Int J Comput Vis (2014) 110:128–140
DOI 10.1007/s11263-013-0686-0

Ultra-fast Lensless Computational Imaging through 5D Frequency Analysis of Time-resolved Light Transport

Di Wu · Gordon Wetzstein · Christopher Barsi · Thomas Willwacher · Qionghai Dai · Ramesh Raskar

Received: 26 February 2013 / Accepted: 30 November 2013 / Published online: 28 December 2013

Abstract Light transport has been analyzed extensively, in both the primal domain and the frequency domain. Frequency analyses often provide intuition regarding effects introduced by light propagation and interaction with optical elements; such analyses encourage optimal designs of computational cameras that efficiently capture tailored visual information. However, previous analyses have relied on instantaneous propagation of light, so that the measurement of the *time dynamics* of light–scene interaction, and any resulting information transfer, is precluded. In this paper, we relax the common assumption that the speed of light is infinite. We analyze free space light propagation in the frequency domain considering spatial, temporal, and angular light variation. Using this analysis, we derive analytic expressions for information transfer between these dimensions and show how this transfer can be exploited for designing a new lensless imaging system. With our frequency analysis, we also derive performance bounds for the proposed computational camera architecture and provide a mathematical framework that will also be useful for future ultra-fast computational imaging systems.

Keywords Computational photography · Light transport · Frequency analysis · Lensless imaging

1 Introduction

With the invention of the electronic stroboscope, Harold "Doc" Edgerton captured the first photograph of a bullet piercing an apple about half a century ago. Today, ultra-fast sensors (e.g., Hamamatsu 2012) can produce visual information at effectively one trillion frames per second, roughly one million times faster than Edgerton's electronic strobes. This technology allows for the capture of light as it propagates through space and facilitates a wide variety of applications in computer vision. Indeed, the common assumption that the speed of light is infinite becomes deprecated. Imaging systems with such high time resolution provide a new degree of freedom available to computer vision applications, namely, time-resolved ultra-fast imaging.

Whereas ultra-fast imaging techniques, such as time-of-flight capture or light detection and ranging (LiDAR), have been successful in measuring the travel time of short light pulses for applications in depth estimation, measuring the full time profile of sensor pixels over an extended period of time has started only recently to be explored in the computer vision community. The key insight of this paper is to treat light propagation among emitters, the scene, and receivers in five dimensions, with space, angle, and time all coupled. This leads to joint re-examination of fundamental inverse problems and solutions in computational imaging. Our paper explores ultrafast computational imaging in four aspects: theory, capture, analysis, and applications.

Electronic supplementary material The online version of this article (doi:10.1007/s11263-013-0686-0) contains supplementary material, which is available to authorized users.

D. Wu · G. Wetzstein (✉) · C. Barsi · R. Raskar
MIT Media Lab, Cambridge, MA, USA
e-mail: gordonw@media.mit.edu

D. Wu · Q. Dai
Department of Automation, Tsinghua University, Beijing, China

T. Willwacher
Department of Mathematics, Harvard University, Cambridge, MA, USA

D. Wu
Graduate School at Shenzhen, Tsinghua University, Beijing, China

	space-angle $x-\nu$	space-time $x-t$	space-angle-time $x-\nu-t$
forward propagation (primal domain)	shear in x	convolution with depth-dependent hyperbola	shear in x hyperbolic curvature in t
forward propagation (frequency domain)	shear in f_ν	multiplication with Fourier transform of hyperbola	shear in f_ν convolution in f_ν
inverse propagation (primal domain)	shear back in x	single depth: **deconvolution** multiple depths: **tomography**	**tomographic reconstruction**
inverse propagation (frequency domain)	shear back in f_ν	single depth: **division** multiple: **filtered backprojection**	**filtered backprojection**

Fig. 1 Overview of operators for forward and inverse light field transport. While conventional analysis (e.g., Durand et al. 2005) consider only time-independent transport (left column), we analyze time-resolved forward and inverse light transport in this paper (*center* and *right columns*) and show that the operations required to inverse-propagate a light field in space, time, and angle correspond to computed tomographic reconstructions. Based on these insights, a new lensless computational imaging system is proposed

In particular, we analyze light propagation in free space in both the primal and frequency domains, considering space, angle, and time for a finite speed of light. We discuss cross-dimensional information transfer and derive upper bounds for ultra-fast streak sensors, which record only subsets of the high-dimensional space. Facilitated by this analysis, we propose a new computational camera architecture that requires no optical elements besides a time-resolved sensor and a pulsed light source. Relevant scene information is recovered via computed tomographic reconstruction in post-processing. Summarized in Fig. 1, our contributions are as follows:

- We analyze free space light propagation in the frequency domain considering spatial, temporal, and angular light variation. We show that the propagation in frequency space can be modeled analytically as a combination of a shear in the light field and a convolution along the angular frequencies.
- We show that propagation in free space has unique information-preserving properties by coupling information among different dimensions. We derive upper bounds for how much of the information contained in one dimension is transferred to those that are measured.
- Based on this analysis, we introduce a novel, lensless camera. This approach exploits ultra-fast imaging combined with iterative reconstruction, while removing the need for optical elements, such as lenses or masks.
- Using the proposed frequency analysis, we derive upper bounds on the depth of field of the proposed camera and verify them with simulations.
- With synthetic scenes and an experimental prototype camera, we demonstrate and evaluate the proposed computational imaging approach.

1.1 Overview

This work is organized as follows. In Sect. 2, we review relevant literature. Section 3 introduces a detailed analysis of transient light propagation in free space in both primal and frequency domains. This analysis is used in Sect. 4 to derive upper bounds for the amount of spatial image information that is preserved after propagation. We introduce a new approach to lensless computational imaging in Sect. 5. The proposed technique captures spatial and temporal image information using a bare sensor and uses computed tomography for reconstruction. We evaluate the performance of the CT-based reconstruction algorithm with resolution charts in Sect. 6 and with simulated and physical experiments in Sect. 7. Section 8 discusses features of the technique and suggests future work.

2 Related Work

Frequency analyses of light transport (Durand et al. 2005) have been of critical importance in computer vision and graphics for evaluating the performance of imaging systems and identifying novel computational camera designs (Edward et al. 1995; Levin 2009; Veeraraghavan et al. 2007) or implementing faster (Ng 2005) and better (Levin and Durand 2010) computational processing schemes. Time-resolved forward and inverse light transport, however, has begun to be studied only recently (Raskar and Davis 2008).

Lensless imagers have been studied and include designs that employ attenuating layers with control of transmittance in space and time (Zomet and Nayar 2006), template-based approaches using micro sensors (Koppal et al. 2011), Fresnel zone plates or photon sieves for imaging using diffractive optics (Andersen 2005), angle-sensitive pixels that require no off-chip optics (Gill et al. 2011), and dynamic LED illumination combined with computed tomographic image reconstruction (Isikman et al. 2011). We introduce a new computational camera design that requires only a single image captured by a time-resolved ultra-fast sensor and does not require additional refractive, attenuating, or diffractive optical elements.

Ultra-fast sensors with the ability to measure temporal impulse responses provide a rich source of imaging

information in seismology, sonology, and optics. Reflection seismology and medical ultrasound use such measurements for performing non-invasive volumetric imaging, whereas sonar, radar, and LiDAR technologies determine an object's range, among other properties, from the echos of sound, radio waves, and light pulses reflecting from a target object (Saleh 2011). Unlike LiDAR, which measures the time of arrival of direct reflections of light, the proposed computational camera employs an ultra-fast sensor that exploits the full time-resolved profile, measuring photons as they arrive at the sensor over an extended period of time, for imaging both albedo and depth.

Imaging at speeds fast enough to observe radiometric changes over time has recently unveiled new applications for scene analysis that cannot be performed with standard cameras. These ultra-fast imaging methods have the ability to reconstruct geometry (though not its albedo) of occluded objects "around a corner" (Velten et al. 2012), acquire in-the-wild BRDFs without using encircling equipment for scenes with known geometry (Naik et al. 2011), recover an occluded target's motion using multipath analysis (Pandharkar et al. 2011), and decompose global light transport (Wu et al. 2012). In this paper, we present a 5D frequency analysis of time-resolved light transport in free space that produces an intuitive explanation of forward light transport within all of the above works. We hope to provide the fundamental mathematical tools for developing and understanding emerging ultra-fast computational camera designs and new applications enabled by this technology.

3 Frequency Analysis of Transient Light Transport

In this section, we analyze time-resolved light transport in free space. While an analysis of steady-state light transport is intuitive in revealing where information is localized in the frequency domain (Durand et al. 2005; Levin 2009), considering the time-resolved case, to the best knowledge of the authors, has not been discussed in the literature. We analyze first the flatland, three-dimensional case and then generalize the analysis to five dimensions.

