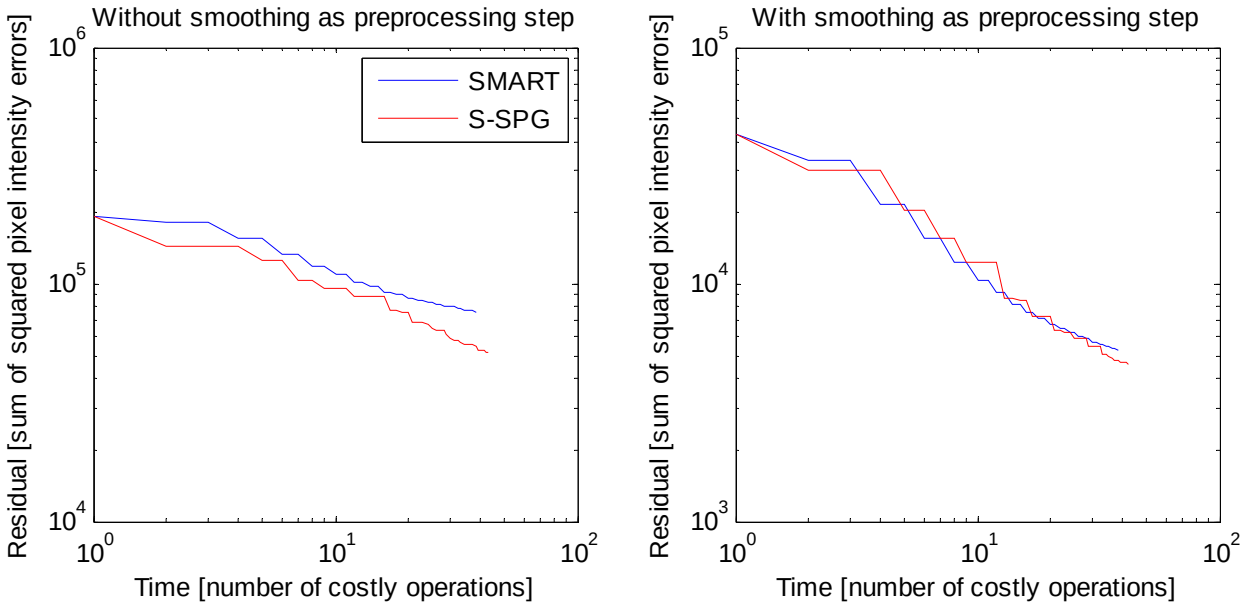

Fig. 11 Residual curves (average over two frames) for both preprocessing cases and algorithms

This effect can also explain the development of the residuals over time, see Fig. 11. In the first case (no smoothing as preprocessing step, left) the S-SPG approach is clearly superior in terms of residual minimization. This can be attributed to the “aggressive” behavior of SMART when it comes to computing and applying the correction terms which seems to be counterproductive in cases where the system of equations is inconsistent.
In the preprocessing case this is less pronounced possibly due to more nonzero pixel intensities surrounding particles. The performances of SMART and S-SPG are similar. There is hardly a noticeable difference between the projections of both reconstructed volumes in this case.

7. Conclusions

Tomographic reconstruction is a vital part of the Tomo-PIV flow measurement technique. High quality reconstructions are just as desirable as fast reconstructions. In this work we presented alternatives to MART and SMART with promising results. Our L1S-SPG implementation requires about twice as much working memory compared to SMART-based solvers but gives better results with respect to our evaluation criteria after a few iterations in cases of high particle density or noisy images. For low particle densities the SMART-based algorithms perform better during early iterations which suggests that a combination of both, a SMART variant for initial iterations followed by L1S-SPG iterations, can be beneficial.

S-SPG and L1S-SPG are also applicable when optical transfer functions have been calibrated to optimize the weights of the equation system (see Schanz et al, 2010). The use of altered weights affects the condition number of the linear equation system and in cases of blurry imaging of particles the condition number is expected to increase. This typically slows down convergence of iterative solvers. In preliminary tests with these kinds of equation systems the SPG-based solvers managed to reduce the residual more quickly than SMART which can be attributed to SPG exploiting information about the objective function's curvature.

We would also like to point out that the generality of the optimization algorithms allows the inclusion of more sophisticated regularizations as well as new degrees of freedom of a more realistic imaging model. For example, additional parameters could be introduced to compensate for unaccounted inter-camera intensity differences or sub-pixel shifts due to vibrations of the optical system.

References

Elsinga G, Scarano F, Wieneke B, van Oudheusen B W (2006) Tomographic Particle Image Velocimetry. Exp Fluids 41:933–947

Mishra D, Muralidhar K, Munshi P (1999) A Robust MART Algorithm For Tomographic Applications. Numerical Heat Transfer, Part B: Fundamentals, 35:4, 485-506

Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image Velocimetry. Exp. Fluids, 47:563-578

Birgin E G, Martinez J M, Raydan M (2000) Nonmonotone spectral projected gradient methods on convex sets. SIAM J. Optimiz. 10:1196-1211

Petra S, Popa C, Schnörr C (2008) Enhancing Sparsity by Constraining Strategies: Constrained SIRT versus Spectral Projected Gradient Methods. In: 7th Workshop on Modelling of Environmental and Life Sciences Problems (WMM 08), Ed Acad Romane, Bucharest, Romania

Petra S, Schröder A, Schnörr C. (2009) 3D Tomography from Few Projections in Experimental Fluid Dynamics. In: Imaging Measurement Techniques in Flow Analysis Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 106. Springer-Verlag Berlin Heidelberg, pages 63-72, ISBN 978-3-642-01105-4

Petra S, Schnörr C (2009) TomoPIV meets compressed sensing. University of Heidelberg, Germany, to appear in Pure Mathematics and Applications: Special issue on Discrete Tomography

Schanz D, Gesemann S, Schröder A, Wieneke B, Michaelis D (2010) Tomographic reconstruction with non-uniform optical transfer functions. Proceedings of 15th Int Symp on Applications of Laser Techniques to Fluid Mechanics, paper #1709 (submitted)