

## Recent advances in transient imaging: A computer graphics and vision perspective

Adrian Jarabo<sup>\*</sup>, Belen Masia, Julio Marco, Diego Gutierrez<sup>\*</sup>

Universidad de Zaragoza - I3A, Zaragoza 50018, Spain

### ARTICLE INFO

#### Article history:

Available online 19 January 2017

#### Keywords:

Transient imaging  
Ultrafast imaging  
Time-of-flight

### ABSTRACT

Transient imaging has recently made a huge impact in the computer graphics and computer vision fields. By capturing, reconstructing, or simulating light transport at extreme temporal resolutions, researchers have proposed novel techniques to show movies of light in motion, see around corners, detect objects in highly-scattering media, or infer material properties from a distance, to name a few. The key idea is to leverage the wealth of information in the temporal domain at the pico or nanosecond resolution, information usually lost during the capture-time temporal integration. This paper presents recent advances in this field of transient imaging from a graphics and vision perspective, including capture techniques, analysis, applications and simulation.

© 2017 Zhejiang University and Zhejiang University Press. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

### 1. Introduction

In 1964, MIT professor Harold Edgerton produced the now-iconic *Bullet Through Apple* photograph (see Fig. 1, left). His work represented an unprecedented effort to photograph events too fast to be captured with traditional techniques. He invented a new stroboscopic flash light (which he termed the stroboscope), which would shine for about 10 microseconds: bright enough, and short enough, to effectively freeze the world and capture ultrafast events such as the bullet bursting through the apple, a splash of a drop of milk, or the flapping wings of a hummingbird. Almost fifty years later, inspired by these images, the technique known as *femto-photography* (Velten et al., 2013) was introduced; it took Edgerton's vision to a whole new level, by allowing to capture movies of light in motion, as it traversed a macroscopic scene (Fig. 1, right).

This fifty-year span provides a clear example of the progress in ultrafast imaging. Many techniques have appeared in the last few years, some inspired by femto-photography, others following completely different approaches. They share the common goal of trying to make visible the invisible: Whether it is due to the speed of the event being captured, to the presence of scattering media, to the lack of photons, or to an occluding object, ultrafast imaging aims to leverage the wealth of information usually lost during

the capture-time temporal integration. This has revolutionized the fields of imaging and scene understanding, opening up new possibilities, but also discovering new challenges.


In this paper, we provide an in-depth overview of the most significant works in this domain. We concern ourselves mostly with works in the areas of computer graphics and computer vision; as such, we deal only with visible light and infrared. For other techniques that make use of different wavelengths (such as microwaves, or techniques operating in the terahertz domain), we refer the reader to other excellent sources such as the recent survey by Satat et al. (2016). Similarly, another recent survey (Bhandari and Raskar, 2016) offers an overview of the field from a signal-processing perspective. From our graphics and vision view, we adopt the commonly used term *transient imaging*, referring to imaging techniques fast enough to capture transient information of light transport, as opposed to traditional techniques that capture steady-state information (such as regular images).

We have structured our work as follows: First, we introduce **capture** techniques in Section 2, separating techniques that directly obtain transient information (such as the previously mentioned femto-photography, or the recent interferometry-based works), from techniques that *reconstruct* that information from a sparse set of measurements, usually sacrificing temporal resolution (such as recent approaches based on time-of-flight (ToF) cameras). In Section 3 we proceed to discuss works whose main goal is to **analyze** transient light transport, both in the primary and frequency domains. We additionally discuss techniques involving spatio-temporal coding and modulation. In Section 4 we offer a cross section of existing techniques from an **application's** point

<sup>\*</sup> Corresponding authors.

E-mail addresses: [ajarabo@unizar.es](mailto:ajarabo@unizar.es) (A. Jarabo), [diegog@unizar.es](mailto:diegog@unizar.es) (D. Gutierrez).

Peer review under responsibility of Zhejiang University and Zhejiang University Press.



**Fig. 1.** Left: In 1964, Harold Edgerton captured the iconic *Bullet Through Apple* image (© MIT Museum). The bullet traveled at about 850 m/s, which translated into an exposure of approximately 4–10 millionth of a second. Right: Almost 50 years later, the femto-photography technique was introduced (Velten et al., 2013), capable of capturing *light in motion*, with an effective exposure time of one trillionth of a second. The large split image is a composite of the three complete frames shown in the insets. Source: The complete videos of this and other scenes can be downloaded from <http://giga.cps.unizar.es/~ajarabo/pubs/femtoSIG2013/>.

of view. Again with a focus on graphics and vision, we subdivide this section in geometry reconstruction, motion estimation, and material estimation; a common problem in most of the applications discussed is the *multipath interference (MPI) problem*, which is tackled from many different angles. With the establishment of transient imaging, the **simulation** of time-resolved light transport is becoming an increasingly important tool, which we cover in Section 5. Last, Section 6 offers some final conclusions and discussions.

## 2. Capture

The interaction between light and matter is described as a linear operator by the light transport equation (Ng et al., 2003):

$$\mathbf{i} = \mathbf{T}\mathbf{p}, \quad (1)$$

where  $\mathbf{i}$  is the 2D image (as a column vector of size  $I$ ) captured by the camera,  $\mathbf{p}$  is the vector of size  $P$  representing the scene illumination, and  $\mathbf{T}$  is the scene transport operator encoded as a  $I \times P$  matrix. Equation (1) assumes that the light transport has reached steady-state. In its transient form (O'Toole et al., 2014), incorporating the temporal domain to the light transport equation yields:

$$\begin{aligned} \mathbf{i}(t) &= \int_{-\infty}^{\infty} \mathbf{T}(\tau)\mathbf{p}(t-\tau)d\tau \\ &= (\mathbf{T} * \mathbf{p})(t), \end{aligned} \quad (2)$$

where  $\mathbf{i}(t)$  stores the light arriving at time  $t$ ,  $\mathbf{p}(t)$  is the time-resolved illumination function at instant  $t$ , and  $\mathbf{T}(t)$  is the transport matrix describing the light transport with a time-of-flight of exactly  $t$ . Note that from here on all terms are time-dependent. The second equality represents the convolution in the temporal domain between  $\mathbf{T}$  and  $\mathbf{p}$ . In practice, the transient image cannot be captured at instant  $t$  directly, given physical limitations of the sensor. Instead, the signal is also convolved by the temporal response of the sensor  $\mathbf{s}(t)$  centered at  $t$  as:

$$\begin{aligned} \mathbf{i}(t) &= \int_{-\infty}^{\infty} \mathbf{s}(t-\tau)(\mathbf{T} * \mathbf{p})(\tau)d\tau \\ &= (\mathbf{s}(t) * \mathbf{T} * \mathbf{p})(t). \end{aligned} \quad (3)$$

For transient imaging, we are interested in computing the transient image  $\mathbf{i}(t)$  corresponding to the impulse response of the transport matrix  $\mathbf{T}$  (i.e. Eq. (2)). This would effectively mean that the illumination  $\mathbf{p}(t) = \delta_0(t)$  and sensor response  $\mathbf{s}(t) = \delta_t(t)$  are Dirac deltas centered on 0 and  $t$  respectively. In order to capture

this impulse response  $\mathbf{T}$ , several approaches have been presented, depending on the type of illumination and sensor response used. If we focus only on illumination, the main lines of work have used either impulse illumination, or coded illumination. In the case of the former, techniques have used either ultrafast imaging systems to directly record light transport (Section 2.1), or phase interferometry to recover the propagation of light (Section 2.2). In the case of the latter, the coded illumination has been usually correlated with the coded sensor response, allowing to recover the time-resolved response by means of post-capture computation (Section 2.3). A comparison of selected capture systems, including their spatio-temporal resolution, is summarized in Table 1.

### 2.1. Straight temporal recording

In theory, the most straightforward way to capture the impulse transport matrix  $\mathbf{T}$  is to use an imaging system with impulse illumination, and extremely short exposure times. However, this is very challenging in practice. First, the signal-to-noise ratio (SNR) is extremely low, since very few photons arrive during the exposure time. On top of this, ultrashort illumination pulses are required, to avoid the effect of the convolution on the transport matrix. Further, ultrafast imaging systems (in the order of nano to picosecond resolution) do not exist for high-resolution, two-dimensional imaging.

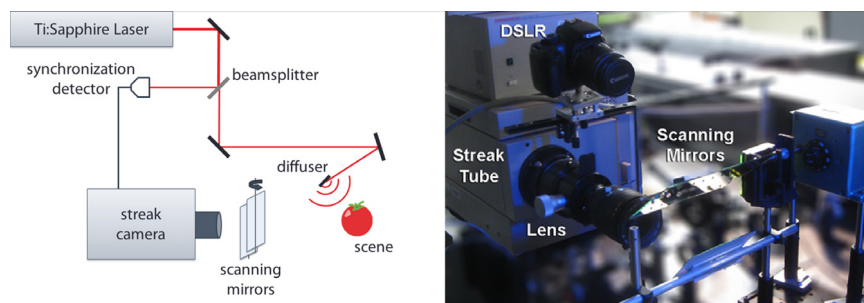
Ultrashort (impulse) illumination is in general achieved by using laser-based illumination, such as femtolasers (Velten et al., 2013). Several different approaches have been proposed to capture transient light transport, and to mitigate the challenge imposed by the extremely low SNR. In the following, we categorize these works according to the imaging system they are based on. Note that we concern ourselves with ultrafast imaging techniques focusing on time-resolved light transport, recording either a single bounce (e.g., for LIDAR applications), or multiple scattering. Other impulse-based ultrafast imaging techniques, e.g., based on pulse stretching (Nakagawa et al., 2014; Lau et al., 2016), are not discussed in this work.

