Volume Graphics (2006)
T. Möller, R. Machiraju, T. Ertl, M. Chen (Editors)

Volumetric Density Capture From a Single Image

Christian Fuchs[+] Tongbo Chen[+] Michael Goesele[*] Holger Theisel[+] Hans-Peter Seidel[+]

MPI Informatik[+] University of Washington[*]

Abstract

We propose a new approach to capture the volumetric density of scattering media instantaneously with a single image. The volume is probed with a set of laser lines and the scattered intensity is recorded by a conventional camera. We then determine the density along the laser lines taking the scattering properties of the media into account. A specialized interpolation technique reconstructs the full density field in the volume. We apply the technique to capture the volumetric density of participating media such as smoke.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation — Digitizing and Scanning; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling — Physically Based Modeling

1. Introduction and Previous Work

The acquisition of time-varying volumetric density fields has recently attracted the attention of computer graphics researchers. Hasinoff and Kutulakos [HK03] presented a multi-view method to volumetrically reconstruct flames using a photo-consistency approach. Ihrke and Magnor [IM04, IM05] used sparse view tomography to reconstruct flames and optically thin smoke from a small set of camera views. Hawkins et al. [HED05] captured time-varying smoke density by rapidly scanning a laser plane through a volume and imaging the light scattered by individual smoke particles from a lateral view with a high-speed camera (see Figure 1, left). Physical measurement systems such as Yip et al. [YLWL87] or laser induced fluorescence (LIF) [DD01] follow a similar approach and capture the whole volume *sequentially* from a single view. This allows [HED05] to sample locations in the moving light plane with high spatial resolution yielding high quality renderings of the captured model.

Our method is inspired by single view techniques but takes a fundamentally different sampling approach: The volume is illuminated with a grid of laser lines in order to sample the whole volume *simultaneously.* Essentially, the 2D laser plane is discretized and spread out to discrete locations in space. The volume is captured with a standard camera. As illustrated in Figure 1, this trades (potentially) continuous sampling in one spatial domain (Δy) against continuous sampling in the time domain (Δt). Careful placement of camera and light sources avoids occlusions when the laser lines are projected onto the image plane and the full but sparser sampled 3D information is captured with a single image. The sampling density can be increased by projecting multiple grids of different colored illumination into the volume.

This new sampling paradigm has several consequences:
Decoupling of spatial and temporal sampling: The system enables continuous sampling in the time domain allowing both integration over long time intervals for weak signals and extremely short acquisition times for fast-changing datasets.
Increased time resolution: Using the same camera hardware, frame rates can be increased by 1–2 orders of magnitude. Furthermore, the system contains no mechanically moving parts, i.e., the achievable time resolution is only limited by the imaging hardware and the illumination intensity.
Limited spatial resolution: Spatial resolution is fundamentally limited by the number of laser lines that can be resolved by the imaging system. Sampling is denser along the direction of the laser lines but spatial detail between sampling lines is lost yielding a potential for aliasing.
Density instead of photo-consistency: The actual density of scatterers is reconstructed (up to a scale factor) instead of a photo-consistent representation of the volume as in Hasinoff and Kutulakos [HK03].

In the remainder of the paper we first describe the basic concepts underlying our capture approach (Section 2). We then analyze its properties and study the resulting er-

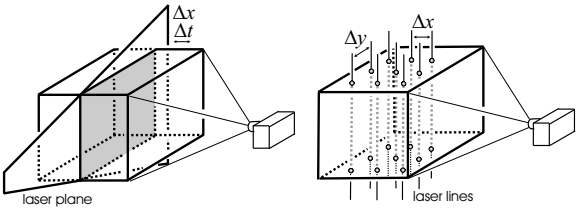

Figure 1: *Different acquisition approaches.* Left: *Sequential scanning of the volume with a plane of light as implemented by [HED05] with discrete sampling in time (Δt) and one spatial dimension (Δx).* Right: *New approach with simultaneous capture of volume density using line grids as illumination and discrete sampling in two spatial dimensions ($\Delta x, \Delta y$).*

