Fast volume reconstruction for 3D PIV

Noname manuscript No.
(will be inserted by the editor)

Fast Volume Reconstruction for 3D PIV

Abhishek Bajpayee · Alexandra H. Techet

the date of receipt and acceptance should be inserted later

Abstract Presented is a memory-efficient and highly parallelizable method for reconstructing volumes, based on a homography fit (HF) synthetic aperture (SA) refocusing method. This technique facilitates rapid processing of very large amounts of data, such as that recorded using high speed cameras, for the purpose of conducting 3D particle imaging velocimetry (PIV) and particle tracking velocimetry (PTV).

Keywords synthetic aperture imaging · light field imaging · homography fit method · particle image velocimetry

1 Introduction

1.1 Motivation

Given the unsteady, three -dimensional nature of fluid flows encountered in both research and industry, there is significant demand for advanced experimental techniques that fully resolve velocity fields in time and space. PIV has been widely and successfully used to resolve 2D velocity fields. However, 2D velocity fields come with certain limitations. For example, since 2D velocity fields do not provide out of plane velocity gradient information, a pressure field calculated from a 2D velocity field contains errors. 2D fields are often insufficient for studying complex, non-symmetric flows such as highly 3D, turbulent flows.

Tomographic PIV (TomoPIV), pioneered by Elsinga et al. [7], is one widely adopted 3D PIV technique that has been used in a wide range of fluids experiments. The TomoPIV technique relies on images from about 4 or 5 cameras looking at a scene from different viewpoints to facilitate tomographic

Massachusetts Institute of Technology, Cambridge, MA 02139, USA
E-mail: ab9@mit.edu, ahtechet@mit.edu

reconstruction of the investigation volume using a multiplicative algebraic reconstruction technique (MART). Synthetic aperture (SA) PIV (SAPIV) was introduced by [4] blending light field imaging concepts into a 3D quantitative velocimetry method. SAPIV uses an array of up to 10 cameras and can spatially resolve densely seeded velocity fields. Because SA imaging facilitates focusing at any arbitrary depth in a volume of interest, accurate intensity volumes could be reconstructed for use in 3D cross correlation based PIV. However, the original implementation of SA reconstruction took considerable computational time.

In an entire 3D PIV process, reconstruction often accounts for a significant portion of the total computation time (see figure 1). Since TomoPIV was introduced in 2006 [7], several improvements to tomographic reconstruction time have been made by Worth et al (MFG-MART) [11], Atkinson et al (MLOS-SMART) [1], Discetti et al (MLOS-MR-SMART) [6] and Lynch et al (SMTE-MART) [9] in 2008, 2009, 2012 and 2015 respectively. However, reconstruction time still remains a significant hurdle to processing large data sets, often acquired using high speed cameras. In addition, computation time for SA reconstruction is in the same of order of magnitude as that for tomographic reconstruction. This drawback and the need for processing large data sets motivates the development of the HF method.

This paper introduces a memory-efficient, faster reconstruction technique, based on a new homography fit (HF) algorithm, for synthetic aperture (SA) refocusing. This new algorithm is highly parallizable and can be implemented readily on an off-the-shelf graphics processing unit (GPU) for further improvement in computational speed.

1.2 Synthetic Aperture Imaging and 3D PIV

Synthetic aperture (SA) imaging projects images of a scene captured from multiple cameras (different view points) onto a virtual focal plane. SA imaging, in a way, simulates images captured via a lens of an arbitrary sized aperture. In principle, as the aperture size of a lens increases its depth of field becomes smaller. SA imaging allows one to simulate a lens with an arbitrarily large aperture, which would not be feasible to manufacture, using an array of multiple cameras looking at the same scene from slightly different viewpoints. By increasing the distance between cameras, the depth of field of such an array can be reduced sufficiently such that it can be used to image particles in a flow field essentially lying only on the focal plane of the calculated images. Particles away from the focal plane are significantly blurred and can be removed owing to the characteristic discrete blur, as discussed by Bajpayee [2]. This allows one to reconstruct the entire volume of interest, in a particle laden tank in order to conduct 3D PIV, as demonstrated by Belden [4] or PTV as demonstrated by Bajpayee [2].