3.1 3D Frequency Analysis

Starting from first principles, space–angle light transport is commonly denoted as (Isaksen et al. 2000; Chai et al. 2000)

$$l^{(S)}(x, v) = l_0^{(S)}(x - vd, v), \tag{1}$$

where $l^{(S)}(x, v)$ is the steady-state light field after transport by a distance d, x is the spatial variable, and v is the relative coordinate on a plane at unit distance (Chai et al. 2000) (Fig. 2, left). Including the speed of light, we recognize that different ray paths arrive at different times. In order to accurately model such systems, the time-resolved forward light transport must consider the temporal dimension t, incorporating the finite speed of light c as

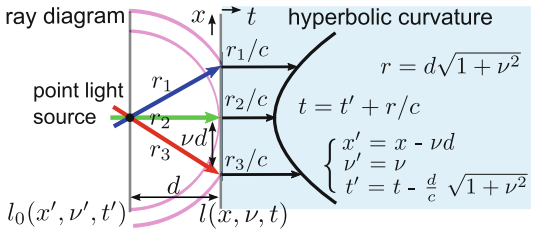

Fig. 2 Ray diagram illustrating the two-plane parameterization of light fields. Rays arrive at a plane a distance d from the source at times that depend on their propagation angle. The left is object space, and the right (*shaded blue*) is the corresponding space–time image (Color figure online)

$$l(x, v, t) = l_0(x - vd, v, t - \frac{d}{c}\sqrt{1 + v^2}). \tag{2}$$

Physically, rays propagating at large angles arrive later than low-angled ones. The delay can be calculated geometrically (Fig. 2). Note that for a pulsed point source, the light field $l_0(x, v, t) = \delta(x)\delta(t)$ becomes $l(x, v, t) = \delta(x - vd)\delta(t - \frac{d}{c}\sqrt{1 + v^2}) \sim \delta(\frac{x}{d} - v)\delta(t - \frac{d}{c}\sqrt{1 + (\frac{x}{d})^2})$ after a propagation distance d. From the second δ-factor, we see that time-resolved propagation produces a hyperbolic space–time curve (Fig. 2, right). The curvature and vertex of the hyperbola vary with the origin of the pulse. For example, a 3 mm change in path length results in a time shift of 10 ps. A summary of space–time propagation is shown in Fig. 3.

In the Fourier domain, propagation in free space for the steady-state case is a shear along the angular frequencies (Durand et al. 2005):

$$\tilde{l}^{(S)}(f_x, f_v) = \tilde{l}_0^{(S)}(f_x, f_v + f_x d), \tag{3}$$

where $\tilde{}$ denotes the Fourier transform of a quantity and f_x and f_v are the frequency variables of space and angle, respectively. For this case, the space–angle shear in the primal domain (Eq. 1) intuitively translates to a shear along the opposite dimension in the frequency domain. As illustrated in Fig. 4, the time-resolved light field spectrum remains sheared in the f_x-f_v dimension, but is now blurred in the f_t dimension, mainly along two branches that form an $\times$-shape. To quantify this result, we calculate the Fourier transform of the time-resolved light field after propagation in free space (both Eqs. 4 and 5 are derived in the Supplement):

$$\begin{aligned}\tilde{l}(f_x, f_v, f_t) &= \iiint e^{-2i\pi f_x x - 2i\pi f_v v - 2i\pi f_t t} \\ &\quad \times l_0\left(x - vd, v, t - \frac{d}{c}\sqrt{1 + v^2}\right) dx\, dv\, dt \\ &= \tilde{l}_0(f_x, f_v + f_x d, f_t) \underset{f_v}{*} F_{\frac{f_t d}{c}}(f_v),\end{aligned} \tag{4}$$

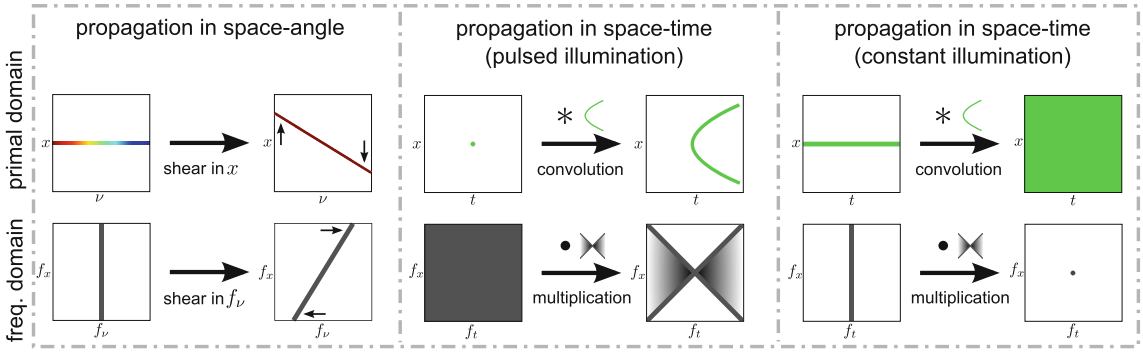

Fig. 3 Without time resolution, free space propagation shears the light field in both the primal and frequency domains (*left*). Due to relative time delays of different rays, a pulsed point source produces a space–time hyperbola, which acts as a convolution kernel on time-resolved light fields (*center*). In frequency space, this hyperbolic convolution is transformed to multiplication with its spectrum. Importantly, without a time-resolved source, only the *DC* component is propagated (*right*). Please note that the space–angle shear is maintained, though not shown, in the time-resolved case

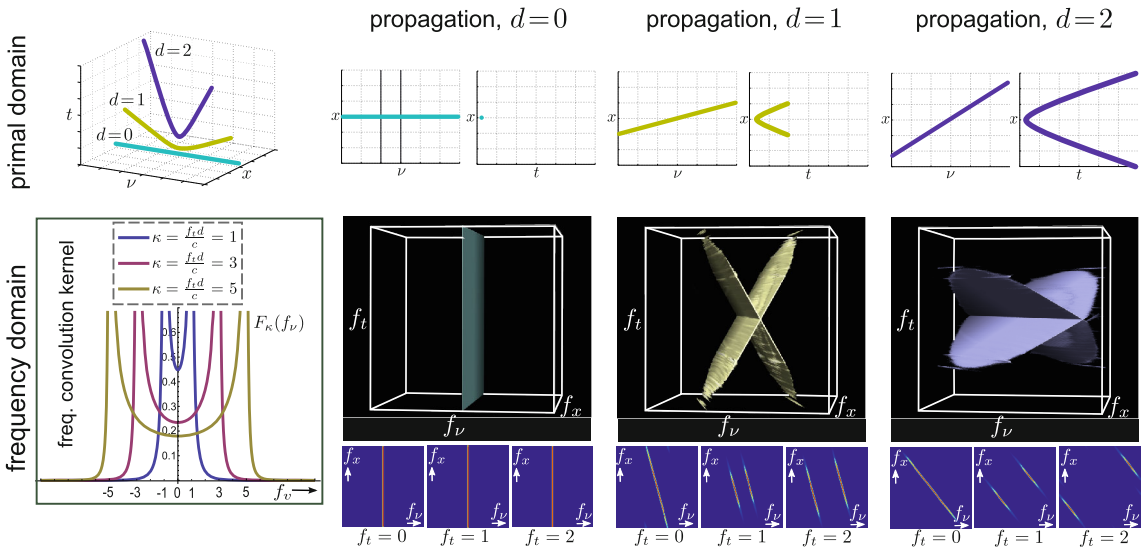

Fig. 4 Visualization of light field propagation. *Top row*: primal domain propagation of light pulse ($l_0(x, \nu\, t) = \delta(x)\delta(t)$). As the propagation distance d increases, the x-t hyperbola increases in curvature, and the x-ν shear becomes steeper. *Bottom row*: corresponding frequency domain representations. The steady state ($f_t = 0$) cross section is the standard Fourier shear, which splits into an $\times$-shape off-plane. (*Left*): cross section of spectrum at various off-axis positions. The time range of this figure is 0–50 ps

where

$$F_\kappa(f_\nu) := \int d\nu e^{-2i\pi f_\nu \nu - 2i\pi\kappa\sqrt{1+\nu^2}}$$
$$= \begin{cases} -\frac{\pi\kappa}{\sqrt{\kappa^2 - f_\nu^2}} H_1^{(2)}(2\pi\sqrt{\kappa^2 - f_\nu^2}) & \text{for} \quad \kappa^2 > f_\nu^2 \\ \frac{2i\kappa}{\sqrt{f_\nu^2 - \kappa^2}} K_1(2\pi\sqrt{f_\nu^2 - \kappa^2}) & \text{for} \quad \kappa^2 < f_\nu^2 \end{cases} . \quad (5)$$

Here, $H_1^{(2)}$ is a Hankel function of the second kind, K_1 is a modified Bessel function of second kind, and $\kappa = f_t d/c$. Cross sections of this function are plotted in Fig. 4. Equation 4 mathematically expresses the effects of propagation on the light field spectrum through (1) shearing the space–angle dimension, then (2) convolving the result with $F_{\frac{f_t d}{c}}(f_\nu)$ along the f_ν dimension. Defining the bandwidth of $F_\kappa(f_\nu)$ as the separation distance between the two delta functions, we see that the spectral width is 2κ and therefore increases linearly with the temporal frequency and the propagation distance. At a time resolution of 10 ps and a total propagation distance of 3 mm, for example, the angular bandwidth is 2 rad^{-1}. A non-dimensional value of $\kappa = 2$ corresponds to $f_t d = 6$ mm/ps.

