

Tomographic 4D Reconstruction of Gas Flows in the Presence of Occluders

Kai Berger, Ivo Ihrke, Bradley Atcheson, Wolfgang Heidrich, Marcus Magnor

Computer Graphics Lab, TU Braunschweig

Imager Group, UBC Vancouver

Email: {berger,magnor}@cg.cs.tu-bs.de, {ihrke,atcheson,heidrich}@cs.ubc.ca

Abstract

We present a method that allows for reconstructing non-stationary, time-varying gas flows around moving objects. Our work extends the background oriented Schlieren tomography (3D-BOS) acquisition technique to capture gas flows also in the presence of occluding objects. An algorithm is presented that exploits the unique properties of BOS background patterns to robustly segment occluding objects. Numerical issues in the refractive index field reconstruction are addressed and successfully solved by the new method.

1 Introduction

In past decades gas flows were visualized with the so-called *Schlieren* technique. It is based on a collimated light source shining light through a gas flow, and a spatial filter to attenuate non-parallel light before projection onto an image plane. Due to refractive index variations in the gas flow, deflected light rays result in darker streaks in the image plane. A good overview of classical Schlieren measurement setups can be found in Settle's book [17].

Schlieren imaging is quite difficult to set up and calibrate. In recent years, another method based on image processing techniques, the so-called *Background-Oriented Schlieren* (BOS) method has been developed to simplify these measurements [7, 13, 14]. For BOS acquisition, it is sufficient to place a textured background pattern behind the gas flow and record with a camera. A still image of the scene without the refractive index field is captured and used as a reference image. Recording the gas flow, every newly captured frame is compared with the reference image by computing the optical flow for every pixel [9, 3]. The physical principle causing the apparent flow is depicted in

Fig. 1. It is obvious that both, classical Schlieren measurement setups, and the BOS method are only capable of qualitative measurements, because volumetric effects are visible only as projections to a plane. Recently, a method for tomographic reconstruction of gas flows using the BOS method was introduced [4]. With this so-called *3D-BOS* method it is possible to perform a quantitative measurement of the underlying refractive index field of the gas flow under observation.

While this method is capable of reconstructing gas flows accurately, it is currently not possible to measure interactions between gas flows and objects in the flow. Measuring such interactions yields interesting insight into the behavior of gas flows and the evolution of e.g. turbulent flow structures around object boundaries as occurring for instance in turbine blades. Our contribution is the extension of the 3D-BOS method to robustly handle occlusions in order to quantitatively measure the surrounding gas flow. After reviewing related work and giving a brief introduction into the 3D-BOS technique in Sect. 2, in Sect. 3, the necessary changes for each step of the existing method [4] are described. In Sect. 4, the alterations are integrated into the method and its performance in reconstructing interactions with occluders is examined.

2 Related Work and Background

While previous methods have only been able to reconstruct rotationally symmetric [2] or stationary gas flows [16], Atcheson et al. [4] describe a method, that is capable of tomographically reconstructing time-varying and non-stationary gas flows. Their method is based on the BOS technique, where a high-frequency background pattern is placed behind the gas flow under observation and recorded with a camera. A wavelet noise pattern [6] has been shown to increase BOS performance compared to




Figure 1: Virtual displacement caused by ray deflection. Under normal circumstances, points A_b and B_b on the background are imaged at A_i and B_i respectively. However, when refraction takes place, B_b appears at A_i and we see a virtual displacement of δ .

previously used patterns [3]. The image of the current frame is then compared with a previously captured reference image of the background pattern without the gas flow. Between these two images the per-pixel *optical flow* is computed [9, 11, 5]. The measured optical flow corresponds to the 2D projection of the deflections caused by the gas flow. It has been experimentally shown, that Horn-Schunck [9] optical flow is the best method to measure the 2D deflections [3].

Based on these 2D deflection vectors the method in [4] aims to reconstruct the 3D refractive index field that causes the measured 2D deflections. Therefore, the volume under observation has to be captured from multiple views. In Fig. 2 the processing pipeline of this method is shown. At first, the 2D vectors I_{oflow} are computed. Afterwards, the method generates silhouettes M_{gas} around the significant vectors in the image to mark the projected area of the gas plume in stage 2. Based on these silhouettes a visual hull [10] V_{gas} of the gas volume is computed by backprojecting the silhouettes M_{gas} into a voxelized grid.