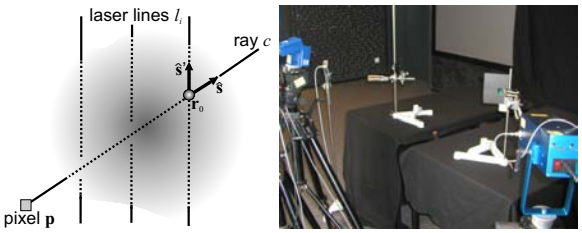

Figure 2: Left: *Situation for an individual ray c.* Right: *Image of the actual acquisition setup with camera on the left and two laser sources (left near camera and blue box on the right). When active, smoke is generated between the two tables.*

rors using a synthetic smoke dataset (Section 3). We describe our prototypical acquisition system and show results for real captured datasets (Section 4). The paper concludes with a discussion of future work.

2. System Description

Figure 1, right, gives an overview over the measurement setup. We assume that the measurement volume contains a spatially varying density of scattering particles that we would like to measure. Apart from their density, the scattering properties of the individual particles should be constant (or be well approximated by a constant). Depending on the size of the particles, scattering will either be dominated by Mie scattering (for larger particles such as typically found in smoke) or by Raleigh scattering [BH83]. In the remainder of this section, we first describe the principles of radiative transfer (Section 2.1) and develop our image formation model (Section 2.2). Section 2.3 shows how we can recover the density of scatterers along the laser lines illuminating the volume. We finally describe in Section 2.4 how we can recover the full density field from this information.

2.1. Radiative Transfer

We start our analysis with the equation of radiative transfer [Ish78] which describes the change in specific intensity $I(\mathbf{r},\hat{\mathbf{s}})$[†] for a small volume element ds at a position $\mathbf{r}$ in space and in a direction $\hat{\mathbf{s}}$:

$$\frac{dI(\mathbf{r},\hat{\mathbf{s}})}{ds} = -\rho\sigma_t I(\mathbf{r},\hat{\mathbf{s}}) + \frac{\rho\sigma_t}{4\pi}\int_{4\pi} p(\hat{\mathbf{s}},\hat{\mathbf{s}}')I(\mathbf{r},\hat{\mathbf{s}}')d\omega' + \varepsilon(\mathbf{r},\hat{\mathbf{s}}). \quad (1)$$

The first term models a decrease in specific intensity due to absorption and outscattering which is proportional to the number of particles per unit volume ρ and the extinction cross section σ_t. The second term describes the increase in specific intensity due to inscattering which depends additionaly on the phase function $p(\hat{\mathbf{s}},\hat{\mathbf{s}}')$. $\varepsilon(\mathbf{r},\hat{\mathbf{s}})$ is the emission from within ds.

In our measurement system (see Figure 1, right, and Figure 2 for the principle of the approach and the notation) we assume that light is scattered from homogeneous particles inside the volume. The extinction cross section σ_t and the phase function $p(\hat{\mathbf{s}},\hat{\mathbf{s}}')$ are constant throughout the volume but the number of particles per unit volume $\rho = \rho(\mathbf{r})$ varies. We furthermore assume that the laser lines l_i are generated outside the observation volume so that there is no emission inside, i.e., $\varepsilon(\mathbf{r},\hat{\mathbf{s}}) = 0$. The goal of the measurements is to recover $\rho(\mathbf{r})$ up to a scale factor, i.e., it is sufficient to recover $D(\mathbf{r}) = \sigma_t\rho(\mathbf{r})$. Equation 1 simplifies then to the first order differential equation

$$\frac{dI(\mathbf{r},\hat{\mathbf{s}})}{ds} = -D(\mathbf{r})I(\mathbf{r},\hat{\mathbf{s}}) + \frac{D(\mathbf{r})}{4\pi}\int_{4\pi} p(\hat{\mathbf{s}},\hat{\mathbf{s}}')I(\mathbf{r},\hat{\mathbf{s}}')d\omega'. \quad (2)$$

which we would like to solve under the boundary conditions given by our setup and assumptions.