Further, note that in particular that as $\kappa \to 0$, the kernel approaches the delta function for all f_ν, so that the final propagated light field approaches the time-independent case. This simplification occurs for either d or f_t vanishing (corresponding to no propagation or considering only the steady-

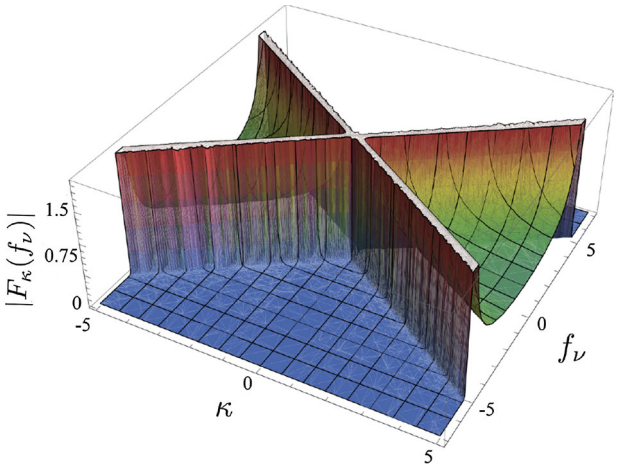

Fig. 5 Three-dimensional visualization of the frequency propagation kernel F_κ (also see Fig. 4, *lower left*). This kernel models the convolution along the angular light field dimension induced by propagation of a time-resolved light field in free space. The kernel depends on the propagation distance, the speed of light, and the temporal frequency, as modeled by the parameter $\kappa = f_t d/c$. The time range of this figure is 0 - 50ps

state component $f_t = 0$, respectively), or for the speed of light c going to infinity (corresponding to neglect of time-resolution). With any of these assumptions, the system (Eq. 4) reduces to the commonly used steady-state Fourier propagator (Eq. 3) (derived in the Supplement).

The convolution operation for the time-resolved case, however, generally "blurs" information in the convolved dimension. We analyze this behavior in more detail in Sect. 4 and demonstrate in Sect. 5 that this effect can be directly exploited to design a novel, lensless computational imaging system. Here, however, the kernel profile depends on not only the convolved dimension, but also the parameter $f_t d/c$. As shown from Eq. 5 and Fig. 5, there are two regimes to consider. Whereas the decay of F_κ is exponential for $|\kappa| < |f_\nu|$ (so that information located there cannot be detected to first order), the falloff is much slower for $|\kappa| > |f_\nu|$ (where information can be detected). At the boundary, $\kappa = f_\nu$, and F_κ becomes singular. We call this boundary the "light cone[1]" and use it to distinguish between these two regimes. Light propagation of a point light source in the primal and frequency domain is illustrated in Fig. 4.

3.2 5D Frequency Analysis

Consider now a fully five-dimensional light field, with lateral variables x and y, respective angular variables ν_x and ν_y, and time t. Free-space propagation is generalized straightforwardly:

$$l(\mathbf{r}_\perp, \boldsymbol{\nu}_\perp, t) = l_0(\mathbf{r}_\perp - \boldsymbol{\nu}_\perp d, \boldsymbol{\nu}_\perp, t - \tfrac{d}{c}\sqrt{1+\nu^2}), \tag{6}$$

where $\boldsymbol{\nu}_\perp = (\nu_x, \nu_y)$, $\mathbf{r}_\perp = (x, y)$, and $\nu^2 = \nu_x^2 + \nu_y^2$. Here, we will use the same notation, l and $\tilde{l}$, for the light field and its spectrum as we did in the 3D case, distinguishing them by their arguments. The 5D Fourier transform is calculated in the same way (see the Supplement).

$$\begin{aligned}\tilde{l}(\mathbf{f}_x, \mathbf{f}_\nu, f_t) &= \int \dots \int d\mathbf{r}_\perp d\boldsymbol{\nu}_\perp dt e^{-2\pi i(\mathbf{f}_x\cdot\mathbf{r}_\perp + \mathbf{f}_\nu\cdot\mathbf{r}_\perp + f_t t)} \\ &\quad \times l_0(\mathbf{r}_\perp - \boldsymbol{\nu}_\perp d, \boldsymbol{\nu}_\perp, t - \frac{d}{c}\sqrt{1+\nu^2}) \\ &= \tilde{l}_0(\mathbf{f}_x, \mathbf{f}_\nu + \mathbf{f}_x d, f_t) \underset{f_{\nu_x}, f_{\nu_y}}{*} F^{(2)}_{\frac{f_t d}{c}}(|\mathbf{f}_\nu|),\end{aligned} \tag{7}$$

where the convolution is two-dimensional, over both angular frequency dimensions, with $\mathbf{f}_\nu = (f_{\nu_x}, f_{\nu_y})$. Here, $F^{(2)}_{\frac{f_t d}{c}}(|\mathbf{f}_\nu|)$ is defined as (derived in the Supplement)

$$F^{(2)}_\kappa(|\mathbf{f}_\nu|) = \begin{cases} (2\pi i + \frac{1}{\sqrt{\kappa^2 - |\mathbf{f}_\nu|^2}}) \frac{\kappa e^{-2\pi i\sqrt{\kappa^2 - |\mathbf{f}_\nu|^2}}}{2\pi(\kappa^2 - |\mathbf{f}_\nu|^2)}, & \text{for } \kappa^2 > |\mathbf{f}_\nu|^2 \\ (2\pi + \frac{1}{\sqrt{|\mathbf{f}_\nu|^2 - \kappa^2}}) \frac{i\kappa e^{-2\pi i\sqrt{|\mathbf{f}_\nu|^2 - \kappa^2}}}{2\pi(|\mathbf{f}_\nu|^2 - \kappa^2)}, & \text{for } \kappa^2 < |\mathbf{f}_\nu|^2 \end{cases}. \tag{8}$$

Note here that $F^{(2)}_\kappa(|\mathbf{f}_\nu|)$ is axially symmetric; however, the convolution in Eq. 7 cannot be reduced to a single, axial integral for a general light field unless it is radially symmetric. Further, the 5D kernel is structurally similar to the 3D version: two regions, one decaying exponentially and one algebraically, separated by a singularity along $\kappa^2 = |\mathbf{f}_\nu|^2$. However, because the 5D function is axially symmetric, this singularity is a ring of radius κ in the (f_{ν_x}, f_{ν_y}) plane.

Thus, time-resolved imaging can be placed naturally in the frequency analysis of light transport. Similar to steady state analysis (Durand et al. 2005), transient frequency analysis can be built upon the results here so as to include, e.g. diffuse reflections and shading, occlusions and interactions with BRDFs.

4 Cross-dimensional Information Transfer and Space–Time–Angle Bandwidth Analysis

The Fourier analysis in the previous section shows that there is a transfer of information between the different dimensions; space, time, and angle are all coupled. In this section, we explore the details of this process. Because current ultrafast cameras in the picosecond range have only one spatial dimension, we focus on the 3D subspace of the 5D kernel. We exploit this information transfer by proposing an ultrafast lensless sensor that captures only space and time. The

[1] The concept of light cones is commonly used in space–time physics; see e.g. www.phy.syr.edu/courses/modules/LIGHTCONE/minkowski.html.