Conceived for range imaging, *laser gated viewing* (Busck and Heiselberg, 2004) exploits the repeatability of light transport in a static scene by sequentially imaging a set of frames. An ultrashort laser pulse is synchronized with an ultrafast camera equipped with a highly sensitive CCD, which images photons arriving during very small temporal windows (in the order of a few hundred picoseconds). Each frame is computed independently, by sliding the imaging window. This system was later extended to use area impulse illumination in the context of non-line-of-sight imaging (Laurenzis and Velten, 2014). In order to improve the SNR, hundreds of

**Table 1**

Selected representative capture setups. Each different technology presents a variety of spatial and temporal resolutions, as well as a range of capturing times. Note that the work by [Laurenzis et al. \(2007\)](#) is targeted for long range depth reconstruction, which imposes a very long exposure time. Note also that [Li et al.'s approach \(2012\)](#) is able to reconstruct  $240^2$  pixels from a single captured one.

Work	Technology	Spatial Res.	Temporal Res.	Cost
<a href="#">Gkioulekas et al. (2015)</a>	Interferometry	$655 \times 648$	33 fs	several hours
<a href="#">Heshmat et al. (2014)</a>	Streak camera	$56 \times 12$	0.3–5.7 ps	6 s
<a href="#">Velten et al. (2013)</a>	Streak camera	$672 \times 600$	0.3–5.7 ps	2 h
<a href="#">Tadano et al. (2015)</a>	AMCW	$160 \times 120$	10 ps	Linear with temporal resolution
<a href="#">Gao et al. (2014)</a>	Streak camera	$150 \times 150$	10/20 ps	$10^{-11}$ s
<a href="#">Laurenzis and Velten (2014)</a>	Laser-gated	$1360 \times 1024$	66.7 ps	?
<a href="#">Garipey et al. (2015)</a>	SPAD	$32 \times 32$	67 ps	240 s
<a href="#">Peters et al. (2015)</a>	AMCW	$160 \times 120$	$\sim 100$ ps	0.05 s (real-time)
<a href="#">Kadambi et al. (2013)</a>	AMCW	$160 \times 120$	100 ps	4 s
<a href="#">Busck and Heiselberg (2004)</a>	Laser-gated	$582 \times 752$	100 ps	32K shots
<a href="#">Heide et al. (2013)</a>	AMCW	$160 \times 120$	1 ns	90 s (+ hours postprocess)
<a href="#">Li et al. (2012)</a>	Laser-gated	$240 \times 240$ (*)	20 ns	76K shots
<a href="#">Laurenzis et al. (2007)</a>	Laser-gated	$696 \times 520$	40 ms (*)	240 shots
<a href="#">Dorrington et al. (2007)</a>	AMCW	$512 \times 512$	0.5 mm (depth)	10–20 s
<a href="#">Hebert and Krotkov (1992)</a>	AMCW	$256 \times 256$ (scan)	0.98 cm (depth)	0.5 s



**Fig. 2.** Left: Diagram of the *femto-photography* setup. A laser, synchronized with the streak camera by means of a beam splitter and a synchronization detector, illuminates the scene after hitting a diffuser; the photons scattered towards the camera are imaged by the streak sensor, which captures a 1D video at picosecond resolution. In order to capture the full 2D time-resolved video, a set of rotating mirrors are used to scan the scene along the  $y$ -axis. Right: Photography of the setup, the DSLR is used to capture a conventional image of the scene.

Source: Image from [Velten et al., 2013](#)

measurements are required for each frame. Since each frame needs to be computed independently, gated imaging scales linearly with the number of frames. In order to improve convergence on range imaging (i.e. focusing on single scattering), approaches such as range-gates coding were developed, where the gate response is continuously modulated over a long time window, then reconstructed using intensity analysis ([Laurenzis et al., 2007](#); [Zhang and Yan, 2011](#); [Laurenzis and Bacher, 2011](#)), compressed sensing (CS) on the temporal domain by random temporal gating ([Li et al., 2012](#); [Tsagakatakis et al., 2012, 2013, 2015](#)), or hybrid approaches combining both techniques ([Zhang et al., 2012](#); [Dai et al., 2013](#)). This can reduce the number of measurements to just two, for 13-bit range images. An in-depth comparison of these approaches was done by [Laurenzis and Woiselle \(2014\)](#). Also relying on CS, [Li et al. \(2012\)](#) obtained full 2D transient images using a gated approach with a single pixel detector.


Systems based on arraying *avalanche photodetectors (APD)* ([Charbon, 2007](#)), and in particular *single photon avalanche diodes (SPAD)* ([Kirmani et al., 2014](#); [Garipey et al., 2015](#)), allow reducing significantly both capture times, and the required power of the light source, due to the single-photon sensitivity of the photodetectors. These systems are relatively simple and low-cost, and use eye-safe illumination. On the down side, they yield lower spatial resolutions, and require higher exposure times (in the order of tens of picoseconds).

In order to improve the temporal resolution, [Velten et al. \(2012b, 2013, 2016\)](#) used a streak camera ([Hamamatsu, 2012](#)) as imaging device. A streak tube sacrifices one spatial dimension (the  $y$ -axis) of the sensor, and uses it to encode the time of arrival of photons. This is done by transforming photons into electrons using a photocathode. The electrons are then deflected at different

angles as they pass through a microchannel plate, by means of rapidly changing the voltage between the electrodes. The CCD finally records the horizontal position of each pulse and maps its arrival time to the vertical axis. This effectively records a 1D video of transient light transport, with a temporal resolution of about two picoseconds. In order to record a 2D video, a rotating mirror progressively scans the scene along the vertical axis. [Fig. 2](#) shows a diagram of the setup, which they called *femto-photography*. As opposed to gated imaging, acquisition times are no longer linear with time resolution, but they scale linearly with the vertical resolution, although it requires repeated captures to get a decent SNR. In order to capture the three  $x$ – $y$ – $t$  dimensions simultaneously, [Heshmat et al. \(2014\)](#) encoded the  $x$ – $y$  spatial domain into a single dimension on the streak sensor, by using a tilted lenslet array. [Gao et al. \(2014\)](#), on the other hand, added a digital micromirror device with a pseudo-random binary pattern encoding the spatial dimension lost in the streak sensor. The transient image is then recovered using sparse reconstruction techniques.

## 2.2. Interferometry-based imaging

Interferometry-based imaging techniques rely on creating interference between electromagnetic fields. Not very many works exist in the fields of graphics and vision that take advantage of this methodology. Although its path length resolution is very high (higher than femto-photography by at least an order of magnitude), it presents limitations regarding field of view, depth of field and time of capture. Moreover, it is extremely sensitive to even micron-scale vibrations. Abramson introduced the first light-in-flight visualizations using holography, by illuminating a flat surface and a hologram plate with short pulses of light ([Abramson,](#)



**Fig. 3.** Michelson interferometer. The beamsplitter sends an input beam to two mirrors  $M_r$  and  $M_s$  at distances  $d_r$  and  $d_s$ . The split beams reflect back and recombine at the beamsplitter, then are imaged by the camera. Insets i and ii show two results with interference ( $d'_r \approx d_s$ ), and without ( $d_r \neq d_s$ ). Inset iii shows an alternative setup for oblique illumination.

Source: Image from Gkioulekas et al., 2015.

1978, 1983). Recently, Gkioulekas et al. (2015) presented an optical assembly capable of achieving a resolution of 10 microns ( $\sim 33.35$  femtoseconds) by using an interferometry approach inspired in optical coherence tomography (see Fig. 3). At this scale, the authors can visualize effects such as dispersion or birefringence. Kadambi et al. (2016a) introduced macroscopic interferometry, relying on frequency sampling of the ToF data. A key advantage of this novel approach is its implicit resolution of the MPI problem (Section 4.1.2); instead of having to disentangle phases as in traditional ToF (which is a hard, non-linear inverse problem, see Sections 2.3 and 4.1.2), the authors recast the problem as a summation of varying frequencies. In addition, the technique is more robust at low SNR levels. On the downside, the resolution of the system is limited to meter-scale ranges.

### 2.3. Phase time-of-flight

Phase-based time-of-flight (P-ToF) imaging, also called correlation-based time-of-flight (C-ToF) imaging or simply ToF imaging, cross-correlates emitted modulated light with frequency  $g_{\omega_T}$ , and the impulse response of a pixel  $\alpha_p$ , modulated and integrated at the sensor with frequency  $f_{\omega_R}$  (see Fig. 4). In its most typical continuous form (also known as *amplitude modulated continuous wave (AMCW)* systems<sup>1</sup>), the camera computes the cross-correlation as:

$$c(t) = \mathbf{s}(t) * \mathbf{p}(t), \quad (4)$$

with  $\mathbf{s}(t)$  the radiance received at the sensor, and  $\mathbf{p}(t)$  the emitted signal. These are in general modeled as:

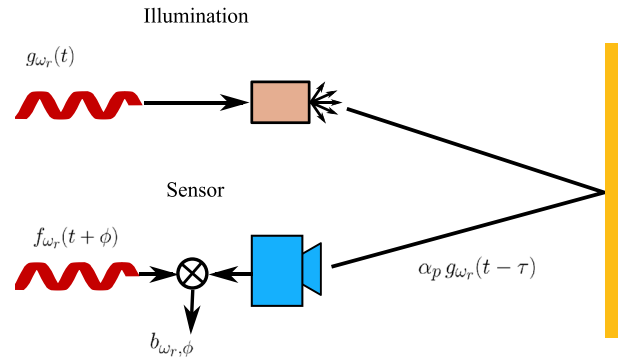
$$\mathbf{s}(t) = \alpha_p \cos(f_{\omega_R} t + \phi) + \beta, \quad (5)$$

$$\mathbf{p}(t) = \cos(g_{\omega_T} t), \quad (6)$$

where  $\phi$  is the phase shift at the sensor, and  $\beta$  the ambient illumination. Capturing a set of different phase shifts  $\phi$  allows to retrieve phase differences between the emitted and the received signals. These per-pixel phase differences correspond to light travel time, thus encoding distance (depth), and other possible sources of delay.