In stage 3, an equation system is set up which is solved for each voxel inside the visual hull by standard numerical algebra techniques. The equation system is solved for the best-fitting set of 3D refractive index gradients which causes the measured 2D deflection vectors in every captured image inside the silhouette. After solving the equation system, each voxel inside the visual hull of the gas volume is assigned a 3D vector, which is the gradient of the unknown refractive index function at this voxel.

In stage 4, a Poisson-Integration [1] is performed to

compute the refractive index field from the 3D vector field. This refractive index field is the best-fitting solution for the measured 2D deflection vectors in the input images of the camera and is a tomographic reconstruction of the captured gas flow.




Figure 2: The main stages of the reconstruction method of Atcheson et al. [4]. At first the 2D deflection vectors I_{oflow} are computed for each camera using optical flow (S1, color-coded flow vectors overlaid on captured image). Then, the flow vectors are used to determine consistent silhouettes M_{gas} for computing a visual hull V_{gas} (S2). Within this hull an equation system is solved for 3D gradient vectors (S3), which are finally integrated to retrieve the refractive index volume (S4).

The acquisition method [4], however, is only capable of reconstructing gas flows from setups, where the camera view of the background pattern through the gas volume is not occluded by objects. In the presence of occluders, on the other hand, the optical flow computation in stage 1 results in erroneous per pixel deflection data in the projected area of the occluding object, so that the following stages cannot be performed correctly.

3 Obstruction 3D-BOS

As the method presented in [4] is not capable of handling objects inside the gas flow which partially occlude the background pattern, certain modifications to the method have to be implemented to enable measurement setups with occluding objects in the gas flow:

- In stage 1, the method has to ensure that only the deflection vectors outside the projected area of the occluder are used for the optical flow reconstruction. This is necessary, because the flow between occluded pixels and background pixels from the reference image would result in incorrect information. We discuss this in more detail in Sect. 3.1.
- In stage 2, having masked out the occluder, the silhouettes of the gas flow contain only pixels outside the projected area of the occluding object. However, the refractive index volume has to be fitted tightly around the volume of the occluder. We discuss this in Sect. 3.2.
- In stage 3, the 3D gradients can only be computed in the voxels which are projected to pixels containing 2D deflection information in the input images. Hence, gradient information is missing in the voxels inside the visual hull of the occluding volume. Since the final integration is a global operation over the voxelized grid, the missing information in the voxel volume of the occluder corrupts the refractive index estimates of the surrounding voxels. We discuss this in Sect. 3.3.

3.1 Masking Occluders

At the first stage, the occluding objects in the recorded camera image have to be segmented. This is necessitated by the fact that in the image regions occluded by an object no sensible optical flow between the camera image and the reference image can be computed. Furthermore, the pixel region comprised by the occluder can corrupt the optical flow values of neighboring pixels. In the case of BOS imaging we can state several properties useful for image segmentation:




Figure 3: *High-Pass-Filtering an image containing occluders. The high-frequency areas of the background pattern exhibit many edges while the occluder is out of focus and appears as one blurred area (left). High-pass filtering the image results in high values in background areas and low values in the occluded areas (marshmallow, middle). The final mask M_{occ} is shown in the right image.*

1. The background pattern always consists of high-frequency noise, e.g. wavelet noise,
2. The occluding object is out of focus and thus low-pass filtered,
3. The background pattern has a well-defined intensity distribution, e.g. a normal distribution, and
4. The occluder moves consistently over time in the input images. Sudden topological changes of the mask are unlikely.

The first and second property can be exploited by applying a high-pass filter to the current image, which results in high values in the area of the background pattern, Fig. 3. In the occluder area, low values dominate because the occluder is out of focus. We blur the resulting values and apply a threshold to obtain an occluder mask M_{occ} . We employ a search for connected components [15] to clean up the resulting mask.

The third property suggests that the background pattern has a consistent intensity histogram, Fig. 4. Comparing the histograms of small patches of the image with this intensity histogram and measuring the distance with a suitable distance function (e.g. Mahalanobis-distance [12]) results in a likelihood image I_{likely} . The higher a pixel value in I_{likely} , the more likely a patch is belonging to the background pattern. Thresholding the image results in a similar occluder mask M'_{occ} to filter out the occluder area.