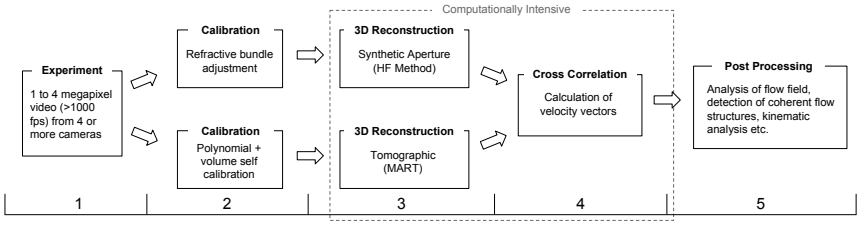

Fig. 1: Schematic outlining the various steps in the PIV process. As indicated, steps 3 and 4 are the most computationally intensive and this paper addresses the acceleration of step 3.

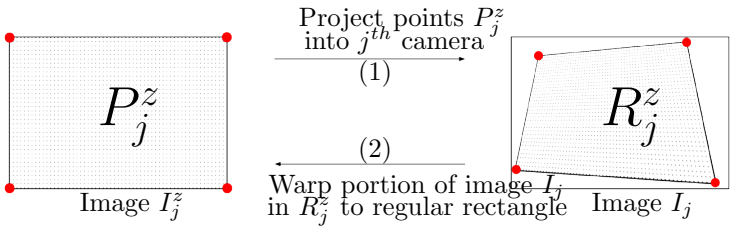

Fig. 2: Schematic outlining the back projection process used for synthetic aperture refocusing. The matrix H for the homography fit method is calculated such that $P_j'^z = HR_j'^z$ where $P_j'^z$ and $R_j'^z$ are the corner points of sets P_j^z and R_j^z respectively.

2 Synthetic Aperture Refocusing

Once a multi-camera system looking at the same scene has been calibrated, images from each camera can be used to calculate images that are sharply focused at arbitrary depths in the scene spanning the entire volume of interest. This is achieved by back projecting images from each camera to a given depth followed by merging these images. This process is called SA refocusing. The SA refocusing process is covered in more detail by Isaksen et al [8] and Vaish et al [10] and later formalized for PIV reconstruction purposes as the map-shift-average algorithm by Belden et al [4]. SA refocusing, when there are no refractive interfaces between the objects of interest and the cameras, is computationally inexpensive. However, refractive interfaces, such as a tank wall and water interface, are present in most PIV/PTV experiments, and thus refractive SA refocusing must be used.

2.1 Refractive SA Refocusing

Refractive SA refocusing takes into account the bending of light rays at each refractive interface when back projecting images. This would correspond to a modified map-shift-average algorithm in which the map and shift operations are nonlinear. Multi-camera systems in scenes that contain refractive interfaces can be calibrated to a sub-pixel accuracy using a refractive bundle adjustment process outlined by Belden [3] which is also able to take lens distortion into account.

To refocus in a case with refractive interfaces, let P_j^z be a set of points representing the location of each pixel of the final back projected image I_j^z and let I_j be the image recorded by the j^{th} camera for the time step being considered. Note that since we know the z coordinate of the back projected image we wish to calculate, we know[1] the points in set P_j^z. In order to implement the desired back projection, the points in P_j^z are projected into the j^{th} camera through the water and glass to a set of points R_j^z using the same method that is used to project a 3D point into a camera during the calibration process [3]. Points in R_j^z are then used to calculate I_j^z using an inverse mapping function via which the portion of image I_j contained in the convex hull of R_j^z is warped to a regular rectangle of the same size as the image to obtain I_j^z. A schematic of this process is shown in figure 2. This is repeated for all cameras and then the back projected images can either be combined by averaging as:

$$I_z = \frac{1}{N} \sum_j I_j^z \tag{1}$$

or by multiplying as:

$$I_z = \prod_j (I_j^z)^\alpha \tag{2}$$

to calculate the refocused image I_z at depth z. In equations 1 and 2, N is the number of cameras in the array and α is an exponent each image is raised to prior to multiplying.