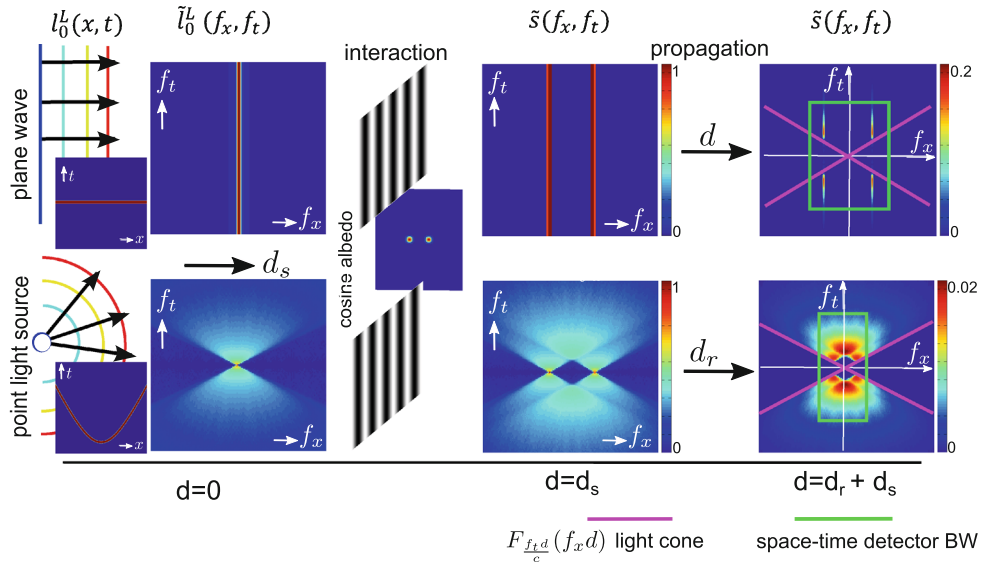

Fig. 6 Time-resolved bandwidth analysis of a Lambertian scene of single spatial frequency f_0. *Top*: planar impulse illumination. Immediately after interaction, the spectrum copies the incident spectrum at $\pm f_0$. After propagation, the signal is attenuated along the f_x direction, but is preserved in time. The green line indicates the space–time bandwidth of the sensor, and the pink line represents the propagation MTF light cone (slopes are $\pm c$). Here, spatial resolution must be greater than f_0 for detection. *Bottom row*: point pulse illumination. After propagation, significant information is located along the f_t direction, so that spatial bandwidth can be relaxed, for instance information can be detected even if the spatial sensor bandwidth f_x^{max} is low. (Please note that we use the side view of the cosine patten just to show a perspective view of the scene plane, and looks like the light is going through it for better understanding of the scene.)

recorded data, called a streak image $s(x,t)$, contains angular information that can be used to infer image and depth of the scene. Such streak sensors are currently commercially available (e.g., Hamamatsu 2012).

We demonstrate and evaluate the proposed computational camera architecture in Sect. 5 and derive its limits in Sect. 6. We also analyze the bandwidth and reconstruction limits of any streak imaging system in this section. In particular, we write $s(x,t)$ (Fig. 2, right) as:

$$s(x,t) = \int l(x,v,t)\,dv = \int l_0(x - vd, v, t - \frac{d}{c}\sqrt{1+v^2})\,dv. \quad (9)$$

Equivalently, the streak spectrum is found by setting the DC component (f_v) to zero.

To simplify the analysis, we consider a Lambertian scene, so that the light field becomes $l_0(x,v,t) = l_0^L(x,t)$, i.e. it is angle-independent. Using Eq. 4, we can write the measured streak image as (derived in the supplement):

$$\tilde{s}(f_x, f_t) = \tilde{l}_0^L(f_x, f_t) F_{\frac{f_t d}{c}}(f_x d). \quad (10)$$

Physically, the final spectrum is the original light field spectrum multiplied by a space–time-dependent weighting factor, whose characteristic width is determined by the propagation distance d. Thus, the kernel F_κ acts as an optical transfer function (OTF). This equation allows us to derive an upper bound on the spatial frequencies f_x of a Lambertian plane at distance d that can be recovered from a streak image $s(x,t)$. As an illustration of Eq. 10, we see in Fig. 6 (bottom) the initial transient point light source $\tilde{l}_0^L(f_x, f_t)$ propagating a distance d_s in the frequency domain under Lambertian assumptions. The effects of different time scales could be understood simply by varying window height in the f_t axis.

To gain some intuition, consider the simplest case of a cosine albedo (of spatial frequency f_0) that is illuminated by a planar impulse: $l_0^L(x,t) = cos(2\pi f_0 x)\delta(t)$, as shown in Fig. 6. (Experimentally, this setup could be effected by illuminating the object with a pulsed point source whose origin is far from the object so that the object sees an approximately planar wavefront.) Note that here, and in the theory below, we assume the scene to be planar, but this is not a necessary condition and can be relaxed. Immediately after the mask, the Fourier transform is a pair of vertical delta functions, centered at $\pm f_0$. After propagation, the spectrum is multiplied by the OTF above. For $c|f_x| > |f_t|$, $F_{\frac{f_t d}{c}}(f_x d)$ decays exponentially, so that, to first order, information in these regions is lost. For $c|f_x| < |f_t|$, the information is preserved. This behavior is shown in Fig. 6 (top right).

In order to detect the signal, its space–time bandwidth must lie within that of the camera. Because there is no cross-dimensional information transfer, i.e. spatial frequencies are

not sheared into time, the resolution criteria is the same as for conventional imaging: the detector's spatial bandwidth f_x^{max} must exceed the signal frequency: $f_0 < f_x^{max}$, regardless of the time bandwidth.

In contrast, consider the same scene illuminated by a spherical wavefront pulse. For simplicity, we assume the point source is located on the z-axis a distance d_s from the scene (Fig. 6). The initial light field becomes $l_0^L(x,t) = cos(2\pi f_0 x)\delta(t-(d_s/c)\sqrt{1+(x/d_s)^2})$. Now, the light field's Fourier transform is the convolution of the transforms of the factors. After propagating a distance d_r, the resulting streak spectrum (derived in the supplement) is

$$\tilde{s}(f_x, f_t) = F_{\frac{f_t Z}{c}}(Z(f_x + d_s f_0/Z)) + F_{\frac{f_t Z}{c}}(Z(f_x - d_s f_0/Z)), \quad (11)$$

where $Z \equiv d_s + d_r$. When plotted in the f_x-f_t plane (see Fig. 6, bottom), we see a strong signal in the time domain (i.e. along the vertical axis), as we would expect from a spherical pulse, which illuminates different scene points at different times. In this case, the scene frequency is recovered if f_t^{max} is large enough (roughly $f_t^{max} > \frac{d_s c}{Z} f_0$), *regardless of the spatial bandwidth*. Thus, a sufficiently fast camera can super-resolve object features.

Note that in the limiting case of $f_t^{max} \to 0$, the detector no longer detects time-resolution. In this case, one will measure a signal only if its space bandwidth is greater than the signal's bandwidth *and* if the propagation distance is sufficiently small so as to have the signal fall within $\sim Z^{-1}$ of the propagator singularity. For $f_t^{max} \sim 0$, the argument of the propagator in Eq. 11 is $Z(f_x - d_s f_0/Z)$, so that spatial frequencies beyond $(1 + f_0 d_s)/Z$ cannot be resolved. For spatial features of 1 mm, a total propagation distance of 1 cm, and a d_s value of 0.5 cm, this cutoff is approximately 0.6 mm^{-1}. Essentially, this is the configuration for conventional (steady state) white light imaging of out-of-focus images. With the sensor placed a short distance away, the features of an image are blurred into the noise.

Using these considerations, we expect that time-resolved measurements can be used to extract scene information otherwise unavailable. In particular, using only a single pixel, which corresponds to the $f_x = 0$ axis in Fig. 6, we still see information as a function of time, so we could potentially recover the (spatial) scene texture, though a first principles calculation has not yet been.

5 Lensless Computational Imaging Using Inverse Time-Resolved Light Transport

While the previous sections analyzed time-resolved forward light propagation and its cross-dimensional correlations, including upper bounds on preserved information, we now proceed with the inverse problem. The frequency analysis and bandwidth will help us define a modulation transfer function with which to compare our reconstruction results to ground truth values. Given a streak image, that is space–time scene information, captured by a bare sensor with no additional optical elements, such as lenses or attenuators, we aim at reconstructing image and depth of the scene. This is an ambitious goal; conventional cameras can only resolve a two-dimensional image without any depth information. A light field camera, recording space and angle x-ν, provides sufficient information to refocus the photograph in post-processing (Isaksen et al. 2000); unfortunately, this usually comes at the cost of a significant drop in spatial image resolution due to the employed lenslet arrays. As derived in the previous section, even a low-resolution bare sensor is theoretically capable of resolving image information with a higher spatial resolution, i.e. with super-resolution. For this application, we make the following assumptions for the remainder of this section: a Lambertian scene (the scene is not necessarily planar), exhibiting no significant global illumination effects or occlusions, is illuminated with a pulsed point light source and recorded with a bare sensor that provides a temporal resolving power in the lower picosecond range.