Early works demonstrated the applicability and performance limitations of this principle for range imaging in robotic environments (Hebert and Krotkov, 1992; Adams and Probert, 1996).

<sup>1</sup> Note that we use the term *AMCW* when referring to these specific sensors, whereas we use *ToF* for general phase-based time-of-flight sensors.



**Fig. 4.** Basic operation principle of a time-of-flight emitter-sensor setup. Light is amplitude-modulated at the source, constantly emitted towards the scene, and each pixel modulates the impulse response of the observed scene point. By performing cross correlation (integration) of both modulated emitted and received signals, phase differences can be estimated to reconstruct light travel time.

Source: Image from Heide et al., 2014b.

Due to hardware characteristics, these approaches were limited to a single range detection per shot, requiring systematic and time-consuming scanning of the scene to obtain a full depth map. The first prototype that allowed simultaneous scene capture with modulated array sensors was introduced by Schwarte et al. (1997), coined under the denomination of *photonic mixer device (PMD)*. Lange et al. (2000) and Lange and Seitz (2001) independently introduced a new type of ToF devices based on demodulation “lock-in” pixels, operating on CCD technology with modulation frequencies of a few tens of MHz, and allowing real-time range measurements. These technologies opened new avenues of research on applications and challenges imposed by hardware characteristics.

An important operational aspect of ToF setups resides in how the emitter and sensor frequencies are paired. Homodyne configurations use the same frequency at both emitter and sensor ( $f_{\omega_R} = g_{\omega_T}$ ), while heterodyne ones use slightly different frequency pairs. While being more complicated computationally, heterodyne setups have been demonstrated to provide better ranging precision (Conroy et al., 2009), allowing up-to sub-millimeter resolution (Dorrington et al., 2007). Additionally, proper calibration of ToF cameras was demonstrated to play a significant role when mitigating systematic errors on range estimation (Fuchs and Hirzinger, 2008; Lindner et al., 2010).

Beyond traditional range imaging, Heide et al. (2013) demonstrated that by correlating a set of sensor measurements with different modulation frequencies and phase shifts, a discrete set of per-pixel light travel times and intensities could be reconstructed through optimization, leading to an inferred transient image of the scene. However, the number of frequencies and phases required for this reconstruction is significantly higher than the default set provided by ToF devices (a few default frequencies and phases vs. hundreds of them). They work around this issue by substituting the built-in light source, signal generator and phase triggering by external elements. This ToF-based setup is much cheaper than femto-photography (Velten et al., 2013); however, it only reaches nanosecond resolution (compared to picoseconds in femto-photography), the signal is reconstructed as opposed to directly captured, and tweaking the off-the-shelf devices requires a significant amount of skilled work.<sup>2</sup>

Successive works aimed to overcome different ToF devices limitations that affect the viability of subsequent reconstruction methods. Kadambi et al. (2013) reconfigured the emitter modulation with custom-coded illumination, which improved conditioning on

<sup>2</sup> <http://www.pulsr.info/>.

the optimization by supporting sparsity constraints. This allowed them to recover per-pixel transient responses using a single frequency, instead of hundreds. Recent work by Peters et al. (2015) introduced a way to generate robust sinusoidal light signals, which allowed them to obtain up to 18.6 transient responses per second using a closed-form reconstruction method.

ToF sensor noise, together with limited emitted light intensity due to safety and energy issues, make sensor exposure time and lens aperture the two main factors to achieve an acceptable SNR. To support real-time applications, exposure times must be kept short, so the aperture is usually large to capture as much available light as possible. This introduces a shallow depth of field that blurs scenarios with significant depth changes. Additionally, the low resolution of these sensors (e.g.  $200 \times 200$  for PMDs) affects the spatial precision of the captured data. Godbaz et al. (2010) provided a solution to the shallow depth of field by using coded apertures and explicit range data available in the ToF camera in order to perform defocus, effectively extending the depth of field. Xiao et al. (2015) leveraged the amplitude and range information provided by the ToF devices to recover the defocus blur kernel and regularized the optimization in those amplitude and range spaces, allowing for defocus and increased resolution.

Regardless of wide apertures, exposure times need to be much longer than a single modulation period  $T_{\omega_R} = 1/f_{\omega_R}$ , in order to mitigate sensor noise. This causes a pathological problem known as *phase wrapping*. Since light travel time is encoded in the phase shift between emitted and received light, the modulation period  $T_{\omega_R}$  determines the maximum light path length  $c T_{\omega_R}$  that can be disambiguated, with  $c$  the speed of light. Any light path starting at the emitter that takes longer than this distance to reach a pixel in the sensor will *phase-wrap*  $T_{\omega_R}$ , falling into the same phase shift than shorter paths within subsequent modulation periods. These phase-wrapped light paths produce interference in the measured data, leading to errors in the reconstruction. A straightforward way to solve this is to lower the modulation frequency, thus increasing the maximum unambiguous path length. However, this decreases the accuracy obtained for the reconstructed path lengths, leading to less precise depth measurements. Jongenelen et al. (2010) demonstrated how to extend unambiguous maximum range while mitigating precision degradation, by exploring different dual combinations of simultaneous high and low modulation frequencies. Recently, the work by Gupta et al. (2015b) generalized the use of multiple high frequencies sequentially for this purpose in what they denominate micro-ToF imaging. Phase-wrapping is closely related to the widely-studied problem of MPI, where light from multiple light paths is integrated in the sensor resulting in signal interference and thus reconstruction errors. However, this is related to how some physical phenomena (e.g. interreflections, scattering) affect certain applications – actually affecting other capture methods too – rather than to operational limitations of the ToF devices themselves. Please refer to Section 4.1.2 for a more detailed discussion.

Recent works explore novel hardware modifications: Tadano et al. (2015) increased temporal resolution beyond the limit of current ToF devices (around 100 picoseconds), by using arrays of LED emitters spatially separated by 3 mm. This effectively corresponds to time shifts of 10 picoseconds. Shrestha et al. (2016) explored imaging applications synchronizing up-to three multi-view ToF cameras. To achieve this, they addressed interference problems between the light sources of the cameras, showing how they can be mitigated by using different sinusoidal frequencies for each sensor/light pair. The authors demonstrated applications such as improved range imaging for dynamic scenes by measuring phase images in parallel with two cameras, doubling single-camera frame rate, and mitigating motion artifacts.

## 2.4. Discussion

In general, transient imaging systems present a trade-off between cost, ease of use, acquisition time, and quality (spatial and temporal resolution), as summarized in Table 1. The impulse-based techniques described in Section 2.1, while equivalent in terms of the main imaging principle, exhibit several differences regarding the mentioned trade-off. Avalanche photodetectors allow for cheap and fast imaging of single-scattered photons, suitable for LIDAR applications; the time resolution achieved is in the order of tens of picoseconds. Gated systems, on the other hand, are more costly in terms of measurements, since they image each frame independently, although this cost can be alleviated when focusing on range sensing. In addition, the temporal resolution of these systems ranges from hundreds of picoseconds to milliseconds, allowing for a large variety of application domains. The systems based on streak cameras offer the highest temporal resolution, in the order of hundreds of femtoseconds, at the expense of longer exposure times to achieve a workable SNR. Last, systems based on interferometry (Section 2.2) yield the highest temporal resolution, but present many shortcomings including motion sensibility, extremely shallow depth of field, and inherent complexity, which make them hard to use in general scenarios.

On the other hand, phase-based time-of-flight systems (Section 2.3) provide a more affordable hardware alternative, with shorter acquisition times. However, these devices (e.g. PMD, ToF Kinect) cannot be used off-the-shelf, requiring many modifications in terms of hardware components, electronics, signal generation, or even multi-device configurations, to achieve acceptable results. Despite considerable hacking, they yield a much lower temporal resolution – usually in the range of nanoseconds – compared to the pico or femtoseconds obtained with impulse-based and interferometry-based approaches, respectively. Moreover, these systems do not directly acquire temporal information; instead they require post-capture computations to reconstruct the signal, which may incur in errors.

## 3. Analysis of transient light transport

The data captured by transient imaging devices is affected by the MPI problem: light from different paths arrives at the same pixel in the sensor (see Section 4.1.2). This makes analysis difficult, and limits the range of applications for transient imaging.


To overcome this, different approaches have been proposed, which can be categorized in two main directions: the first one seeks to exploit the sparsity (or compressibility) of the light transport data in the temporal domain; the second analyzes transient light transport in the Fourier domain. In addition, other works exploit the separability of light transport by spatio-temporal coding during capture.

### 3.1. Sparsity and compressibility

A common approach for transient light transport analysis leverages the sparsity of light transport, either in the temporal domain or its derivatives, or by modeling it as a *compressive signal* in some alternative basis, according to the particular scene and the targeted application. Depending on these forms of sparsity, a number of phenomenological models have been developed, from explicitly modeling physically-based priors (e.g.  $K$ -sparse models in the primal domain), to other models based on observation of the captured phenomena.

#### *K*-sparse responses model.

One of the most common approaches involves considering the time-profile as a mixture of  $K$  Dirac  $\delta$ -functions in the temporal



**Fig. 5.** Decomposition on different components of transient light transport. *Left:* A Gaussian profile closely approximates the convolution of the sensor’s temporal PSF and the direct illumination impulse response. *Middle:* For a point where subsurface scattering is dominant, the time profile decays exponentially over time. *Right:* A time profile showing both direct and subsurface light transport, which results in the sum of a Gaussian impulse (direct illumination) and an exponential decay (SSS). Source: Image from Wu et al. (2014).

domain, where each light interaction with surfaces is represented as an impulse response:

$$\mathbf{t}(t) = \sum_{i=1}^K \alpha_i \delta_{t_i}(t), \quad (7)$$

where  $\mathbf{t}(t)$  is the time-resolved transport response at a single pixel,  $\alpha_i$  is the intensity reaching the sensor, and  $\delta_{t_i}$  is the impulse response at instant  $t_i$ .