The fourth property can be utilized after having computed the occluder masks for a full image se-




Figure 4: Low-frequency occluder areas exhibit a different intensity distribution than the high-frequency background. While the histogram of the background pattern is Gaussian-shaped (bottom histogram), the histograms of the occluder (middle histogram) and the white wall (top histogram) have a significantly different shape.

quence. Randomly occurring outliers can be removed by applying a temporal blur over a time-varying sequence of masks $M_{occ}^{t-n, \dots, t, t+n}$. When the temporally blurred masks are thresholded, only pixels remain, that are consistently marked as belonging to an occluder over a small period of time. However, boundary pixels of a moving occluder are likely to be filtered out with this approach. Thus, the resulting mask has to be dilated to allow for a reasonable speed of motion in the sequence. All three presented masking approaches, M_{occ} , M'_{occ} , M''_{occ} have advantages and drawbacks, which are discussed in Sect. 4.

3.2 Computing 3D Gradients

At this stage the occluder masks M_{occ} have been computed. Applying M_{occ} to the input image and computing the optical flow I_{oflow} results, after integration, in the non-occluded gas silhouette M_{gas} with

$$M_{gas} \cap M_{occ} = \emptyset.$$

The computed 2D deflection vectors in M_{gas} are then used for determining the index gradient field.

The index gradients can only be solved for in $V_{gas} = VH(M_{gas})$, the visual hull of the gas plume that surrounds the occluder ($VH(\cdot)$ indicates




Figure 5: The Occluder M_{occ} (b) is masked out and M_{gas} is computed from the remaining valid optical flow data (c). To fit V_{gas} (f) tightly around the occluder, we compute V_{union} (d) from M_{union} (a). The index gradients are then computed in the voxels in $V_{gas} = V_{union} \setminus V_{occ}$ (e).

the visual hull computation from the input masks of each camera). Only these voxels project to plausible optical flow data in the input images. In order to fit V_{gas} tightly around the occluder we compute the union mask M_{union} , with

$$M_{union} = M_{gas} \cup M_{occ}.$$

Next we compute the union visual hull V_{union} from M_{union} with

$$V_{union} = VH(M_{union})$$

The union visual hull V_{union} consists of both the gas visual hull V_{gas} and the occluder visual hull V_{occ} . The 3D vector computation is then performed in $V_{gas} = V_{union} \setminus V_{occ}$, Fig. 5.

Note that alternatively performing a direct mask subtraction in image plane and a visual hull generation of the remaining mask $VH(M_{gas})$ would result in an incomplete visual hull, Fig. 6, red regions. For the remaining voxel volume V_{gas} three $n \times m$ equation systems are set up with n being the number of measured 2D deflection vectors and m being the voxels of V_{gas} . The equation systems are solved for each component of the 3D gradients of these voxels.




Figure 7: When the 3D gradients are Poisson-integrated to obtain the final refractive index field, the missing gradients in the occluder volume influence the refractive index estimate in the nearby voxels. The refractive index field of a test data set [8] (a) is compared to the resulting index field after applying a rod-shaped occluder to the scene (b). Note that the estimated refractive indices of voxels near the occluder are higher (c) than in the original data set.

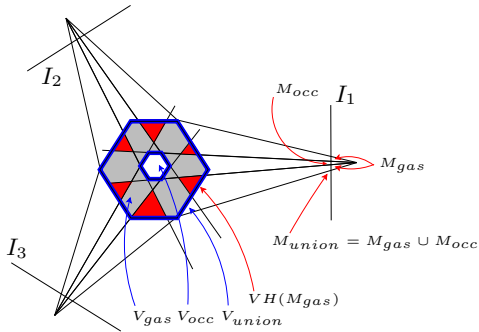


Figure 6: The visual hull V_{gas} (red and gray regions) is approximated best by computing V_{occ} and V_{union} separately and subtracting V_{occ} from V_{union} . Performing, instead, a mask subtraction in each image plane I_1, I_2, I_3 and a visual hull generation afterwards results in an incomplete visual hull, red regions.