Under these conditions, the time-resolved rendering equation (Kajiya 1986; Smith et al. 2008) can be written as

$$s(\mathbf{x},t) = I_0 \int \frac{1}{(r_s\, r_l)^2} \rho\left(\mathbf{x}'\right) \cos\theta_s \, \cos\theta_l \, \delta\left(t - \frac{r_s + r_l}{c}\right) d\mathbf{x}', \quad (12)$$

where the recorded streak image s is restricted in space $\mathbf{x} = (x_x, x_y, x_z)^T$ to the plane $(x_x, x_y, 0)^T$ and $\rho(\mathbf{x})$ describes the diffuse albedo of a point.

As illustrated in Fig. 7, the distances between a scene point and a pixel, as well as between the source and a scene point, are, respectively, r_s and r_l, both of which depend implicitly on $\mathbf{x}$. The cosine terms model the angle between surface point normal and incoming or outgoing light direction, whereas the δ function describes the time-dependency of the system introduced by the pulsed point light source.

Equation 12 is a tomographic reconstruction problem (Kak and Slaney 2001), as each pixel in the streak image measures a weighted integral over the unknown surface albedos in the scene. As opposed to conventional line tomography, the integration surfaces in this application are elliptical surfaces as illustrated by the dashed red line in Fig. 7. This is similar to range tomography in seismic imaging (Saleh 2011) with the significant difference that no phase information is captured in our application, making the inversion problem more difficult.

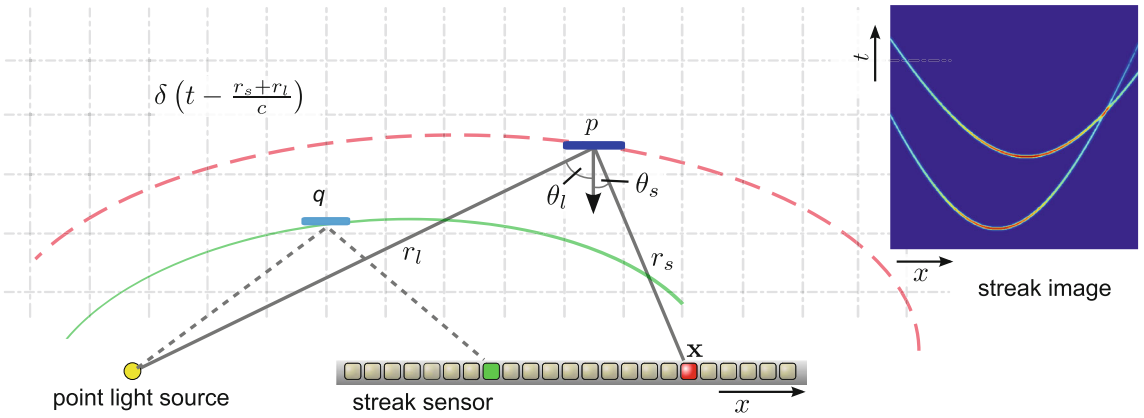

Fig. 7 One-dimensional illustration of tomographic scene reconstruction. A lensless streak sensor (*center*) records the time of arrival of photons emitted by a pulsed point light source (*left*). The captured streak image for two white patches in the scene is illustrated on the right. For an unknown scene, however, the contribution at each sensor pixel (x, t) is ambiguous as it may contain contributions from any point in the scene that is located along an ellipse whose locus of points have identical path lengths from source to detector, illustrated by the dashed red line (left)

While filtered backprojection is one of the most popular techniques for computed tomography, the algorithm requires projections that vary over the full hemisphere. Our application is more similar in spirit to limited baseline tomography, where projections over only a limited set (in our case the finite area of the sensor) are available. For these problems, algebraic reconstructions techniques are commonly applied. Follow (Herman 1995), we employ iterative algebraic reconstruction methods.

For this purpose, we model the scene containing an unknown geometry and surface albedos as a basis expansion:

$$\rho\left(\mathbf{x}\right) = \sum_{k=1}^{N} \rho_k \phi_k(\mathbf{x}), \tag{13}$$

where ρ_k are the coefficients and the basis functions are $\phi_k(\mathbf{x}), k \in 1 \ldots N$, i.e. the basis functions are discrete voxels. This notation allows us to write Eq. 12 as

$$s(\mathbf{x}, t) = I_0 \sum_{k=1}^{N} \rho_k \int \frac{1}{\left(r_s\, r_l\right)^2} \phi_k\left(\mathbf{x}'\right) \cos\theta_s \cos\theta_l \times \delta\left(t - \frac{r_s + r_l}{c}\right) d\mathbf{x}'. \tag{14}$$

In practice, this formulation is discretized into a linear system $\mathbf{s} = \mathbf{P}\boldsymbol{\rho}$, where each element of $\mathbf{s}$ is,

$$s_{xt} = \sum_{k=1}^{N} \rho_k P_{xt}^{(k)}, \tag{15}$$

and we approximate the elements of matrix $\mathbf{P}$ as

$$P_{xt}^{(k)} = \frac{I_0}{\left(r_{xk}\, r_{lk}\right)^2} dA_k \cos\theta_{xk} \cos\theta_{lk}\, \zeta_t\left(\frac{r_{xk} + r_{lk}}{c}\right). \tag{16}$$

Inspired by surflets (Chandrasekaran et al. 2004), we model the unknown scene geometry as a collection of planar patches of size dA_k, facing the bare sensor. The distance from patch k to sensor pixel x is r_{xk}, r_{lk} is the distance from patch k to the light source, and $\zeta_t(\cdot)$ maps the travel time of a photon to the nearest time slot in the streak image. Using this formulation, the final optimization problem is formulated as

$$\underset{\rho}{\text{minimize}}\ \|\mathbf{s} - \mathbf{P}\boldsymbol{\rho}\|_2^2, \ \text{subject} \quad \text{to}\ 0 \le \rho \le 1, \tag{17}$$

which can be solved with standard approaches of constrained linear optimization as discussed in Sect. 7.1.

6 Depth of Field Analysis

In this section, we derive the upper bounds of the proposed optical acquisition system and evaluate the performance of the combined optical and computational camera design.

In conventional optical systems, the performance of a lens or image quality measured by a sensor through some optical assembly is mostly derived as the modulation transfer function (MTF). The MTF is defined as the magnitude of the optical transfer function and provides a measure for contrast and resolution of an imaging system. Intuitively, the MTF gives the relative loss of energy of a diffuse target for a specific spatial frequency. The desired bound of any optical system is the diffraction limit; in practice, however, this is not always achieved, and oftentimes energy is optically lost for higher spatial frequencies, resulting in a decreased image contrast.

The proposed optical setup consists of a bare sensor without any refractive optical elements. Intuitively, we expect a planar target to be perfectly reconstructed when located

directly on the sensor. As the target moves away from the sensor, the image information is computationally extracted from the measurements, but possibly at a lower quality. The final image quality depends on both the optical system and the reconstruction algorithm—we proceed by evaluating the upper bounds provided by the proposed optical system using a modified MTF, followed by an evaluation of the reconstructed image quality.

Motivated by modulation transfer functions in conventional imaging systems, we define an MTF for the proposed lensless imager as

$$\mathrm{MTF}_{d,f_0} = \int \left| \frac{\tilde{s}\,(f_0, f_t)}{\tilde{l}_0^L\,(f_0, f_t)} \right| df_t = \int \left| F_{\frac{f_t d}{c}}\,(f_0 d) \right| df_t, \quad (18)$$

where the limits of integration correspond to the temporal bandwidth of the sensor, 1/(20 ps). Following Eq. 10, this formulation measures the relative loss of energy for a diffuse plane textured with a pattern that exhibits a single spatial frequency f_0. The Fourier transform (FT) of the light field emitted by the plane is $\tilde{l}_0^L$, whereas $\tilde{s}$ is the FT of the streak sensor image measuring spatial and temporal light variation of $\tilde{l}_0^L$ after a propagation distance d. As discussed in Sect. 4, the special case of a planar object imaged at a distance to the sensor reduces the convolution, in the frequency domain, with the propagation kernel F_κ to a multiplication. This simplification allows us to write the MTF in Eq. 18 (right) as the integral over the magnitudes of the kernel's temporal frequencies.

For finite limits, Eq. 18 cannot be calculated in closed form. Figure 8 plots the MTF for an optical setup with dimensions described in Sect. 7. We observe that, similar to conventional optical systems, higher spatial frequencies of a plane at a fixed distance d exhibit a stronger loss of energy than low frequencies. Similarly, as the distance between the sensor and the plane increases, the MTF decreases, indicating that the expected reconstruction quality of objects at larger distances would decrease. In particular, for a time bandwidth lower than the spatial frequency under investigation, the MTF falls off strongly, because the integral in Eq. 18 includes only the exponentially small region of F_κ. Furthermore, the MTF does not take into account the finite sensor size, which would act as a low pass filter, or the reconstruction algorithm. Each point in the scene generates a curved wavefront, but due to the finite sensor size, only a portion of it is detected. For smaller sensors, the tails of the hyperbolas are lost, so that it becomes challenging to separate different hyperbolas.