This model forms the basis of gated-laser (Section 2.1) and correlation-based (Section 2.3) range methods, which assume a single delta response ( $K = 1$ ) on the first scattering event with the nearest surface. This single-bounce assumption is however far from robust, since it assumes scattering on opaque surfaces with no indirect illumination, and thus suffers from the MPI problem. To partially alleviate this, other approaches include additional sparse responses for other interreflections, including an arbitrary large number of responses ( $K > 1$ ) (Kadambi et al., 2013; Kirmani et al., 2013; Bhandari et al., 2014a; Qiao et al., 2015; Kadambi et al., 2015; Peters et al., 2015). For the particular case of  $K = 2$ , this approach allows for a very fast reconstruction (Dorrington et al., 2011; Godbaz et al., 2012; Adam et al., 2016; Naik et al., 2015).

Wu et al. (2012a, 2014) noted that in streak images this  $K$ -sparse model is actually convolved by the sensor temporal point spread function (PSF) (Fig. 5, left), which presents a Gaussian shape (Velten et al., 2013). Following this observation, the authors model equation (7) as a sum of Gaussians instead:

$$\mathbf{i}(t) = \sum_{i=1}^K \alpha_i G_{t_i}(t, \sigma) \quad (8)$$

where  $G_{t_i}$  is the Gaussian centered at  $t_i$ , with standard deviation  $\sigma$  dependent on the imaging system. Note that Eq. (8) does not model the transport operator  $\mathbf{t}(t)$ , but the imaged pixel at the transient image  $\mathbf{i}(t)$ . Using Gaussians instead of delta functions breaks the (primal) sparsity assumption; however, as noted by the authors, the signal is still sparse in the gradient domain for direct transport, as well as specular reflections and transmissions. This allowed them to separate direct and global transport components (interreflections plus subsurface scattering). Further work used this model to reconstruct the full transient image based on correlation-based sensors (Heide et al., 2013), while Hu et al. (2014) improved Wu’s technique by using a more robust method based on convolutional sparse coding.

The models discussed so far assume impulse light transport, where all scattering events occur between perfectly opaque or transmissive surfaces and light travels freely through empty space. Since indirect illumination is a continuous signal, it would be impractical to model it using the  $K$ -impulse model (or its Gaussians-based version). Two other models, discussed below, address subsurface (volumetric) scattering, as well as diffuse interreflections.

#### Exponential volumetric models.

In the presence of translucent objects or participating media, indirect illumination plays an important role. For the case of

translucent, optically thick objects, Wu et al. (2012a, 2014) empirically observed that a single exponential decay accurately models subsurface scattering as:

$$\mathbf{v}(t) = \exp(\log(L_i) + \gamma t), \quad (9)$$

where  $L_i$  is the incident illumination and  $\gamma$  is the scattering coefficient of the object (Fig. 5, middle and right).

For light transport in more general participating media (i.e. relaxing the assumption of optically thick media), Heide et al. (2014b) used a similar exponential decay assumption; however, since most sensors have a Gaussian PSF, they approximated this decay by using an exponentially modified Gaussian. This allowed them to model transient light transport in participating media as a mixture model of these types of Gaussians, which provide a compressive base for transient light transport in such media. Interestingly, a similar set of time-resolved exponentially modified Gaussians, integrated in the temporal domain, lies at the core of the quantized diffusion model (D’Eon and Irving, 2011) for rendering high-quality subsurface scattering.

#### Exponential-based diffuse illumination.

Freedman et al. (2014) observed that the Lambertian indirect temporal response  $\mathbf{I}(t)$  also presents a smooth exponential shape, although different from Wu’s. They modeled it as:

$$\mathbf{I}(t) = A t^\alpha e^{\beta t}, \quad (10)$$

where  $A$ ,  $\alpha$  and  $\beta$  depend on the geometry and reflectance of the underlying scene, and  $t \geq 0$ . For  $t < 0$ , we have  $\mathbf{I}(t) = 0$ .

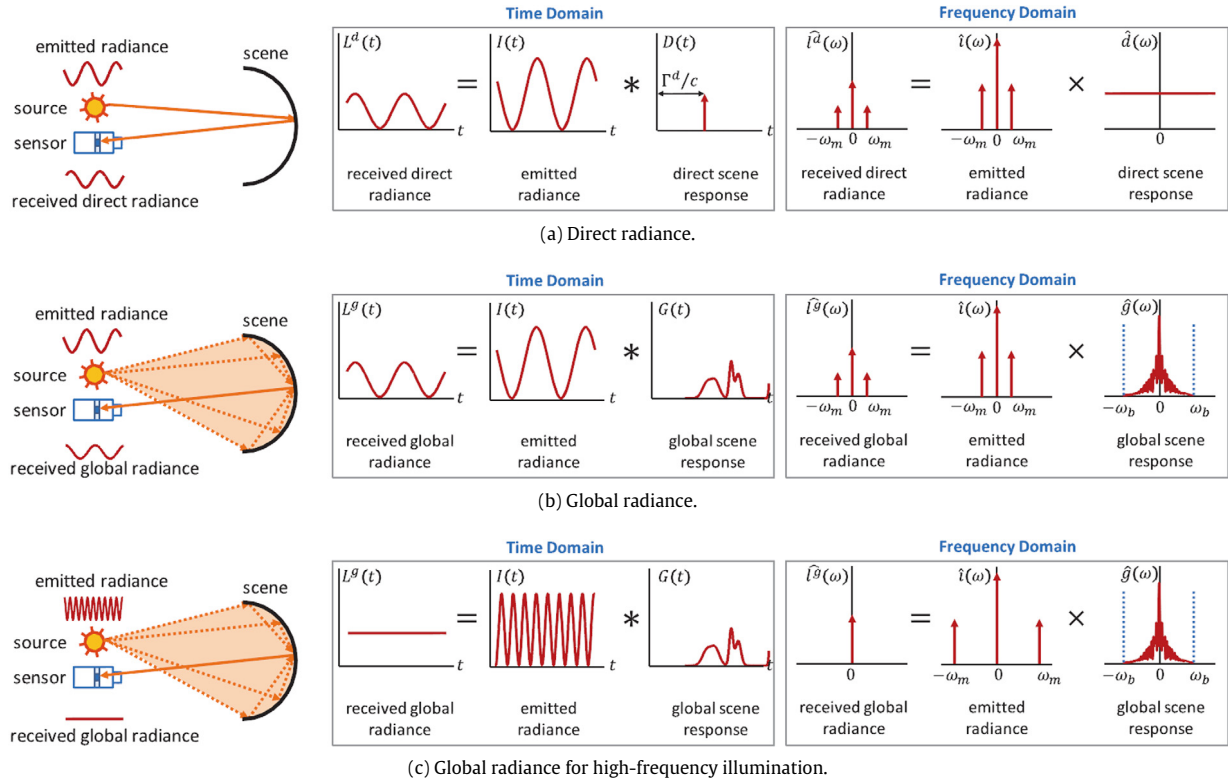
With this approach, the authors modeled transient light transport as a sum of the impulse-like transport ( $K$ -sparse model) for direct and specular paths, plus the Lambertian indirect illumination as:

$$\mathbf{t}(t) = \sum_{i=1}^K \alpha_i \delta_{t_i}(t, \sigma) + \sum_{j=1}^{K_l} \mathbf{I}_j(t - t_j), \quad (11)$$

with  $t_j$  the shortest light path of the diffuser. Unfortunately, in this formulation transient light transport is no longer sparse neither in the primal nor in the gradient domains. However, the impulse functions and the exponential decays form again a compressive base.

### 3.2. Frequency domain

Fourier-based analyses and techniques have been developed to gain fundamental insights about the information encoded in the temporal domain, and as a tool for effective capture and processing of transient data. Wu et al. (2012b) extended the frequency analysis on the incoming steady-state light field (Durand et al., 2005) by analyzing the full 5D time-resolved light field. Their analysis revealed a cross-dimensional information transfer between domains. They further demonstrated potential applications of this



**Fig. 6.** Frequency analysis of light transport using AMCW systems, modeled using a phasor representation of the sinusoidal signal. In the primal domain, the response is obtained by convolving the emitted radiance with the response of the scene, which is a delta response for direct transport (a), and might get arbitrarily complex in the presence of multiple interreflections (b, c). This response is in general smooth, and therefore bandlimited in the frequency domain. This means that for high-frequency modulated emitted radiance, the global response contains only the DC component. Source: Image adapted from Gupta et al. (2015b).

analysis by prototyping a transient-based bare sensor imaging system. Lin et al. (2014, 2016) observed that capturing transient light transport using AMCW sensors with a homodyne setup is equivalent to sampling transient light transport in the frequency domain. Based on this observation, they recovered the temporal radiance profile by means of the inverse Fourier transform of the captured data. Kadambi et al. (2015) showed that the formulation used in AMCW (based on the correlation of sinusoid waves, encoded as phasors), and frequency-domain Optical Coherence Tomography (OCT) (Huang et al., 1991) are analogous. This allowed them to use the well-studied methods in OCT to capture transient light transport. One of the advantages of these techniques is that they do not require sampling in phase; as the authors demonstrate, this allows to use any sensor (including traditional CMOS-based cameras), and might eventually lead to sub-picosecond resolution captures with standard, low-cost time-of-flight sensors.

Finally, Gupta et al. (2015b) used frequency analysis on top of their phasor-based model of light transport. The authors demonstrated that for high-frequency correlation-based sensing, diffuse interreflections vanish due to the band-limited nature of such interreflections. This allowed them to avoid the multipath interference problem by capturing direct light transport with high-frequency sparse impulses, while lower modulation frequencies were used to capture indirect transport (Fig. 6).