2.2. Image Formation Model

We now take the specific situation of our measurement setup into account: The incident intensity I_i at a position $\mathbf{r}_0$ where the backprojected ray c of a pixel $\mathbf{p}$ with direction $-\hat{\mathbf{s}}$ intersects with a laser line l_i with direction $\hat{\mathbf{s}}_{l_i}$ can be computed as

$$I_i(\mathbf{r}_0,\hat{\mathbf{s}}) = \frac{D(\mathbf{r}_0)}{4\pi} I_{l_i}(\mathbf{r}_0,\hat{\mathbf{s}}_{l_i}) p(\hat{\mathbf{s}},\hat{\mathbf{s}}_{l_i}). \quad (3)$$

We assume in the following that $I_{l_i}(\mathbf{r}_0,\hat{\mathbf{s}}_{l_i}) = I_{l_i}$ is constant along each laser line l_i and can be calibrated in a separate step. Following Ishimaru [Ish78] we can now split the intensity along the ray c into two terms:

$$I(\mathbf{r},\hat{\mathbf{s}}) = I_{\mathrm{ri}}(\mathbf{r},\hat{\mathbf{s}}) + I_{\mathrm{d}}(\mathbf{r},\hat{\mathbf{s}}). \quad (4)$$

The reduced intensity $I_{\mathrm{ri}}(\mathbf{r},\hat{\mathbf{s}})$ models the decay of I_i due to absorption and outscattering along c according to the first term in Equation 2:

$$I_{\mathrm{ri}}(s) = I_i(\mathbf{r}_0,\hat{\mathbf{s}})e^{-\tau} = I_i(\mathbf{r}_0,\hat{\mathbf{s}})e^{-\int_0^s D(\mathbf{r})ds}. \quad (5)$$

[†] Note that specific intensity can be converted into radiance by integrating over the spectrum of the radiation.

s measures the distance from $\mathbf{r}_0$ along the ray c to the pixel $\mathbf{p}$. The remaining contributions caused by inscattering from the volume are accumulated in the diffuse intensity $I_d(\mathbf{r},\hat{\mathbf{s}})$. pixel $\mathbf{p}$ can therefore be described as:

$$I_\mathbf{p} = I_{ri}(s) + I_d(s). \tag{6}$$

2.3. Recovering $D(\mathbf{r})$ along Laser Lines l_i

The goal of this section is to recover the scaled density values $D(\mathbf{r})$ along the laser lines l_i. We rewrite Equation 6 using Equations 3 and 5:

$$I_\mathbf{p} = \frac{D(\mathbf{r}_0)}{4\pi} I_{l_i} p(\hat{\mathbf{s}},\hat{\mathbf{s}}_{l_i}) e^{-\int_0^s D(\mathbf{r})ds} + I_d(s). \tag{7}$$

This can be rewritten as

$$D(\mathbf{r}_0) = 4\pi \frac{I_\mathbf{p} - I_d(s)}{I_{l_i} p(\hat{\mathbf{s}},\hat{\mathbf{s}}_{l_i})} \cdot e^{\int_0^s D(\mathbf{r})ds}. \tag{8}$$

The phase function $p(\hat{\mathbf{s}},\hat{\mathbf{s}}')$ can either assumed to be isotropic, theoretically derived from the properties of the scattering media [BH83], or be measured experimentally [HED05]. Note, however, that our acquisition setup evaluates $p(\hat{\mathbf{s}},\hat{\mathbf{s}}')$ only over a small range of angles (all possible intersection angles of a laser line l_i and a ray c in Figure 2). Locally approximating $p(\hat{\mathbf{s}},\hat{\mathbf{s}}')$ with a constant yields therefore only a small error. I_{l_i} can be determined in a calibration step for each laser line l_i. The diffuse intensity $I_d(s)$ can be estimated from the set of camera pixels $N(\mathbf{p})$ in the spatial neighborhood of $\mathbf{p}$ whose backprojected ray does not intersect with any laser line l_i:

$$I_d(s) \approx \tilde{I}_d(s) = \frac{1}{||N(\mathbf{p})||} \sum_{\mathbf{p}' \in N(\mathbf{p})} I_{\mathbf{p}'}. \tag{9}$$