2.2 Fast Refractive Refocusing: Homography Fit Method

The reconstruction step in all PIV experiments is preceded by a calibration step. This takes the form of polynomial calibration and volume self calibration for tomographic PIV and refractive bundle adjustment for SAPIV. Furthermore, calibration has to be conducted only once for an experiment and takes a

[1] The z depth we wish to calculate the refocused image at is the z depth in the calibrated world coordinate system. Note that conversion of P_j^z from world units to pixel units (such that the extent of P_j^z in pixel units is the size, in pixels, of the refocused image being calculated) and vice versa will be required as needed.

few minutes. The reconstruction and cross correlation steps are the most computationally intensive and the technique discussed herein addresses improving performance of step 3, as shown in figure 1.

SA reconstruction is computationally intensive because back projection through refractive interfaces involves an iterative step for each pixel in order to ensure that Snell's law is satisfied for each ray. During the calibration process, the glass wall of the tank is assumed to have flat interfaces. Therefore, all light rays are being refracted at planar interfaces. For a simple scene without refractive interfaces, the back projection process can be reduced to a homography transform since points in one image are mapped to another such that straight lines remain straight (property of projective transforms in the absence of distortion). It can be shown that collinear 3D points when refracted through planar interfaces and then projected into a camera remain collinear in the image. As a result, the projection of 3D points into a camera through planar refractive interfaces can also be approximated with a simple projective transform. The back projection step can thus be achieved by reducing the required inverse mapping to a simple homography transform, as in the case of pinhole back projection. A homography matrix H for this approximation can be computed using a minimum of 4 non collinear points in R_j^z.

Let $P_j'^z$ be a subset of P_j^z containing only the corner points of P_j^z (shown in figure 2 as red dots). Similarly, let $R_j'^z$ be a set containing the corner points of R_j^z. The points $R_j'^z$ can be calculating by projecting the points $P_j'^z$ in a similar fashion as the original method. Following which, a matrix H can then be calculated such that:

$$P_j'^z = H R_j'^z \tag{3}$$

and H can then be used to apply a homography transform to image I_j to calculate I_j^z.

This homography fit (HF) method drastically reduces the computational cost as now only 4 points have to be projected through refractive interfaces instead of as many points as the number of pixels (512×512 in this study) in the images being used. It must be noted that, as long as the experimental setup contains only planar refractive interfaces, the HF method is an exact simplification of the original SA reconstruction method. Therefore, it does not affect the reconstruction quality Q which will follow the same trends as reported by Belden et al [4] for all reconstructions in this study. Most PIV experiments contain either flat refractive interfaces or negligible refractive effects through index matching, flat view boxes etc. For experiments containing non-planar refractive interfaces, the HF method will not be valid.

2.3 Benchmarking of the HF method

The code used for SA reconstruction in this study is implemented in C++ and also written to be parallelized on CUDA capable GPUs. Table 1 lists code benchmarking results, on 2 different machines with and without using

Table 1: Time taken to reconstruct a $512 \times 512 \times 128$ voxels volume; Desktop specs: Intel Core i7 CPU (2.93 GHz), NVIDIA GeForce GTX 650 Ti GPU; Server specs: Intel Xeon CPU (2.6 GHz), NVIDIA Titan X GPU; None of the code is CPU parallelized so all results shown utilize a single CPU core.

	Original Method	HF Method	Speedup (over original method)
Desktop: CPU only	14.19 mins	4.09 s	208x
Server: CPU only	13.50 mins	7.41 s	109x
Desktop: CPU + GPU	1.14 mins	0.45 s	152x
Server: CPU + GPU	8.13 s	0.38 s	21x

a GPU, for reconstructing (and thresholding) a single $512 \times 512 \times 128$ voxels volume. Regardless of whether a GPU is used or not, it can be seen that the HF method significantly improves computation time and the use of a GPU improves computation time by another order of magnitude.

Table 2: Reconstruction computation time (T) to reconstruct a $512 \times 512 \times 128$ voxels volume with seeding density $N_{ppp} = 0.1$ on the a desktop (specs reported in table 1) machine.

Hardware	Algorithm	T [min]
CPU only	MFG-MART, 5 iterations	2.06
CPU only	HF method	0.068
CPU + GPU	HF method	0.0075

In order to compare the preformance of the HF method with tomographic reconstruction, the authors ran tomographic reconstruction simulations using their own extension of code developed by Clark [5] written in FORTRAN and MATLAB. Table 2 shows time to reconstruct a $512 \times 512 \times 128$ voxels volume with seeding density N_{ppp} of 0.1 using both the HF method and MFG-MART based tomographic reconstruction. As shown in table 2, the best SA reconstruction case (desktop with off-the-shelf GPU) translates to a 270x improvement over tomographic reconstruction.