### 3.3. Spatially-coded capture

A third approach involves spatial light coding, coupled with temporal modulation during capture. These strategies allow to analyze (decompose) light transport in scenes, without explicitly knowing its temporal profile. Naik et al. (2015) exploited the

relationship between direct and indirect illumination when using high-frequency spatial illumination patterns. They leveraged the work of Nayar et al. (2006) to decouple both components, significantly improving range imaging. This work however needs an external projector to perform the separation. To overcome this problem, Whyte et al. (2015) derived a theoretical framework to perform spatially-modulated separation using phases and amplitudes, allowing their use in AMCW systems. They additionally obtained an optimal set of projection spatial patterns. Also based on spatial modulation, O'Toole et al. (2014) used optical probing of light transport (O'Toole et al., 2012), which allows not only direct-indirect separation, but also decomposing high- and low-frequencies of indirect transport.

### 3.4. Discussion

Analyzing transient light transport is a key step towards producing practical applications. Exploiting sparsity in the primal domain has been a common, flexible approach. However, although the  $K$ -sparse models have physical meaning, they become impractical for diffuse interreflections, including volumetric and subsurface scattering. To overcome this, phenomenological models based on empirical observations have been developed, and shown to provide a good base for compressive decomposition. On the other hand, moving from the primal to the Fourier domain helps extend the range of potential analyses, allowing more principled decompositions, or using tools from well-established fields (e.g. inverse Fourier transform, or frequency-based OCT). This helps reduce the need for heuristic bases to represent light transport. Interestingly, some other works follow a different approach, making use of spatial light modulation together with temporal modulation, leading to very robust light transport decomposition at the cost of additional effort and machinery during capture.



**Fig. 7.** (a) Range imaging using a ToF sensor, where the phase delay of the emitted modulated radiance encodes light time of flight. This formulation works for direct reflection, but presents problems when light from multiple paths is integrated in the sensor, as happens in the presence of specular reflections (b) or transmission (c). This problem gets even more complicated in the presence of diffuse transport, due to scattering in media (d) or Lambertian diffuse interreflections (e). Source: Image after Bhandari et al. (2014b).

## 4. Applications

The temporal resolution offered by transient data has opened up a wealth of possibilities in terms of applications. Long-standing vision and graphics problems, such as depth recovery in the presence of interreflections, or light component separation, have received renewed attention in light of this new data (Raskar and Davis, 2008). Here we delve into the areas which have benefited most from transient imaging in the realm of graphics and vision, but applications extend to fields such as medical imaging, surveillance, or atmospheric sciences, to name a few.

### 4.1. Geometry reconstruction

The most prevalent application of transient imaging is the recovery of depth information from the scene. More recently, ultrafast transient data has opened the door to recovering not only depth but full geometry information, i.e., including non-line-of-sight areas (due to occlusions or the presence of a dense medium). This requires being able to drop the assumption that light only bounces once before reaching the camera; effectively, this means having the ability to separate the paths followed by the different light rays, that is, solving the multipath interference problem. While a large body of work on traditional range imaging via ToF exists (Hansard et al., 2012; Remondino and Stoppa, 2013), we place the focus here on the recent approaches exploiting ultrafast data. The reader may also refer to specialized tutorials on ToF imaging (Kolb et al., 2010; Gupta et al., 2015a).

#### 4.1.1. Range imaging

Range imaging, that is, obtaining depth information from a scene, has been traditionally achieved via two different methods: those based on multiple viewpoints of the scene<sup>3</sup> and obtaining correspondences between them, and those based on time of flight. The principle behind ToF sensors has been explained above (please refer to Section 2.3 for a detailed description), and the computation of depth information (i.e. distance to the camera,  $d$ ) from that temporal data is as conceptually as simple as applying  $d = vt$ , where  $t$  is the time light has taken to travel the distance  $d$ , given its velocity in the medium  $v = c\eta$  (where  $c$  is the speed of light in a vacuum). More specifically, in the commonly used AMCW cameras, what is measured is the phase shift  $\theta$  in the modulation envelope (as well as the amplitude), from which the distance can be obtained as (Dorrington et al., 2011; Kadambi et al., 2015):

$$d = \frac{v\theta}{4\pi f_{\omega_R}}, \quad (12)$$

where  $f_{\omega_R}$  is the frequency of the modulation signal.

Modern Kinect sensors, for instance, employ this technology to obtain depth information (Bamji et al., 2015). ToF cameras allow obtaining relatively low resolution depth images of the scene. Still,

<sup>3</sup> Structured light approaches, relying on one camera and one projector, can be considered a subset of these (Gupta et al., 2012).

they offer a great advantage over previous scan-based (LIDAR) techniques – which measured the time of flight of a laser beam to a point – the capture speed is in the order of hundreds of frames per second. The main stumbling block of traditional ToF techniques for range imaging is the underlying assumption that the light only bounces once before reaching the camera. While this may be true in some scenes, in many other cases, and particularly in the presence of concavities (such as the corners of a room), purely specular objects, and scattering media (fog, tissue, etc.), this no longer holds and conventional depth recovery methods fail (see Fig. 7). We cover techniques developed to address this problem in Section 4.1.2.

Other means of obtaining transient data, such as direct imagers, have also been applied to the recovery of depth information. Again, if the assumption for each point is that light arriving to the camera has bounced only once, reconstructing depth is trivial, but this is often not the case. Thus, the focus of existing works in the area of direct imagers has been to reconstruct occluded geometry. We cover these works in Section 4.1.3.

#### 4.1.2. The multipath interference (MPI) problem

The MPI problem is common for most transient imaging devices, specially in those with long exposure times: for example, in the context of gated-based LIDAR systems (Section 2.1), where a modulated sensor response (Laurenzis et al., 2007; Laurenzis and Woiselle, 2014) is used to robustly acquire depth. However, is in ToF cameras where the problem is more noticeable. Some early approaches to solving the MPI problem in ToF cameras targeted in-camera light scattering (Kavli et al., 2008; Schäfer et al., 2014); others targeted also indirect illumination but require placing tags in the scene (Falie, 2009), or made severe assumptions on scene characteristics (Jamtsho and Lichti, 2010). For an in-depth discussion about the MPI problem from a signal processing perspective, we refer the reader to a recent article by Bhandari and Raskar (2016). A comparative summary of the techniques discussed in this section can be found in Table 2.

The work of Fuchs (2010) provided a model of MPI for the case in which all distracting surfaces are Lambertian, based on explicitly computing indirect illumination on the estimated depth map and iteratively correcting it. Follow-up works aimed at a more general solution targeting the source of the problem: the separation of the individual components when multiple returns are present (Godbaz et al., 2008, 2009), also called Mixed Pixel Restoration. These techniques, however, cannot be used with off-the-shelf cameras, since they require measuring multiple phase steps per range measurement (as opposed to the usual four). Of large relevance is the work of Dorrington et al. (2011), in which the authors proposed a numerical solution that can be employed in off-the-shelf ToF cameras. Shortly after, Godbaz et al. (2012) proposed two closed-form solutions to the problem. These two works assume, however, that there are two return components per pixel, and work with two or up to four modulation frequencies. This two-component, dual-frequency approach was generalized by Bhandari et al. (2014b). Kirmani et al. (2013) targeted simultaneously phase unwrapping




**Table 2**

Comparison between the different existing techniques addressing the problem of MPI.

Source: Adapted and updated from Naik et al. (2015).

Work	Multipath type	Solution type	Hardware modifications
Fuchs (2010)	Continuous	Iterative	None
Dorrington et al. (2011)	2-sparse	Iterative	Frequency sweep
Godbaz et al. (2012)	2-sparse	Closed-form	Frequency sweep
Kadambi et al. (2013)	$K$ -sparse	Iterative	Custom code
Kirmani et al. (2013)	$K$ -sparse	Iterative	Frequency sweep
Heide et al. (2013)	$K$ -sparse	Sparse regularization	Frequency sweep
Freedman et al. (2014)	$K$ -sparse	Iterative	None
Jiménez et al. (2014)	$K$ -sparse	Iterative	None
O'Toole et al. (2014)	Continuous	None	Extensive
Lin et al. (2014)	Continuous	Closed-form	Frequency sweep
Gupta et al. (2015b)	Continuous	Closed-form	Extensive
Naik et al. (2015)	Continuous	Closed-form	External projector
Peters et al. (2015)	$K$ -sparse	Closed-form	Frequency sweep
Qiao et al. (2015)	$K$ -sparse	Sparse regularization	Frequency sweep
Kadambi et al. (2015)	Continuous	Closed-form	Frequency sweep
Whyte et al. (2015)	Continuous	Closed-form	Custom code



**Fig. 8.** Given a scene with transparent objects (a), regular time of flight cameras fail at reconstructing their depth (b). Kadambi and colleagues' method, assuming that light transport is modeled as a sparse set of Dirac deltas, can correctly recover the depth of the unicorn (c and d).

Source: Figure from Kadambi et al. (2013).

and multipath interference cancellation, using a higher number of frequencies (five or more), but at a lower computational cost than previous approaches, thanks to a closed form solution. Still, they assumed sparsity in the recovered signal, and again restricted their model to two-bounce situations ( $K = 2$ , see Section 3.1).

The use of multiple modulation frequencies was also leveraged by Heide et al. (2013). In their case, they used hundreds of modulation frequencies, and proposed a model that includes global illumination. Freedman et al. (2014) also required multiple frequencies, and proposed a model (not limited to two bounces) which assumes compressibility of the time profile; they solved the problem iteratively via  $\mathcal{L}_1$  optimization. Kadambi et al. (2013) reduced the number of frequencies required to recover a time profile (and thus depth information) to one, by using custom codes in the emission in combination with sparse deconvolution techniques, to recover the time profiles as a sparse set of Dirac deltas. This technique allowed to recover depth in the presence of interreflections, including transparent objects (Fig. 8).