This assumes that $I_d(s)$ varies smoothly over the image plane and corresponds to the removal of multiple scattering in the calibration part of Hawkins et al. [HED05]. We furthermore need to set $\tilde{I}_d(s) = I_\mathbf{p}$ if $\tilde{I}_d(s) > I_\mathbf{p}$ to avoid physically implausible results, e.g., due to noise in the acquisition. Under the assumption of optically thin scattering material, we can furthermore set $\int_0^s D(\mathbf{r})ds = 0$ without introducing too large an error. Equation 8 becomes then

$$D(\mathbf{r}_0) \approx 4\pi \frac{I_\mathbf{p} - \tilde{I}_d(s)}{I_{l_i} p(\hat{\mathbf{s}},\hat{\mathbf{s}}_{l_i})} \tag{10}$$

which allows us to recover the scaled density values along all laser lines.

2.4. Reconstruction of the Entire Density Field $D(\mathbf{r})$

Given the density values D along the laser lines, we can employ several different interpolation techniques to interpolate $D(\mathbf{r})$. Any interpolation technique such as the push-pull algorithm [GGSC96, DCOY03] will, however, preserve the values of D along the laser lines and thus also preserve the non-uniform sampling inherent in our data structure. High frequency details will only be available along the laser lines and yield noticeable artifacts.

We therefore opted for a more general reconstruction approach that approximates the sample values and yields a smoother reconstruction. We formulate our approximation problem as follows: Given are n discrete samples of measured density values $D(\mathbf{p}_i)$ at locations $\mathbf{p}_i$ $(i = 0,\ldots,n-1)$ on the laser lines. We then approximate the field at a position $\mathbf{r}$ as

$$\widetilde{D}(\mathbf{r}) = \frac{\sum_{i=0}^{n-1} D(\mathbf{p}_i) \cdot w(\|\mathbf{r}-\mathbf{p}_i\|)}{\sum_{i=0}^{n-1} w(\|\mathbf{r}-\mathbf{p}_i\|)}. \tag{11}$$

The weighting function w is defined as

$$w(x) = \begin{cases} 0.5 \cdot \cos(\frac{x \cdot \pi}{R}) + 0.5 & \text{for } x < R \\ 0 & \text{else} \end{cases}. \tag{12}$$

The parameter R which determines the width of the reconstruction kernel needs to be manually selected for a given sampling configuration. Note that R has to be chosen such that in the R-neighborhood of every point $\mathbf{r}$ of the domain there is at least one sample point $\mathbf{p}_i$.

3. Simulation

To perform an analysis of a dataset with ground truth we used a 100 frame simulation of smoke emitted from the border of a volume using the technique of Treuille et al. [TLP06]. The data was stored as a 64^3 voxel density field. Assuming that each dataset is defined over the domain $[0,1]^3$ and assuming a trilinear interpolation between the grid points, we have continuous scalar fields f_g for each time frame which act as ground truth.

In order to reconstruct one of them, we consider a bundle of 100 rays starting from $(\frac{1}{2},\frac{1}{2},-3)$ and passing through $(\frac{i}{9},\frac{j}{9},1)$ for $i,j = 0,\ldots,9$. Then we sampled f_g in its domain along the rays and applied our approximation technique with parameter $R = 8$ to get a reconstructed field $f_1^{R=8}$. In a second test, we added a second bundle of 100 rays starting at $(-5,\frac{1}{2},\frac{1}{2})$ and passing through $(1,\frac{i}{9},\frac{j}{9})$ for $i,j = 0,\ldots,9$. The field reconstructed from these 200 rays with $R = 8$ is $f_2^{R=8}$. Figure 3 shows $f_g, f_1^{R=8}$, and $f_2^{R=8}$ for Frame 60 of the dataset. While both reconstructions faithfully represent the overall structure of the field, it is clear that many high-frequency details are lost. We therefore computed a smoothed version of the ground truth field by convolving it with a normalized version of the reconstruction kernel w (Equation 12). Figure 3(d) depicts the smoothed ground truth field $f_g^{R=8}$ for a kernel radius $R = 8$ which is well approximated by both, $f_1^{R=8}$ and $f_2^{R=8}$.