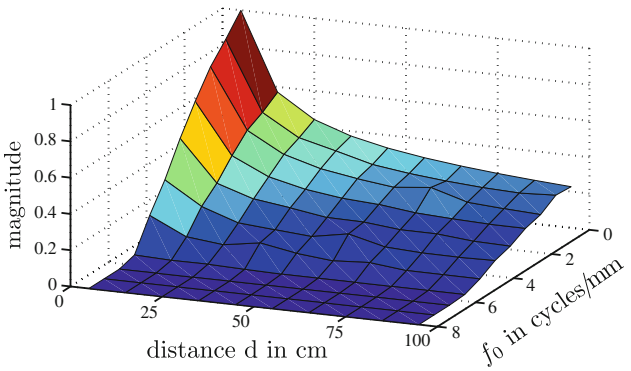

Fig. 8 Modulation transfer function (MTF). Numerically evaluating Eq. 18 allows for an upper bound on the expected reconstruction quality to be placed. The MTF measures the relative preserved energy of a diffuse plane at a distance d to a bare image sensor. The plane is textured with a pattern that exhibits a single spatial frequency f_0. The lower the MTF magnitudes, the more energy is optically lost, making a computational reconstruction increasingly difficult and sensitive to noise. We observe that the MTF magnitudes for a fixed spatial frequency decrease with increasing distance to the sensor. A similar decrease is observed for increasing spatial frequencies at a fixed distance to the sensor

We evaluate the quality of the proposed computational camera design, including optical acquisition and tomographic reconstruction, in Fig. 9. For this purpose, we simulate a resolution chart of size 96 × 96 pixels at various distances d, ranging from 0 to 100 cm, to the bare sensor. Reconstruction is performed by solving Eq. 17 as described in Sect. 7.1. For this experiment, the same dimensions are used as for the MTF in Fig. 8 and for the simulations in Sect. 7. As predicted by the MTF, reconstruction quality measured in peak-signal-to-noise ratio (PSNR) drops with an increasing distance d but also for increasing spatial frequencies at a fixed depth.

We conclude that the proposed computational imaging systems exhibits a finite depth of field, similar to conventional cameras, with the focal plane being located directly on the sensor. We continue in the following section to evaluate the performance of the proposed system using more complex scenes composed of textured objects at varying distances to the sensor. We also show initial results of a prototypical ultrafast computational camera.

7 Experimental Results

In this section, we show four different experiments of the proposed lensless imaging system. The results in Figs. 10 and 11 simulate a streak sensor that captures a three-dimensional volume containing two spatial dimensions as well as temporal light variation. The experiments show reconstructions of only the diffuse albedos for a single plane at a distance to the sensor in Fig. 10 and for two planes at different distances in Fig. 11. In the latter case, the depth is estimated along with the diffuse albedos. In the experiments shown in Figs. 10 and 11, the laser is simulated to be located in the center of the optical axis, whereas the laser is aimed at the center of the diffuser for the real experiment (Fig. 12, left). Figures 12 and 13 show reconstructions of simple two-

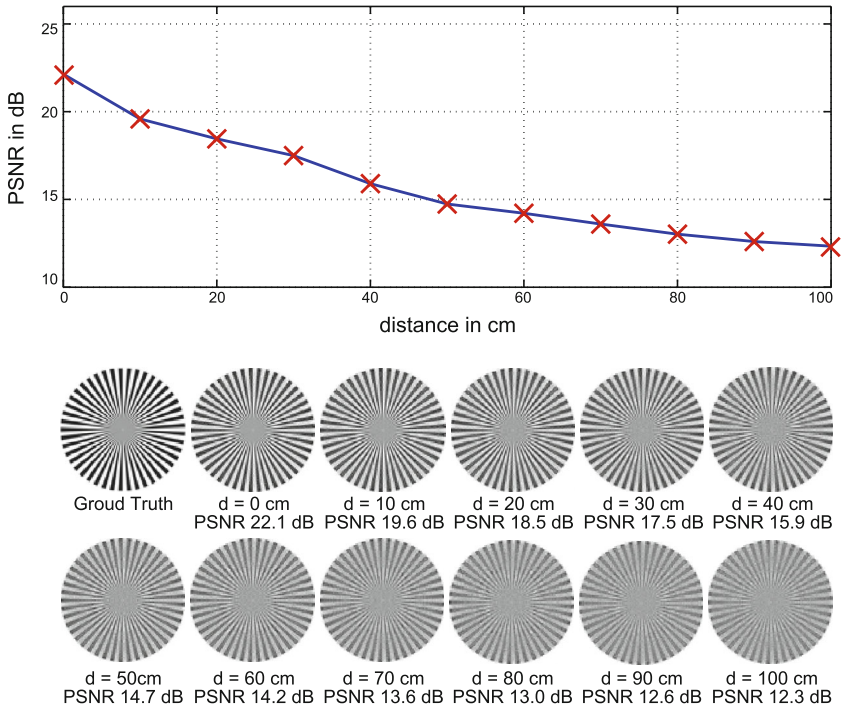

Fig. 9 Depth of field (DOF). This plot evaluates the combined quality of optical acquisition and tomographic computational reconstruction. A synthetic planar resolution chart (*center row, left*) is acquired at different distances to a bare sensor and subsequently reconstructed at the known depth. As predicted by the MTF (Fig. 8), reconstruction quality drops with an increasing distance between resolution chart and sensor

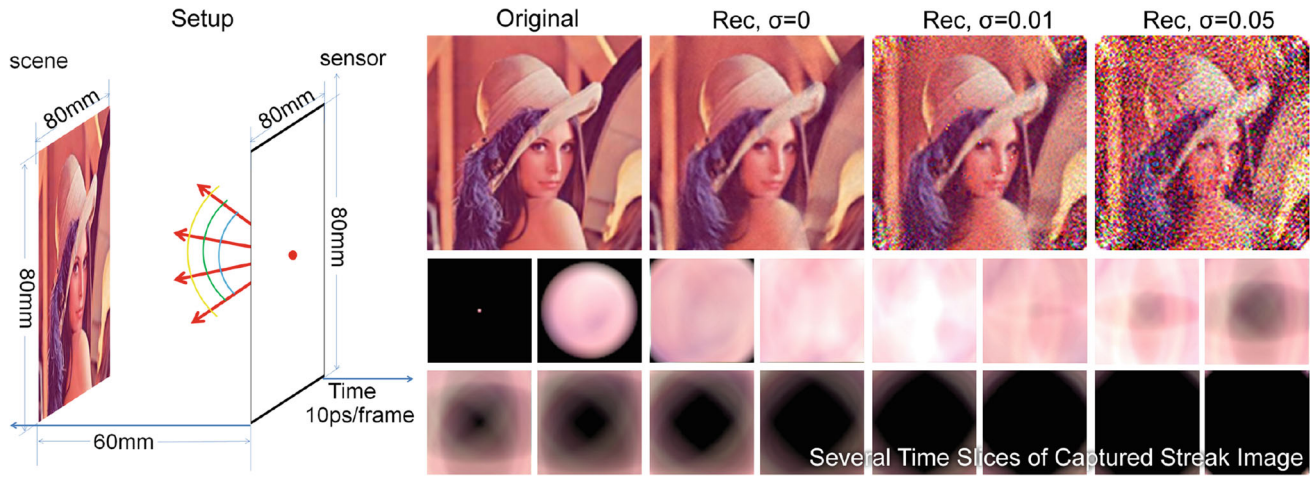

Fig. 10 Simulated reconstructions of a textured plane at a distance to the sensor. A bare sensor placed at a distance to the scene (*left*) captures a streak image, which is a series of 2D images (*bottom row*). The scene is reconstructed using computed tomography for varying levels of additive sensor noise (*top row*). The point source wavefront, colored from *blue* to *green* to *yellow*, represents increasing time (Color figure online)

dimensional scene patches from a streak sensor that only measures one spatial dimension along with temporal variation.

The results illustrated in Fig. 10 are captured with a simulated streak sensor that has a spatial resolution of 64 × 64 pixels and a size of 80 × 80 mm. This sensor records 32 time slots, each integrating over 10 ps (Fig. 10, bottom row). The scene is a diffuse plane located 60 mm away from the sensor and has the same size as the sensor, but a spatial resolution of 128 × 128 pixels. The top row shows reconstructions with varying levels of additive Gaussian sensor noise.