All these works assumed a  $K$ -sparse transport model (Section 3.1). It is worth noting, however, that in the case of scattering media being present, a sparse formulation of the time profile is no longer possible. The problem of scattering media is treated in Section 4.1.4. A slightly different approach was taken by Jiménez et al. (2014), who proposed an optimization framework to minimize the difference between the measured depth, and the depth obtained by their radiometric model. Convergence to a global minimum was not guaranteed, but a number of examples including real scenes were shown. Hardware modifications are not required.

A different means of eliminating or separating global light transport in a scene was presented by O'Toole et al. (2014), who made the key observation that transient light transport is separable in the temporal frequency domain (see Section 3.2). This allowed them to acquire and process only the direct time-of-flight component, by using a projector with light modulated in space and time (note that they do not use correlation-based ToF). Gupta et al. (2015b) built on this idea, and proposed a framework termed *phasor imaging*. A key observation is that global effects vanish for

frequencies higher than a certain, scene-dependent, threshold; this allowed the authors to recover depth in the presence of MPI, as well as to perform direct/global separation, using correlation-based time-of-flight sensors. Neither Gupta et al.'s work, nor O'Toole et al.'s, imposed the restriction of sparsity of the multipath profile. Neither did Naik et al. (2015), who also attempted direct/global separation to obtain correct depth in the presence of MPI. A similar approach was followed by Whyte et al. (2015).

#### 4.1.3. Reconstruction of non-line-of-sight geometry

A recent analysis of the problem of non-line-of-sight (NLOS) geometry recovery and its feasibility with ToF-based devices can be found in the work of Kadambi et al. (2016b). Although this problem has been tackled in other imaging modalities, such as radar, a seminal paper in our area was the work of Velten et al. (2012a), who provided a method to reconstruct NLOS in a controlled laboratory setup (similar to femto-photography) but without the need for scene priors, by using a backpropagation technique, and leveraging the extremely high temporal resolution of the time profiles recovered. This work was later extended by Gupta et al. (2012), with the aim of improving the analysis and robustness of the reconstruction method.

The idea that NLOS reconstruction was possible with a somehow similar setup was raised by Kirmani et al. (2009, 2011a). Later, Laurenzis and Velten (2014), and Buttafava et al. (2015) generalized Velten's backpropagation NLOS reconstruction using data captured with gated systems and SPAD respectively. Using non-impulse illumination for NLOS was later investigated by Kirmani and colleagues (2011b, 2012), which aimed to capture spatially varying reflectance from non-line-of-sight planes. Heide et al. (2014a) presented a technique with less expensive hardware based on AMCW systems (see Section 2.3), although requiring significant modifications to off-the-shelf hardware. Their technique works in the presence of ambient illumination, and relies on scene priors such as sparsity of the geometry to regularize the non-linear optimization problem that arises when the linear image formation model (after (Heide et al., 2013)) is inverted. Hullin

(2014) described an analysis-by-synthesis approach: forward light transport is implemented using radiosity, and posed the problem as an optimization that deforms the geometry (starting from a tessellated spherical blob), until the error between the captured radiance and the output of the forward model is minimized. As the author notes, convexity, and thus convergence of the optimization problem to the global solution, is not guaranteed. A generative analysis-by-synthesis approach was also used by Klein et al. (2016), to detect and track NLOS objects under a number of simplifying assumptions, employing only a laser pointer and a conventional 2D camera.

On a theoretical level, recently, Tsai et al. (2016) developed a framework for shape recovery for the particular case of two-bounce light paths. The problem of separating these paths specifically remains unanswered, but the work provides the foundation for future attempts.

#### 4.1.4. Imaging through scattering media

There are many works regarding imaging through scattering media, traditionally focusing on medical or deep-tissue imaging (Han et al., 2000), or inspection of art (Abraham et al., 2010), to name a few examples. Transient data has also been recently used for this goal. Heide et al. (2014b) introduced a convolutional sparse coding approach using correlation image sensors. In particular, the authors used a modified ToF camera, combined with a physically-based transient image formation model (see Section 3) that improves sparsity in scattering media. Using instead a femto-photography setup (Velten et al., 2013), Raviv et al. (2014) obtained the six degrees of freedom (3D position and orientation) of rigid known geometric shapes, by leveraging scattering information. The authors made the observation that obtaining the full image is not always convenient or necessary; with less information to recover, their single-shot system can be used in dynamic scenes, allowing to track the object being imaged. Femto-photography was also employed by Naik et al. (2014) to recover the spatially-varying reflectance of a scene seen through a scattering medium by means of a numerical inversion algorithm. The method requires the time spent in the scattering medium to be lower than the temporal resolution of the camera. Last, Satat and co-workers (2015b) combined time-resolved information with an optimization framework to image a scene behind a relatively thick scattering medium of 1.5 cm.

Highly related is the problem of recovering the shape (depth) of transparent objects and their backgrounds. An example was shown by Kadambi et al. (2013) (see Fig. 8), in the context of their sparse deconvolution approach to the multipath interference problem. Lee and Shim presented a two-step approach, using a skewed ToF stereo camera (Lee and Shim, 2015). First the transparent object is detected by analyzing inconsistencies between the views from the two cameras; depth is then recovered by means of optimization, minimizing the distance between inconsistent points along their respective light rays. The same authors later proposed a second technique, this time using a single ToF camera (Shim and Lee, 2016). Although lighter on the hardware side, the method requires two different measurements, with and without the transparent object. By directly analyzing the distortions created by a transparent object in ToF profiles, Tanaka et al. (2016) showed that the refractive light path (from which depth can be inferred) can be uniquely determined with a single parameter. This is estimated with the help of a known reference board, moved to two different locations behind the transparent object.

#### 4.2. Motion estimation

Transient data also allows the detection of non-line-of-sight moving objects. First tackled by Pandharkar et al. (2011), the

proposed approach utilized a femtosecond laser as illumination source, and data is obtained from a streak camera at picosecond resolution, with equipment similar to the one shown in Fig. 2, but a different setup. The tracked object is located at each frame by backprojecting the recorded signals, and solving a constrained least squares problem. From these locations, the motion vector is obtained.

A different approach relied on the well-known Doppler effect to obtain the velocity field of a given scene, using a ToF camera (Heide et al., 2015). The shift in frequency of the illumination that occurs when a large radial velocity is present breaks down the conventional ToF formulation in homodyne setups. Drawing inspiration from the communications field, the authors proposed the use of orthogonal illumination and modulation frequencies, requiring two simultaneous measurements (one homodyne, one heterodyne) to recover the velocity field. An additional, phase-offsetted homodyne measurement yields a range image as well.

#### 4.3. Material estimation


Reflectance acquisition is traditionally a time-consuming process, due to the large number of measurements needed to fully capture a generic BRDF. Transient data, with its potential ability to disentangle light paths, has enabled alternative means of reflectance acquisition, including “around the corner” setups. Early work in this regard was carried out by (Pandharkar, 2011), later improved by Naik et al. (2011), based on a three-bounce setup with planar walls. Light from a laser bounces off a diffuse surface, then is reflected off a planar sample of the material, before reaching a third surface, also diffuses, which is observed by a streak camera. The laser needs to be swept over a number of positions on the first surface, but different outgoing directions are captured in a single shot (Fig. 9). Different light paths with the same path length pose a challenge; the authors formulated the reflectance recovery problem as an undetermined linear system, modeling the reflectance with a low-dimensional parametric BRDF model. Subsequent work using the same hardware was able to realize reflectance acquisition for objects behind a diffuser, which has applications in biological and medical imaging (Naik et al., 2014). This requires estimating not only reflectance but also the scattering profile of the diffuser (assumed Gaussian). The authors formulated the forward model and then solved a convex optimization problem minimizing the difference between estimated and measured intensity. Tsai et al. (2016) described reflectance estimation by taking advantage of the coverage of different incident and outgoing ray directions given by two-bounce paths for certain scene geometries. Different from other works, this approach did not require all surfaces to be Lambertian.

##### 4.3.1. Fluorescent lifetime imaging

The technique known as fluorescence lifetime imaging allows to obtain reflectance properties of surfaces through turbid media. It has applications in diagnosis or inspection, and is one of the most classical applications of ultrafast time-resolved imaging using impulse based illumination, with exposure times in the order of a few hundred picoseconds (see e.g. Wu et al., 1995). Using a femtosecond laser and a streak camera, Satat et al. (2015a) utilized time-resolved images to localize and classify the lifetimes of fluorescent probes, which need not be in the line of sight of the camera. Another approach to the problem, which required simpler hardware but is only able to reconstruct directly visible fluorescent samples, is the work by Bhandari et al. (2015). A ToF Kinect sensor was used to obtain lifetime information from samples, without calibration nor known illumination.

##### 4.3.2. Material classification

A related but simpler problem is material classification. Wu et al. (2014) decomposed global light transport into its direct,



**Fig. 9.** (a) Setup for reflectance acquisition using transient imaging: points in the left wall are successively illuminated, creating a range of incoming directions on the back wall; light in turn bounces off it and reflects on the right wall before reaching the camera. (b) Streak image captured with two patches on the back wall. (c) Illustration of the setup: the bright region in the streak image in (b) (red part of the rightmost graph) contains information about the shape of the specular lobe at P. Source: Figure from Naik et al. (2011).

subsurface scattering, and interreflection components, by analyzing the time profile at picosecond resolution. They noted that the time profile of a point lit by subsurface scattering decays exponentially, which allowed them to identify translucent materials in a scene. Further exploring a similar idea, Su et al. (2016) later analyzed four materials with different degrees of translucency, and tested several learning methods on the captured data. The authors achieved varying success rates identifying the materials: Wax was easy to classify, due to its strong subsurface scattering, while progressively less translucent materials, such as paper, styrofoam or towel, became more complex.