3.1. Error Analysis

We define the RMS error between two scalar fields f and f' stored as a voxel densities sampled at a set of identical locations $\mathbf{V}$ as

$$\mathrm{RMS}(f,f') = \sqrt{\frac{\sum_{v_i \in \mathbf{V}} (f(v_i) - f'(v_i))^2}{\|\mathbf{V}\|}}. \tag{13}$$

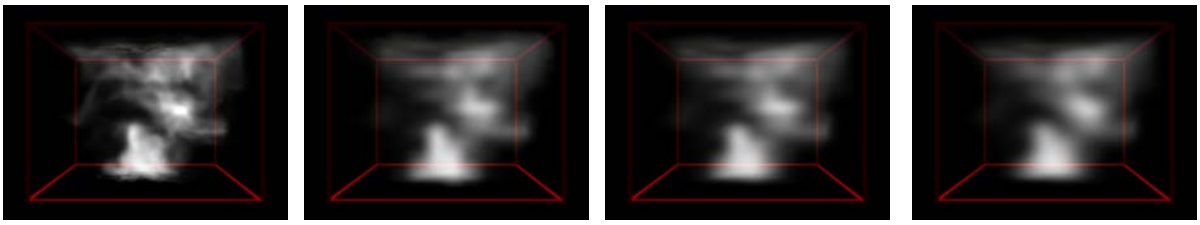

(a) ground truth f_g (b) reconstruction $f_1^{R=8}$ (c) reconstruction $f_2^{R=8}$ (d) smoothed ground truth $f_g^{R=8}$

Figure 3: *Result for the simulated smoke dataset.* (a): *Rendering of Frame 60 of the synthetic smoke dataset.* (b) and (c): *Reconstruction using one bundle and two bundles of 10×10 rays, respectively, and* $R = 8$. (d): *Rendering of the smoothed original volume with* $R = 8$.

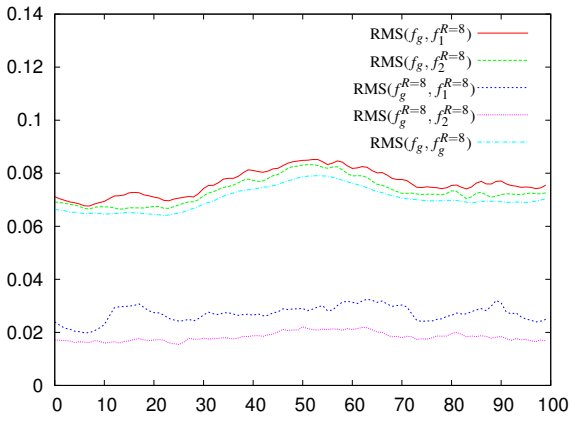

Figure 4: *Various RMS errors for the 100 frames contained in the simulated dataset. Horizontal axis: frame number. Vertical axis: RMS error. The density values in the original dataset vary between 0 and 1.6.*

In practice, $\mathbf{V}$ corresponds to the set of 64^3 voxels defining our field. Figure 4 depicts the RMS error between various versions of the field for all 100 frames in the dataset. The density values in the original dataset vary between 0 and 1.6.