Figure 11 shows simulated results for a scene with a varying depth. As discussed in Sect. 5, this requires the tomographic reconstruction to estimate both diffuse albedo and scene depth simultaneously. In this experiment, the spatial resolution of each of the two depth planes is 96 × 96 pixels, and the number of time slots recorded is increased to 38 (with the same exposure time of 10 ps). The reduced quality of these results can be entirely attributed to the occlusion

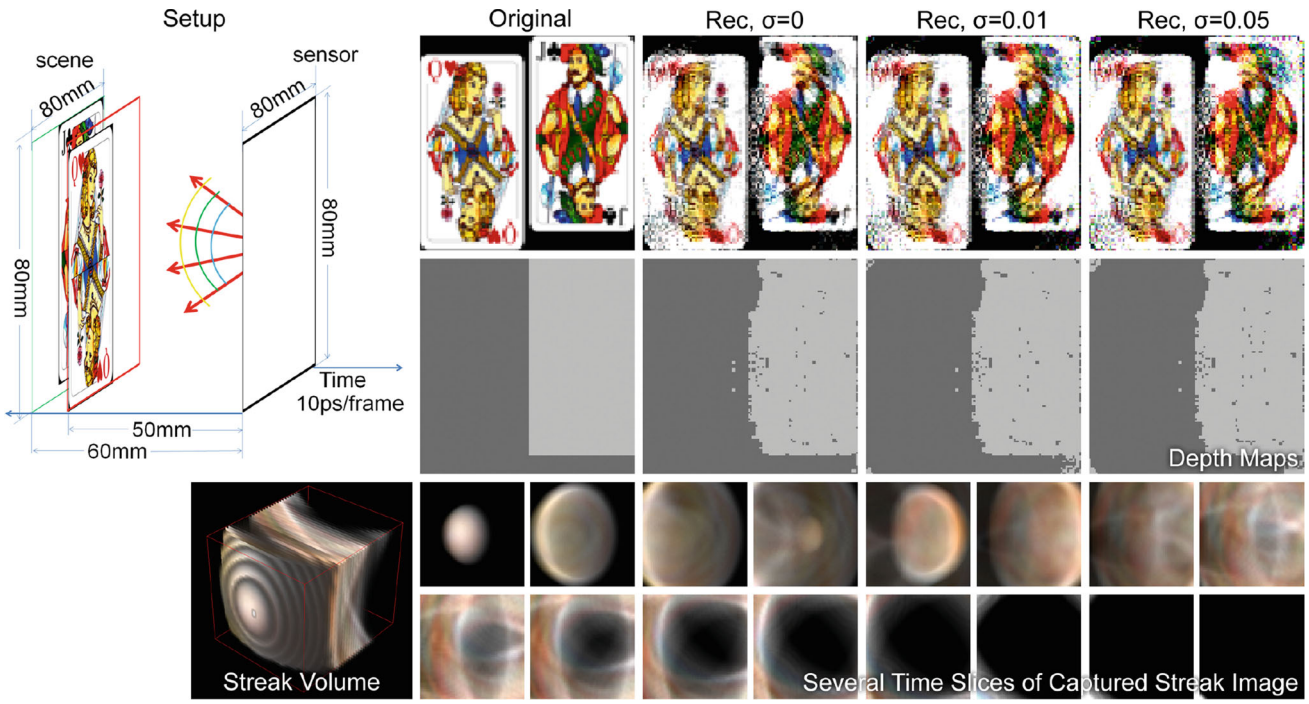

Fig. 11 Simulated results for reconstructing image and depth of an unknown scene captured with a bare sensor. Several time slices and a volumetric rendering of the recorded streak images are shown in the *bottom row*. While the scene depth can be reconstructed relatively accurately (*center row*), even small amounts of occlusion in the scene, as exhibited by the two playing cards, result in global noise in the reconstructed diffuse albedos (*top row*). The point source wavefront, colored from *blue* to *green* to yellow, represents increasing time

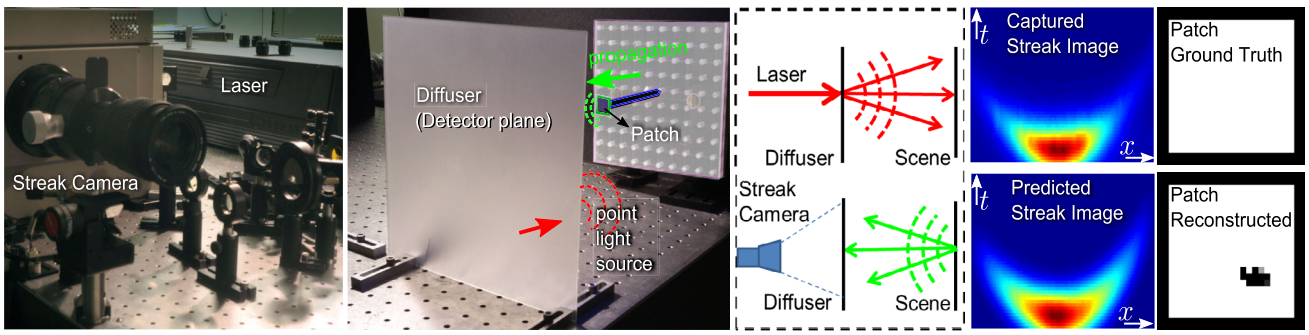

Fig. 12 The prototype camera setup shows the streak camera and the laser (*left*) and the scene (*center left*). Due to physical constraints, we model a bare sensor in space by imaging the light scattered onto an optical diffuser. A schematic illustration of the setup (*center right*). Experimental and simulated streak images agree (*right*). Successful reconstruction of a simple, white 2D patch from a 1D sensor is shown on the right

between the cards in the scene. Even a small amount of occlusion, as exhibited in this experiment, results in a global increase in reconstruction noise. As stated in Eq. 12, the goal of the proposed reconstruction algorithm is to recover a three-dimensional volume with albedos. For scenes containing opaque objects, this directly results in the depth of the scene, with the upper limit being the MTF above.

Real results captured with a one-dimensional streak camera (see Sect. 7.2) are shown in Fig. 12. The intricate design of this camera prevents us from placing the bare sensor directly in the scene; we emulate this setup by recording a scene patch behind an optical diffuser instead. Challenges in this experiment are manifold: intensity variations and temporal jittering of the laser along with vignetting and non-linear temporal distortions in the streak camera result in slight variations of the recorded signal as compared to the predicted data using known scene geometry and albedo for this simple scene. Furthermore, a two-dimensional scene at a calibrated distance is measured with a one-dimensional sensor, which makes the reconstruction particularly difficult. Nevertheless, we show successful reconstruction of a simple scene—a diffuse white patch—in Fig. 12. The white patch is a two-dimensional object imaged with a one-dimensional streak sensor. Although not immediately obvious, this time-

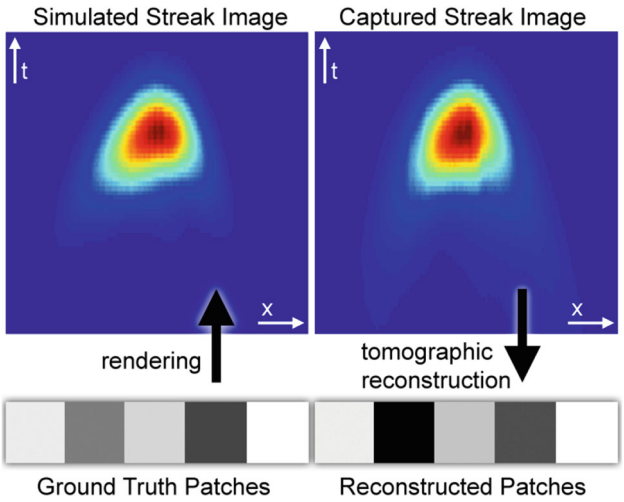

Fig. 13 Experimental results. A one-dimensional barcode consisting of five patches with different albedos (*lower left*) is placed at a distance behind a diffuser and recorded with the streak camera prototype (*upper right*). The tomographic reconstruction of the albedos (*lower right*) exhibits slight errors in the darker patches, which are due to imperfect calibration of the prototype (difference between simulated and captured streak images, *top row*)

resolved 1D image captures sufficient information for the simple 2D scene to be recovered. This is possible because of the cross-dimensional information transfer, from spatial to space–time domains, introduced by light propagation in free space.