## 5. Simulation

Light transport, described using either Maxwell's equations (Born and Wolf, 2002), or the more practical radiative approximation (Chandrasekhar, 1960), is defined in a time-resolved manner. However, since the final goal is usually to compute light transport in steady-state, the practical assumption that the speed of light is infinite becomes a reasonable approximation from a simulation (rendering) perspective. See e.g. Gutierrez et al. (2008) and Křivánek et al. (2013) for an overview on steady-state rendering.

With the establishment of transient imaging in graphics and vision, the simulation of time-resolved light transport is becoming an increasingly important tool. Smith et al. (2008) developed the first framework in the context of the traditional rendering equation (Kajiya, 1986). This was later formalized by Jarabo et al. (2014), extending the path integral (Veach, 1997) to include time-resolved effects such as propagation and scattering delays.

Transient rendering has been used to synthesize videos of light in motion (Jarabo et al., 2014), but is also a key tool to provide ground truth information to develop novel light transport models (O'Toole et al., 2014; Adam et al., 2016), or benchmarking (Nair et al., 2013; Pitts et al., 2014). It can also be used as a forward model for solving inverse problems (Keller et al., 2007; Keller and Kolb, 2009; Fuchs and Hirzinger, 2008; Fuchs, 2010; Jiménez et al., 2012, 2014; Hullin, 2014; Klein et al., 2016).

The key **differences** with respect to steady-state simulation are:

- The speed of light can no longer be assumed to be infinite, so propagation delays need to be taken into account. Note that some works in steady-state rendering also need to account for propagation delays (e.g. rendering based on wave-optics (Moravec, 1981; Musbach et al., 2013), or solving the Eikonal equation for non-linear media (Gutierrez et al., 2005; Ihrke et al., 2007)), although their final goal is to obtain a steady-state image integrated in time.


- Scattering causes an additional delay, due to the electromagnetic and quantum mechanisms involved in the light-matter interaction. These give rise to effects such as fluorescence, or Fresnel phase delays (see Fig. 10).
- The temporal domain must be reconstructed; however, naive reconstruction strategies (i.e. frame-by-frame) are extremely inefficient.
- Motion in the scene (e.g. camera movements) brings about the need to include relativistic effects.

In the following, we discuss the different approaches for effectively reconstructing the temporal radiance profile in simulation; then, orthogonally to reconstruction, we focus on the different algorithms for simulating transient light transport, and on their main target application (see Table 3 for an overview).

### 5.1. Reconstruction of the temporal profile

As discussed by Jarabo et al. (2014), rendering each transient frame independently is highly impractical, given the extremely short exposure times: sampling paths with a given temporal delay is almost impossible, while randomly sampling paths would be extremely inefficient. The most straightforward way to solve this issue and render effectively transient light transport is to reuse the samples for all frames, binning them in the temporal domain (Jarabo, 2012; Marco, 2013; O'Toole et al., 2014; Ament et al., 2014; Adam et al., 2016; Pitts et al., 2014). This is equivalent to a histogram density estimation; although easy to implement, it has a slow convergence of  $O(N^{-\frac{1}{3}})$ , with  $N$  being the number of samples. Jarabo et al. (2014) presented a better alternative, proposing a reconstruction method based on kernel density estimation (Silverman, 1986), which leads to faster convergence ( $O(N^{-\frac{4}{5}})$ ). Interestingly, rendering each frame independently, and using the histogram in the temporal domain, are equivalent to the gate imaging and streak imaging techniques discussed in Section 2.1, respectively.

If the goal is not to generate the full transient profile, but just the modulated response at the sensor as if it were captured by a correlation-based sensor (see Section 2.3), the problem is reduced to generating a single image modulating each sample according to its delay and the sensor response. Thus, while we still need to keep track of the path propagation delays, it can be done within the framework of the traditional path integral, where the sensor response is a function of time. For depth recovery, Keller et al. (2007); Keller and Kolb (2009) proposed a GPU-accelerated rendering system modeling such response. The system is limited to single-bounce scattering, so it assumes no MPI. The sensor response needs accurate sensor modulation models, including temporal behavior and noise. Gupta et al. (2015b) introduced a noise model for AMCW imaging devices, while Lambers et al. (2015) presented



**Fig. 10.** Examples of phenomena observed in transient state: From left to right, wavelength-dependent indices of refraction produce temporal chromatic dispersion; temporal decomposition of ordinary and extraordinary transmission in a birefringent crystal; and energy re-emission after 10 nanoseconds in a fluorescent object.

Source: Image from Jarabo et al., 2014.

**Table 3**

Comparison of selected works on time-resolved rendering, including their output, the type of transport being simulated, and their convergence rate with respect to the number of samples  $N$ . Note that the works based on radiosity (Hullin, 2014; Klein et al., 2016) (marked with an asterisk) have constant convergence, but their cost and accuracy depend on the number of geometric subdivisions in the scene. The work of Meister et al. (2013a) converges with the same rate as traditional steady-state progressive photon mapping.

Work	Output	Algorithm	Global Illumination	Convergence
Keller et al. (2007); Keller and Kolb (2009)	4-Bucket Image	Rasterization	No	$O(1)$
Jarabo (2012)	Transient Image	BDPT	Yes	$O(N^{1/3})$
Meister et al. (2013a)	4-Bucket Image	Photon Mapping	Yes	$O(N^{-2/3})^*$
Marco (2013)	Transient Image	Photon Beams	Media only	—
Jarabo et al. (2014)	Transient Image	BDPT	Yes	$O(N^{-4/5})$
Pitts et al. (2014)	Transient Image	BDPT	Yes	$O(N^{-1/3})$
Hullin (2014); Klein et al. (2016)	Transient Image	Radiosity	2nd bounce	$O(1)^*$
Ament et al. (2014)	Transient Image	Photon Mapping	Yes	$O(N^{-1/3})$
Adam et al. (2016)	Transient Image	BDPT / MTL	Yes	$O(N^{-1/3})$

other physically-based models of the sensor and the illumination, including high-quality noise and energy performance.

## 5.2. Light transport simulation algorithms

Depending on the application domain, existing algorithms to simulate transient light transport trade off accuracy for speed. As a forward model for efficient reconstruction of the geometry of occluded objects, Hullin (2014) and Klein et al. (2016) extended Smith et al.'s (2008) transient version of the radiosity method (Goral et al., 1984) on the GPU. This method is limited to Lambertian surface reflections, and second-bounce interactions.

On the other hand, most works aiming at generating *ground truth* data have used transient versions of Monte Carlo (bidirectional) path tracing (BDPT) (Jarabo, 2012; Jarabo et al., 2014; Pitts et al., 2014; Adam et al., 2016). These are unbiased methods, and support arbitrary scattering functions, including participating media. However, they are in general slow, requiring thousands of samples to converge. To accelerate convergence, Jarabo et al. (2014) introduced three techniques for uniform sampling in the temporal domain targeted to bidirectional methods, while Lima et al. (2011) and Periyasamy and Pramanik (2016) proposed importance sampling strategies in the context of Optical Coherence Tomography. These techniques are designed to work in the presence of participating media; this is a particularly interesting case for transient imaging, since one of its key applications is seeing through such media (fog, murky water, etc.).

Other algorithms aiming to produce ground truth data robustly rely on a photon tracing and gathering approach (Jensen, 2001; Hachisuka et al., 2013). Meister et al. (2013a, b) used a transient version of photon mapping, resulting into a robust estimation of light transport, and allowing to render caustics in the transient domain. Ament et al. (2014) also used transient photon mapping to solve the refractive RTE. However, these techniques are intrinsically biased, due to the density estimation step at the core of the photon mapping algorithm. This bias was reduced by Jarabo et al. (2014), who introduced progressive density estimation along the temporal domain. Targeted to transport in media, Marco (2013)

proposed a transient version of the photon beams algorithm, which was later implemented in 2D by Bitterli (2016).

Last, as mentioned above, camera movements at this temporal resolution bring about the need to simulate relativistic effects in transient light transport. These were simulated by Jarabo et al. (2013, 2015), including time dilation, light aberration, frequency shift, radiance accumulation and distortions on the camera's field of view. The system considered linear motion, as well as acceleration and rotation of the camera.

## 6. Conclusions

Systems capable to obtain scene information from the temporal response of light have existed for a while, mostly focusing on range imaging (e.g. LIDAR systems based on laser impulse or on ToF sensors). However, the numerous problems and limitations of these sensors (low resolution, MPI, etc.) have made them practical only in a limited set of scenarios. Recent advances in transient imaging and ToF technology have, on the other hand, triggered significant improvements on the main, primal application of ToF and LIDAR systems (i.e. range imaging), by providing more robust solutions to the MPI problem. As a result, this has expanded the range of applications and accuracy of this technology, becoming a gold standard for single image depth recovery. Beyond that, transient imaging has opened a vast new field of applications, making traditional and new ill-posed problems in computer vision tractable: Non-line-of-sight vision, or vision through turbid media, are some of the most exciting examples of these new applications. Other problems such as material capture and recognition, motion estimation, dehazing, or bare-sensor imaging have also been tackled successfully using transient imaging. In many cases, the combination of computer graphics, vision, and imaging techniques has played an important role, making transient imaging a truly multidisciplinary field.

Of course, there is still much work to do: While improved hardware has now reached the market (e.g. the Kinect 2 sensor), most other applications are still restricted to controlled laboratory conditions. Limitations in hardware (poor SNR, long capture times),

and software (expensive post-processing, lack of general enough priors) are still too restrictive to move these other applications (e.g. NLOS) into the wild. We hope that by categorizing the current state of the art of the field in a holistic way, including all the pieces we believe crucial for transient imaging (capture, simulation, analysis and applications), will help provide both a clear picture of the main limitations of the existing pipelines, as well as the current development of applications enabled by transient imaging.