3.3 Handling holes in the gradient volume

After determining the integration volume V_{gas} and computing the 3D gradients for the voxels in it, we finally have to take care of the Poisson-integration of the index gradients. Fig. 7 shows the problem that occurs, when a synthetic test volume (a) with a synthetic occluder positioned inside the volume (b) is Poisson-integrated. After straight-forward Poisson-integration the refractive index estimates of




Figure 8: In the volume of the occluder 3D gradients cannot not be computed (a), because no 2D deflection data is available. We diffuse the gradient information from the boundary into the holes (b).

voxels near the occluder differ significantly (c) from the input data set. Thus the influence of the boundary has to be minimized for the voxels around the occluder. In order to solve this problem we perform the Poisson-integration not in V_{gas} , but in the union visual hull V_{union} . We fill the holes in the occluder volume V_{occ} with sensible gradient data from the boundary. This process is a 3D equivalent to image inpainting. The gradient values of V_{gas} are diffused into the the voxels of V_{occ} . V_{occ} acts as the mask for the diffusion process. The result is visualized in Fig. 8. Before the diffusion step, a 2D slice of the 3D vector field shows a hole in the area occupied by the occluder (a). Afterwards, the data from the boundary are propagated into the hole (b). Note that the values of voxels in V_{gas} remain unchanged due to prior masking with V_{occ} .




Figure 9: The histograms of small patches of the background pattern (*top right, bottom right*) are shifted and obtain a narrow peak shape and are not as similar to the average background histogram (*top left*) as expected.

4 Experiments and Results

At first we tested the occluder segmentation in the first stage of the tomographic reconstruction method. We recorded scenes similar to the measurement setup in [4]: We placed a high-frequency background pattern behind the gas volume under observation and additionally moved small occluding objects into the gas plume. The occluding objects were recorded out of focus and thus appear blurred and low-pass filtered in the image sequence.

We found that the histogram-based approach is capable of masking out objects in the recorded sequences, which have a significantly different intensity distribution from the background. However, if filter patches are chosen too small, detection can become unreliable. Fig. 9 depicts the histograms of small patches and of a large area of the background. The histograms of the small patches are shifted and resemble narrow peaks rather than Gaussian shapes. This is due to variation in the average intensity between neighboring patches in the background pattern.

The high-pass filter method is able to determine occluder silhouettes of objects which are out of focus or exhibit a homogeneous intensity. We found that using this method the smallest detectable area size equals the average area between two detectable edges in the background pattern. A higher resolution in the background pattern thus enables the detection of smaller or thinner objects in the input




Figure 10: Masking out occluders. An input image captured from a typical measurement setup contains the high-frequency background, a gas flow and an occluder (*left*). The thresholded high-pass filtered image shows a good separation of occluder and background (*middle*). The resulting conservative mask overlaid onto the input image shows the robustness of the segmentation method (*right*).

image. A typical output of the masking process is shown in Fig. 10.

We also tested the gradient diffusion process with synthetic intensity images. We computed the gradients of the images and removed information by applying different masks, Fig. 11 (far left). The masked out holes were filled with gradient data from the boundary. We integrated the diffused gradients again to retrieve an intensity image, Fig. 11 (middle right), which we compared to the original image, Fig. 11 (middle left). In Fig. 12 the maximum per pixel errors for each applied mask are listed. We found that the amount of the absolute per pixel error is related to the position of the mask. When the mask is positioned in a high-frequency area of the image, the absolute error is higher than when it is positioned in a low frequency area.

Finally we tested the Poisson integration with a volume consisting both of voxels with and without a gradient value. We used a synthetic fuel injection data set with known refractive indices [8] and virtually placed a rod-shaped occluder diagonally inside the volume. After generating the occluder masks for every camera, we subtracted the visual hull of




Figure 11: Absolute per-pixel error for integrating diffused images. A binary mask (*far left column*) is applied to the gradients of an intensity image (*middle left column*). The masked-out gradients are filled out with diffusion. The resulting image is obtained by Poisson-integration (*middle right column*). The absolute per pixel error is depicted in the *far right column*.

the occluder from the visual hull of the fuel injection volume and computed the gradient values only in the remaining voxels.

After straight-forward integration, we found that voxels in the neighborhood of the occluder volume are significantly over-estimated, Fig. 13 (top right). When the occluder voxels are filled with sensible gradient data from the neighboring voxels by the diffusion process, the resulting estimates are much closer to the original values of the test data set, Fig.13 (bottom right).

5 Conclusion

We have presented a new gas flow reconstruction method that enables the measurement of interactions of the index field with occluders. We have identified the problems that occur when placing occluders into the measurement setup.