The RMS errors between the ground truth f_g and the reconstructions from one and two ray bundles ($f_1^{R=8}$ and $f_2^{R=8}$) are almost identical. The RMS error decreases drastically when it is computed against the smoothed version $f_g^{R=8}$. Furthermore, $f_2^{R=8}$ fairs now much better than $f_1^{R=8}$. This suggests that most of the error in the reconstructions is due to the sparse sampling and smooth approximation that suppresses high frequency detail. To verify this, we computed the RMS error RMS$(f_g, f_g^{R=8})$ between the ground truth field and the smoothed version of the ground truth field. Figure 4 shows that this error is only slightly lower than RMS$(f_g, f_2^{R=8})$, i.e., most of the reconstruction error seems to be due to the loss of high frequency detail.

4. Acquisition System

We built a prototype of an acquisition system to test our ideas in practice. The setup contains two collimated laser sources – a 130 mW red laser and a 30 mW blue laser. Each collimated laser beam is transformed into a bundle of perspective rays using a burst grating. A black mask limits the number of rays in each bundle to 5×5 rays for each laser. The generated ray bundles are roughly perpendicular to each other; their intersection defines a measurement volume of about 50×24×50 cm^3 (see Figure 2, right, for an image of the measurement setup). We use a smoke machine to create a spatially varying density of scattering media.

A high quality color CCD camera is used to capture images of the measurement volume. Its placement ensures that no two rays of the same color project to the same location on the image plane. We are thus able to capture the two bundles independently using the camera's red and blue channel. In addition, we illuminate the measurement volume diffusely with green LEDs and use the green channel of the camera to capture simultaneously a ground truth image of the acquired dataset. Figure 6, left, shows an example input image of the system. Although all three light sources have a narrow spectral response, we observe crosstalk between some of the color channels which we remove using standard image processing techniques.

4.1. Calibration and Capture

For geometric calibration of the camera and the lasers we capture several high-dynamic range images [DM97] of each bundle of laser rays illuminating a yellow and white checkerboard pattern. This allows us to reliably detect the centers of the laser spots. After computing the 3D location of these spots using a calibration toolbox [Bou06] we can estimate the position of the laser lines in space relative to the camera. Figure 5 (left) shows a visualization of the spatial sampling in which the two bundles of laser rays are clearly visible. Using the calibration data, we can now extract samples

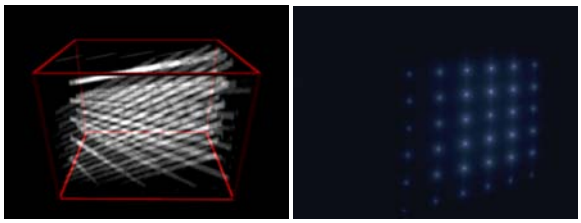

Figure 5: *Geometric and photometric calibration of the setup.* Left: *Visualization of the spatial sampling. The two bundles of laser rays are clearly visible.* Right: *Calibration image for the blue laser to recover the relative intensities I_{l_i} for the individual rays (tone-mapped).*

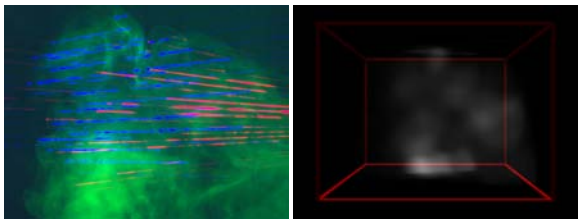

Figure 6: *Acquired image of smoke and reconstructed density field.* Left: *Input image with two independent grids of laser lines (red, blue) sampling the volume and diffuse green illumination (image intensity rescaled for display).* Right: *Density field reconstructed from this single image.*