Similar to HF method based SA reconstruction, the time taken for tomographic reconstruction for constant N_{ppp} grows linearly with volume depth. It must be noted that tomographic reconstruction time also grows with N_{ppp} whereas SA reconstruction time is independent of N_{ppp}. Moreover, the timing results shown in tables 1 and 2 are obtained using 32 bit floating point images. For perspective, consider an experiment conducted using high speed cameras with seeding density N_{ppp} of 0.1, reconstructed volume size of $2048 \times 2048 \times 512$

voxels and 1000 time frames of interest. Tomographic volume reconstruction of all frames would take about 90 days as opposed to the HF method (with GPU) which would take 8 hours. Therefore, the presented technique can make the volume reconstruction time insignificant compared to time taken by cross correlation.

3 Conclusion

This paper shows that HF method based SA imaging can reconstruct volumes in a significantly shorter time as compared to the time taken by algorithms used in tomographic techniques. Such an increase in computation speed can allow researchers to process over two orders of magnitude more data than that which can be practically processed currently. Thus moving closer to real-time implementations of high speed 3D PIV, in highly unsteady fluids applications.

Additionally, since the HF method is a projection technique, it can be used to simplify mapping functions for tomographic reconstruction to reduce computational complexity and eliminate the need for volume self calibration for a setup calibrated using refractive bundle adjustment.

All code used in this study has been released under the Open Source Flow Visualization (OpenFV, openfv.org) library which also includes support for tomographic reconstruction and window deformation PIV. The authors encourage other researchers to use the library and contribute improvements and features to help facilitate free and open access to software in the PIV and flow visualization community.

Acknowledgements Funding for this study was provided by the University of Michigan Naval Engineering Education Center.

References

1. Callum Atkinson and Julio Soria. An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids*, 47(4-5):553–568, 2009.
2. Abhishek Bajpayee and Alexandra H Techet. 3d particle tracking velocimetry (ptv) using high speed light field imaging. In *PIV13; 10th International Symposium on Particle Image Velocimetry, Delft, The Netherlands, July 1-3, 2013*. Delft University of Technology, Faculty of Mechanical, Maritime and Materials Engineering, and Faculty of Aerospace Engineering, 2013.
3. Jesse Belden. Calibration of multi-camera systems with refractive interfaces. *Experiments in fluids*, 54(2):1–18, 2013.
4. Jesse Belden, Tadd T Truscott, Michael C Axiak, and Alexandra H Techet. Three-dimensional synthetic aperture particle image velocimetry. *Measurement Science and Technology*, 21(12):125403, 2010.
5. Thomas Henry Clark. *Measurement of three-dimensional coherent fluid structure in high Reynolds number turbulent boundary layers*. PhD thesis, University of Cambridge, 2012.
6. Stefano Discetti and Tommaso Astarita. A fast multi-resolution approach to tomographic piv. *Experiments in fluids*, 52(3):765–777, 2012.

7. GE Elsinga, F Scarano, B Wieneke, and BW Van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6):933–947, 2006.
8. Aaron Isaksen, Leonard McMillan, and Steven J Gortler. Dynamically reparameterized light fields. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 297–306. ACM Press/Addison-Wesley Publishing Co., 2000.
9. KP Lynch and F Scarano. An efficient and accurate approach to mte-mart for time-resolved tomographic piv. *Experiments in Fluids*, 56(3):1–16, 2015.
10. Vaibhav Vaish, Gaurav Garg, Eino-Ville Talvala, Emilio Antunez, Bennett Wilburn, Mark Horowitz, and Marc Levoy. Synthetic aperture focusing using a shear-warp factorization of the viewing transform. In *Computer Vision and Pattern Recognition-Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on*, pages 129–129. IEEE, 2005.
11. NA Worth and TB Nickels. Acceleration of tomo-piv by estimating the initial volume intensity distribution. *Experiments in Fluids*, 45(5):847–856, 2008.