Figure 13 shows another result captured with the prototype ultra-fast camera. For this experiment, we placed a 1D barcode at a distance behind the diffuser (see setup in Fig. 12). The barcode consists of five patches with varying albedos (Fig. 13, lower left). Again, we observe slight differences between the simulated streak image rendered from the known scene geometry and reflectance (Fig. 13, upper left) and the data recorded with our prototype camera (Fig. 13, upper right). The albedos of the five patches are successfully reconstructed from the captured data (Fig. 13, lower right), yet with some errors in the two darker patches that are due to the mismatch of expected and recorded streak image.

7.1 Software Implementation

We solve Eq. 17 using the simultaneous algebraic reconstruction technique (SART) (Kak and Slaney 2001). This is an iterative approach to solve computed tomography problems. In our Matlab-based implementation, we run 10,000 iterations for each of the reconstructions in Figs. 10, 11, 12, and 13. The large number of iterations is required due to the slow convergence rate of SART. The computations are performed with a sparse representation of matrix $\mathbf{P}$ in Eq. 17. Computing times are approximately one to two hours for all three color channels of each of the results. The computation time is approximately 3 min for 5,000 iterations for the real experimental data.

7.2 Hardware Setup

For the captured result shown in Fig. 12, the illumination source is a femtosecond (fs) Titanium:Sapphire laser, which produced 50 fs long pulses, centered at 795 nm at a repetition rate of 75 MHz. The beam is focused onto the diffuser wall (ground glass). The streak camera is a Hamamatsu C5680, which is used to detect one spatial dimension of the diffuser with a time resolution of 15 ps and quantum efficiency of about 10%. The camera position and viewing direction are fixed.

8 Discussion

In summary, we have presented a frequency analysis and a rigorous derivation of analytical expressions for the most fundamental process in light transport—propagation in free space, considering space, time, and angle. While prior models for the steady-state case intuitively describe light field propagation as a localized shear in both the primal and frequency domain, the hyperbolic curvature in space–time, introduced by differences in travel distances for different angles, delocalizes signal energy in the frequency domain. However, we demonstrate that it can be intuitively expressed as a combination of the well known shear in the light field and a convolution along the angular frequencies. This convolution optically blurs information between the different dimensions; we demonstrate that this information transfer can be computationally exploited using a novel, lensless computational imaging technique which has potential applications in biomedical imaging, thin mobile devices, remote sensing, and surveillance.

While the forward analysis presented is most general in considering space, time, and angle for arbitrary scenes, the inverse problem including bandwidth analysis and tomographic scene reconstruction restrict the imaged scene to be diffuse and to exhibit negligible amounts of global illumination effects and occlusions. The proposed sensing technique requires an ultra-short pulsed illumination, as for instance provided by a laser, and a streak camera capable of resolving time variations in the lower picosecond range. Currently available hardware achieving the required temporal resolution is restricted to sensing one spatial scanline per recorded image. The exposure time of the proposed system is significantly shorter than conventional imaging system; this fact in combination with the removed lens reduces the overall amount of light measured at each sensor pixel. Note that, for the spatial and temporal scales considered here, the geometric model is sufficient for modeling and reconstruction,

as shown previously (Velten et al. 2012; Wu et al. 2012). For features that approach the wavelength of light, however, diffraction must be incorporated through, e.g. transient generalizations of phase space and wave optics. Furthermore, it is unnecessary to resort to the particle (quantum) nature of light at this time scale. (Indeed picosecond optics is treated classically in e.g. fiber optic communications.) The quantum nature of light would come into play in low light conditions, for which photon statistics plays a role.

The main challenges of high-quality image reconstruction are currently dictated by the employed hardware setup. The streak camera is difficult to calibrate as nonlinear effects in laser intensity and image projection, are observed. These challenges limit experimental reconstructions to simple scenes consisting of textured planar patches.

Nevertheless, the work presented in this paper provides the fundamental analytic platform on which to re-examine many computer vision applications in light of a new degree of freedom: time-resolved image information. Further, it provides a method for building a consistent theory of time-resolved light field propagation. Future work will examine effects of optical elements other than propagation in free space, including lenses, diffusers, diffractive elements, and scattering media.

Acknowledgments The work of the MIT affiliated coauthors was funded by the Media Lab Consortium Members. Gordon Wetzstein was supported by an NSERC Postdoctoral Fellowship. Tsinghua University affiliated coauthors were supported by China National Basic Research Project (No. 2010CB731800) and the Key Project of NSFC (Nos. 61120106003 and 61035002).

References

Andersen, G. (2005). Large optical photon sieve. *Optics Letters*, *30*, 2976–2978.

Chai, J. X., Tong, X., Chan, S. C., & Shum, H. Y. (2000). Plenoptic sampling. *ACM Transactions on Graphics (SIGGRAPH)*, *19*, 307–318.

Chandrasekaran, V., Wakin, M., Baron, D., & Baraniuk, R. (2004). Surflets: A sparse representation for multidimensional functions containing smooth discontinuities. In: *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on, IEEE, Chicago* (p. 563).

Durand, F., Holzschuch, N., Soler, C., Chan, E., & Sillion, F. X. (2005). A frequency analysis of light transport. *ACM Transactions on Graphics (SIGGRAPH)*, *24*, 1115–1126.

Edward, R., Dowski, J., & Cathey, W. T. (1995). Extended depth of field through wave-front coding. *Applied Optics*, *34*, 1859–1866.

Gill, P. R., Lee, C., Lee, D. G., Wang, A., & Molnar, A. (2011). A microscale camera using direct fourier-domain scene capture. *Optics Letters*, *36*, 2949–2951.

Hamamatsu: Hamamatsu Streak Camera. (2012). http://learn.hamamatsu.com/tutorials/java/streakcamera/

Herman, G. T. (1995). Image reconstruction from projections. *Real-Time Imaging*, *1*, 3–18.

Isaksen, A., McMillan, L., Gortler, S. J. (2000) Dynamically reparameterized light fields. In: *Proceedings of the ACM Transactions on Graphics (SIGGRAPH), San Diego* (pp. 297–306).

Isikman, S. O., Bishara, W., Mavandadi, S., Yu, F. W., Feng, S., Lau, R., Ozcan, A. (2011). Lens-free optical tomographic microscope with a large imaging volume on a chip. In: *Proceedings of the National Academy of Sciences, Washington*.

Kajiya, J. T. (1986). The rendering equation. *ACM Transactions on Graphics (SIGGRAPH)*, *20*, 143–150.

Kak, A. C., & Slaney, M. (2001). *Principles of computerized tomographic imaging*. Tahoe: Society for, Industrial Mathematics.

Koppal, S. J., Gkioulekas, I., Zickler, T., Barrows, G. L. (2011). Wide-angle micro sensors for vision on a tight budget. In: *Proceedings of the CVPR, New York* (pp. 361–368).

Levin, A., & Durand, F. (2010). Linear view synthesis using a dimensionality gap light field prior. In: *Proceedings of the CVPR, New York* (pp. 1–8).

Levin, A., Hasinoff, S. W., Green, P., Durand, F., & Freeman, W. T. (2009). 4D frequency analysis of computational cameras for depth of field extension. *ACM Transactions on Graphics (Siggraph)*, *28*, 97:1–97:14.

Naik, N., Zhao, S., Velten, A., Raskar, R., & Bala, K. (2011). Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Transactions on Graphics (SIGGRAPH Asia)*, *30*, 171.

Ng, R. (2005). Fourier slice photography. *ACM Transactions on Graphics (SIGGRAPH)*, *24*, 735–744.

Pandharkar, R., Velten, A., Bardagjy, A., Lawson, E., Bawendi, M., & Raskar, R. (2011). Estimating motion and size of moving non-line-of-sight objects in cluttered environments. In: *Proceedings of the CVPR, New York* (pp. 265–272).

Raskar, R., & Davis, J. (2008). *5D Time-Light Transport Matrix: What Can We Reason about Scene Properties?*. Report: MIT Technical.

Saleh, B. (2011). *Introduction to subsurface imaging*. Cambridge: Cambridge University Press.

Smith, A., Skorupski, J., & Davis, J. (2008). Transient Rendering. Technical Report UCSC-SOE-08-26, School of Engineering, University of California, Santa Cruz.

Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., & Tumblin, J. (2007). Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Transactions on Graphics (SIGGRAPH)*, *26*, 69.

Velten, A., Willwacher, T., Gupta, O., Veeraraghavan, A., Bawendi, M., & Raskar, R. (2012). Recovering three-dimensional shape around a corner using ultra-fast time-of-flight imaging. *Nature, Communications*, *3*, 745–758.

Wu, D., O'Toole, M., Velten, A., Agrawal, A., & Raskar, R. (2012). Decomposing global light transport using time of flight imaging. In: *Proceedings of the CVPR, New York* (pp. 1–8).

Zomet, A., & Nayar, S.K. (2006). Lensless imaging with a controllable aperture. In: Proceedings of the CVPR, New York (pp. 339–346).