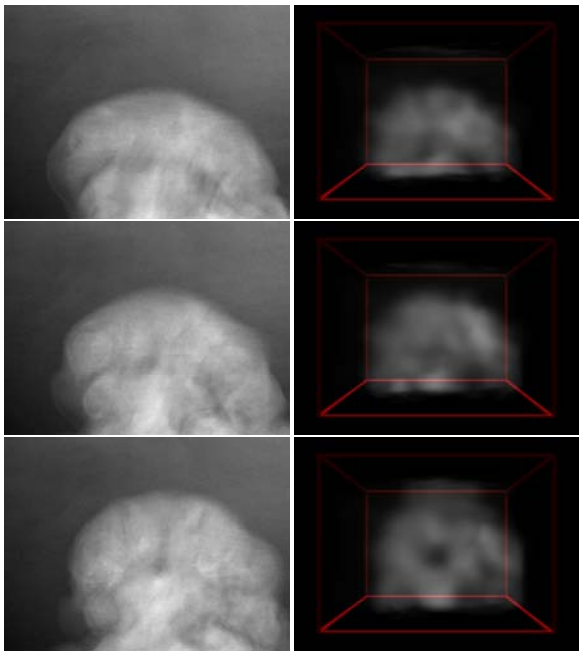

Figure 7: *Results for smoke from a smoke machine. The three images were taken in sequence from top to bottom. The images on the left show ground truth photographs of the scene extracted from the green channel. The images on the right show renderings of our reconstruction.*

$I_{\mathbf{p}}$ from the camera images by marching along the projections of the rays. We improve stability and ensure that the full width of the projected laser line is captured by integrating over a small neighborhood perpendicular to the projected ray direction.

The burst gratings create bundles of rays where each laser line can have a different intensity. We therefore need to recover the relative intensity I_{l_i} of all laser lines in each bundle in the calibration phase. To this end, we capture a high-dynamic range image per laser of a sheet of subsurface scattering plastics illuminated by the corresponding ray bundle (see Figure 5 right). We integrate the laser intensity over a constant area around each laser spot to get I_{l_i}. The two lasers might still behave differently, e.g., due to different scattering properties for different wavelengths. We therefore capture during calibration images of smoke and recover separate density fields D_{red} and D_{blue} for the red resp. the blue laser bundle. We then determine a scale factor k that minimizes the rms error between D_{red} and $k \cdot D_{\text{blue}}$ and scale the input data for the blue channel accordingly before reconstructing the density field.

4.2. Acquisition Results

We captured several datasets with the acquisition system described in Section 4 and depicted in Figure 2 (right). All images were taken with 0.25 s exposure time. Captured results are shown in Figures 6 and 7.The results are rendered using a raytracing-based direct volume rendering approach.

Figure 6 shows an input image capturing smoke from a smoke machine. The intensity variations along the laser lines are clearly visible. In addition, the blue laser lines are tighter focussed than the red lines which is corrected by the intensity scaling described in Section 4. Note in the result image on the right how much detail could be reconstructed from this single input view.

Figure 7 shows comparisons between ground truth photographs and renderings for three images of dataset consisting of 50 images captured at approximately 3 fps. Note that this speed is mainly constrained by the frame rate of the camera and not by the measurement principle. The ground truth photographs were extracted from the camera's green channel as described in Section 4. The reconstructed dataset captured the overall shape of the smoke as well as prominent features. Its resolution is, however, limited due to the sparse sampling using only 5×5 grids of laser lines.

5. Discussion and Conclusion

We presented a new sampling approach to sample time-varying volumetric density fields using grids of laser lines as illumination. This allows us to continuously sample in the

time domain at the cost of sparser sampling on the spatial domain.

Like other measurement systems such as Hawkins et al. [HED05] or Narasimhan et al. [NNSK05], we are assuming that the scattering behavior inside the measurement volume is dominated by single scattering. This limits both, the size of the measurement volume and the density of the scatterer inside the volume, due to two effects: First, the intensity of the laser rays I_{l_i} and the scattered intensity $I_{\mathbf{p}}$ decrease inside the measurement volume due to outscattering and absorption yielding a systematic bias in the reconstructed field $D(\mathbf{r})$. Modeling and inverting this effect for spatially varying densities is difficult even if all scattering parameters are known. Second, the diffuse intensity I_{d} increases hereby severely limiting the signal-to-noise ratio in the measurements.

5.1. Future Work

There are several directions for further research: A different camera system would allow us to operate the whole setup at a higher frame rate removing motion blur from captured images. The number of laser lines in the setup is currently limited by the properties of the gratings. Using individual laser sources (e.g., a set of laser pointers) to generate the lines independently would allow a much denser sampling. The number of lines is then only limited by their projected width on the captured image.