We derived modifications for each stage of the ex-

Mask	Maximal per pixel error	% Difference to error without mask
no	0.2088	0 %
1	0.2091	0.14 %
2	0.2118	1.43 %

Figure 12: The maximum per pixel error after the gradient diffusion. For every mask in Fig. 11 (*left row*) the maximum absolute error (*middle row*) and the percentage of difference to the maximum absolute error in an integrated image without a mask (*right row*) is listed.

isting 3D-BOS method [4] to allow for an index field computation in the presence of occluders. In the first stage we add a robust image segmentation method to mask out the occluder silhouette. In the second stage we define the integration volume to be the union of the visual hulls of the gas flow and the occluder. In the third step finally we diffuse the computed 3D gradient data into the voxels inside the visual hull of the occluder to provide for minimized artifact Poisson integration at the boundary to these voxels.

With the presented method one can now for the first time visualize interaction processes of heated gas flows with stationary and moving occluders in the BOS measurement setup.

6 Acknowledgements

This work has been funded by the German Science Foundation, DFG MA2555/5-1.

Ivo Ihrke was supported by a Feodor-Lynen Fellowship of the Alexander von Humboldt Foundation, Germany

References

- [1] A. Agrawal, R. Raskar, and R. Chellappa. What Is the Range of Surface Reconstructions from a Gradient Field? *Proceedings of ECCV*, 3951:578–591, 2006.
- [2] A. K. Agrawal, B. W. Albers, and D. W. Griffin. Abel Inversion of Deflectometric Measurements in Dynamic Flows. *Applied Optics*, 38(15):3394–3398, 1990.




Figure 13: The gradient diffusion process. **Top row:** Placing a rod-shaped occluder into a synthetic refractive index set [8] (**left**) and performing the altered method without gradient hole filling leads to black holes in the gradient field (**middle**). These holes are treated as surrounding unheated air in the final integration and corrupt the resulting refractive index estimates of nearby voxels (**right**). **Bottom row:** The same synthetic data set (**left**) with the same occluder is processed with gradient diffusion (**middle**). The estimated refractive indices of surrounding voxels are much closer to the original (**right**).

- [3] B. Atcheson, W. Heidrich, and I. Ihrke. An evaluation of optical flow algorithms for background oriented schlieren imaging. *Experiments in Fluids*, 46(3):467–476, 2009.
- [4] Bradley Atcheson, Ivo Ihrke, Wolfgang Heidrich, Art Tevs, Derek Bradley, Marcus Magnor, and Hans-Peter Seidel. Time-resolved 3d capture of non-stationary gas flows. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 27(5):article 132, 2008.
- [5] T. Brox, A. Bruhn, and N. Papenbergh. High Accuracy Optical Flow Estimation Based on a Theory for Warping. *Proceedings of ECCV*, pages 25–36, 2004.
- [6] R. L. Cook and T. DeRose. Wavelet Noise. *Proceedings of ACM SIGGRAPH*, 24(3):803–811, 2005.
- [7] S. Dalziel, G. Hughes, and B. Sutherland. Whole-field Density Measurements by ‘Synthetic Schlieren’. *Experiments in Fluids*, 28(4):322–335, 2000.
- [8] DFG. Fuel injection volumetric data set. DFG SFB 382.
- [9] B. K. P. Horn and B. G. Schunck. Determining Optical Flow. *Artificial Intelligence*, 17(1-3):185–203, 1981.
- [10] A. Laurentini. The Visual Hull Concept for Silhouette-based Image Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162, 1994.
- [11] B. D. Lucas and T. Kanade. An Iterative Image Registration Technique with an Application to Stereo Vision. In *International Joint Conference on Artificial Intelligence*, volume 3, pages 674–679, 1981.
- [12] P. C. Mahalanobis. On the generalized distance in statistics. In *Proc. Nat. Inst. Sci. India*, volume 2, pages 49–55, 1936.
- [13] G. Meier. Computerized Background-Oriented Schlieren. *Experiments in Fluids*, 33(1):181–187, 2002.
- [14] H. Richard and M. Raffel. Principle and Applications of the Background Oriented Schlieren (BOS) Method. *Measurement Science and Technology*, 12(9):1576–1585, 2001.
- [15] Christian Ronse and Pierre A. Devijver. *Connected components in binary images: the detection problem*. John Wiley & Sons, Inc., New York, NY, USA, 1984.
- [16] A. Schwarz. Multi-tomographic Flame Analysis With a Schlieren Apparatus. *Measurement Science and Technology*, 7(3):406–413, 1996.
- [17] G. S. Settles. *Schlieren and Shadowgraph Techniques: Visualizing Phenomena in Transparent Media*. Springer, 2001.