The current data processing approach makes no assumption about the structure of the density field $D(\mathbf{r})$ and yields therefore a smooth reconstruction of the smoke volume (especially in sparsely sampled dimensions). Reconstruction algorithms that make use of prior knowledge of the structure of the data (e.g., [ONOI04]) can improve the visual quality of the reconstructed density field. Alternatively, the structure could be inferred by analyzing the frequency content in denser sampled dimensions in the spirit of Dischler et al. [DGF98].

Acknowledgements

We would like to thank Adrian Treuille for providing the simulated smoke dataset. This work was supported in part by a Feodor Lynen Fellowship granted by the Alexander von Humboldt Foundation, the University of Washington Animation Research Labs, Adobe, and Microsoft.

References

[BH83] BOHREN C. F., HUFFMAN D. R.: *Absorption and Scattering of Light by Small Particles*. Wiley & Sons, 1983. 2, 3

[Bou06] BOUGUET J.-Y.: Camera Calibration Toolbox for MATLAB, 2006. Available online at http://www.vision.caltech.edu/bouguetj/calib_doc/. 4

[DCOY03] DRORI I., COHEN-OR D., YESHURUN H.: Fragment-Based Image Completion. *ACM Trans. Graph. 22*, 3 (2003), 303–312. 3

[DD01] DEUTSCH S., DRACOS T.: Time Resolved 3D Passive Scalar Concentration-Field Imaging by Laser Induced Fluorescence (LIF) in Moving Liquids. *Meas. Sci. Technol. 12* (2001), 188–200. 1

[DGF98] DISCHLER J.-M., GHAZANFARPOUR D., FREYDIER R.: Anisotropic Solid Texture Synthesis Using Orthogonal 2D Views. *Comput. Graph. Forum 17*, 3 (1998), 87–96. 6

[DM97] DEBEVEC P., MALIK J.: Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of SIGGRAPH 97* (1997). 4

[GGSC96] GORTLER S. J., GRZESZCZUK R., SZELISKI R., COHEN M. F.: The lumigraph. In *Proceedings of SIGGRAPH 96* (1996). 3

[HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: Acquisition of Time-Varying Participating Media. *ACM Trans. Graph. 24*, 3 (2005), 812–815. 1, 2, 3, 6

[HK03] HASINOFF S. W., KUTULAKOS K. N.: Photo-Consistent 3D Fire by Flame-Sheet Decomposition. In *Proc. ICCV 2003* (2003), pp. 1184–1191. 1

[IM04] IHRKE I., MAGNOR M.: Image-Based Tomographic Reconstruction of Flames. In *ACM Siggraph / Eurographics Symposium on Computer Animation* (2004), pp. 367–375. 1

[IM05] IHRKE I., MAGNOR M.: Adaptive Grid Optical Tomography. In *Vision, Video, and Graphics* (2005), pp. 141–148. 1

[Ish78] ISHIMARU A.: *Wave Propagation and Scattering in Random Media*. Academic Press, 1978. 2

[NNSK05] NARASIMHAN S. G., NAYAR S. K., SUN B., KOPPAL S. J.: Structured Light in Scattering Media. In *Proc. ICCV 2005* (2005), pp. 420–427. 6

[ONOI04] OWADA S., NIELSEN F., OKABE M., IGARASHI T.: Volumetric Illustration: Designing 3D Models with Internal Textures. *ACM Trans. Graph. 23*, 3 (2004), 322–328. 6

[TLP06] TREUILLE A., LEWIS A., POPOVIĆ Z.: Model Reduction for Real-Time Fluids. *ACM Trans. Graph. 25*, 3 (2006). 3

[YLWL87] YIP B., LAM J. K., WINTER M., LONG M. B.: Time-Resolved Three-Dimensional Concentration Measurements in a Gas Jet. *Science 235* (1987), 1209–1211. 1