## Acknowledgments

We would like to thank Matthias Hullin for discussions and comments on early versions of the work. This research has been partially funded by DARPA (project REVEAL), the European Research Council (Consolidator Grant, project CHAMELEON), and the Spanish Ministerio de Economía y Competitividad (projects TIN2016-78753-P, TIN2016-79710-P and TIN2014-61696-EXP. Julio Marco was additionally funded by a grant from the Gobierno de Aragón.

## References

- Abraham, E., Younus, A., Delagnes, J.C., Mounaix, P., 2010. Non-invasive investigation of art paintings by terahertz imaging. *Appl. Phys. A* 100 (3).
- Abramson, N., 1978. Light-in-flight recording by holography. *Opt. Lett.* 3 (4).
- Abramson, N., 1983. Light-in-flight recording: high-speed holographic motion pictures of ultrafast phenomena. *Appl. Opt.* 22 (2).
- Adam, A., Dann, C., Yair, O., Mazor, S., Nowozin, S., 2016. Bayesian time-of-flight for realtime shape, illumination and albedo. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Adams, M.D., Probert, P.J., 1996. The interpretation of phase and intensity data from amcw light detection sensors for reliable ranging. *Int. J. Robot. Res.* 15 (5).
- Ament, M., Bergmann, C., Weiskopf, D., 2014. Refractive radiative transfer equation. *ACM Trans. Graph.* 33 (2).
- Bamji, C.S., O'Connor, P., Elkhatib, T., Mehta, S., Thompson, B., Prather, L.A., Snow, D., Akkaya, O.C., Daniel, A., Payne, A.D., Perry, T., Fenton, M., Chan, V.H., 2015. A 0.13  $\mu\text{m}$  CMOS System-on-Chip for a  $512 \times 424$  time-of-flight image sensor with multi-frequency photo-demodulation up to 130 MHz and 2 GS/s ADC. *IEEE J. Solid-State Circ.* 50 (1).
- Bhandari, A., Barsi, C., Raskar, R., 2015. Blind and reference-free fluorescence lifetime estimation via consumer time-of-flight sensors. *Optica* 2 (11).
- Bhandari, A., Feigin, M., Izadi, S., Rhemann, C., Schmidt, M., Raskar, R., 2014a. Resolving multipath interference in Kinect: An inverse problem approach. In: *IEEE SENSORS*.
- Bhandari, A., Kadambi, A., Whyte, R., Barsi, C., Feigin, M., Dorrington, A., Raskar, R., 2014b. Resolving multipath interference in time-of-flight imaging via modulation frequency diversity and sparse regularization. *Opt. Lett.* 39 (6).
- Bhandari, A., Raskar, R., 2016. Signal processing for time-of-flight imaging sensors. *IEEE Signal Process. Mag.* 33 (5).
- Bitterli, B., 2016. Virtual Femto Photography. <https://benedikt-bitterli.me/femto.html>.
- Born, M., Wolf, E., 2002. Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light. Cambridge University Press, ISBN: 0 521 64222 1.
- Busck, J., Heiselberg, H., 2004. Gated viewing and high-accuracy three-dimensional laser radar. *Appl. Opt.* 43 (24).
- Buttafava, M., Zeman, J., Tosi, A., Eliceiri, K., Velten, A., 2015. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Opt. Express* 23 (16).
- Chandrasekhar, S., 1960. Radiative Transfer. Dover, ISBN: 0486605906.
- Charbon, E., 2007. Will avalanche photodiode arrays ever reach 1 megapixel. In: *International Image Sensor Workshop*.
- Conroy, R.M., Dorrington, A.A., Künemeyer, R., Cree, M.J., 2009. Range imager performance comparison in homodyne and heterodyne operating modes. In: *IS&T/SPIE Electronic Imaging*.
- Dai, H., He, W., Miao, Z., Chen, Y., Gu, G., 2013. Three-dimensional active imaging using compressed gating. In: *Intl. Symp. Photo. Det. Imag.*
- D'Eon, E., Irving, G., 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph.* 30 (4).
- Dorrington, A.A., Cree, M.J., Payne, A.D., Conroy, R.M., Carnegie, D.A., 2007. Achieving sub-millimeter precision with a solid-state full-field heterodyning range imaging camera. *Measure. Sci. Tech.* 18 (9).
- Dorrington, A.A., Godbaz, J.P., Cree, M.J., Payne, A.D., Streeter, L.V., 2011. Separating true range measurements from multi-path and scattering interference in commercial range cameras. In: *IS&T/SPIE Electronic Imaging*.
- Durand, F., Holzschuch, N., Soler, C., Chan, E., Sillion, F.X., 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24 (3).
- Falje, D., 2009. Improvements of the 3D images captured with time-of-flight cameras, arXiv preprint [arXiv:0909.5656](https://arxiv.org/abs/0909.5656).
- Freedman, D., Smolin, Y., Krupka, E., Leichter, I., Schmidt, M., 2014. SRA: Fast removal of general multipath for ToF sensors. In: *ECCV*.
- Fuchs, S., 2010. Multipath interference compensation in time-of-flight camera images. In: *IEEE Intl. Conf. Patt. Recog.*
- Fuchs, S., Hirzinger, G., 2008. Extrinsic and depth calibration of ToF-cameras. In: *CVPR*.
- Gao, L., Liang, J., Li, C., Wang, L.V., 2014. Single-shot compressed ultrafast photography at one hundred billion frames per second. *Nature* 516 (7529).
- Garipey, G., Krstajić, N., Henderson, R., Li, C., Thomson, R.R., Buller, G.S., Heshmat, B., Raskar, R., Leach, J., Faccio, D., 2015. Single-photon sensitive light-in-flight imaging. *Nature Comm.* 6.
- Gkioulekas, I., Levin, A., Durand, F., Zickler, T., 2015. Micron-scale light transport decomposition using interferometry. *ACM Trans. Graph.* 34 (4).
- Godbaz, J.P., Cree, M.J., Dorrington, A.A., 2008. Mixed pixel return separation for a full-field ranger. In: *IEEE Inter. Conf. Imag. Vis. Comp.*
- Godbaz, J.P., Cree, M.J., Dorrington, A.A., 2009. Multiple return separation for a full-field ranger via continuous waveform modeling. In: *IS&T/SPIE Electronic Imaging*.
- Godbaz, J.P., Cree, M.J., Dorrington, A.A., 2010. Extending amcw lidar depth-of-field using a coded aperture. In: *ACCV*.
- Godbaz, J.P., Cree, M.J., Dorrington, A.A., 2012. Closed-form inverses for the mixed pixel/multipath interference problem in amcw lidar. In: *IS&T/SPIE Electronic Imaging*.
- Goral, C.M., Torrance, K.E., Greenberg, D.P., Battaile, B., 1984. Modeling the interaction of light between diffuse surfaces. *SIGGRAPH Comput. Graph.* 18 (3).
- Gupta, M., Kadambi, A., Bhandari, A., Raskar, R., 2015a. Computational time of flight.
- Gupta, M., Nayyar, S.K., Hullin, M.B., Martin, J., 2015b. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM Trans. Graph.* 34 (5).
- Gupta, O., Willwacher, T., Velten, A., Veeraraghavan, A., Raskar, R., 2012. Reconstruction of hidden 3d shapes using diffuse reflections. *Opt. Express* 20 (17).
- Gutierrez, D., Muñoz, A., Anson, O., Seron, F., 2005. Non-linear volume photon mapping. In: *EGSR*.
- Gutierrez, D., Narasimhan, S.G., Jensen, H.W., Jarosz, W., 2008. Scattering. In: *ACM SIGGRAPH ASIA 2008 Courses*.
- Hachisuka, T., Jarosz, W., Georgiev, I., Kaplanyan, A., Nowrouzezahrai, D., 2013. State of the art in photon density estimation. In: *ACM SIGGRAPH ASIA 2013 Courses*.
- Hamamatsu, 2012. Guide to streak cameras. [http://sales.hamamatsu.com/assets/pdf/catsandguides/e\\_streakh.pdf](http://sales.hamamatsu.com/assets/pdf/catsandguides/e_streakh.pdf).
- Han, P.Y., Cho, G.C., Zhang, X.-C., 2000. Time-domain transillumination of biological tissues with terahertz pulses. *Opt. Lett.* 25 (4).
- Hansard, M., Lee, S., Choi, O., Horaud, R.P., 2012. Time-of-Flight Cameras: Principles, Methods and Applications. Springer Science & Business Media.
- Hebert, M., Krotkov, E., 1992. 3D measurements from imaging laser radars: how good are they? *Image Vis. Comput.* 10 (3).
- Heide, F., Heidrich, W., Hullin, M., Wetzstein, G., 2015. Doppler time-of-flight imaging. *ACM Trans. Graph.* 34 (4).
- Heide, F., Hullin, M., Gregson, J., Heidrich, W., 2013. Low-budget transient imaging using photonic mixer devices. *ACM Trans. Graph.* 32 (4).
- Heide, F., Xiao, L., Heidrich, W., Hullin, M.B., 2014a. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In: *CVPR*.
- Heide, F., Xiao, L., Kolb, A., Hullin, M.B., Heidrich, W., 2014b. Imaging in scattering media using correlation image sensors and sparse convolutional coding. *Opt. Express* 22 (21).
- Heshmat, B., Satat, G., Barsi, C., Raskar, R., 2014. Single-shot ultrafast imaging using parallax-free alignment with a tilted lenslet array. In: *CLEO: Science and Innovations*.
- Hu, X., Deng, Y., Lin, X., Suo, J., Dai, Q., Barsi, C., Raskar, R., 2014. Robust and accurate transient light transport decomposition via convolutional sparse coding. *Opt. Lett.* 39 (11).
- Huang, D., Swanson, E.A., Lin, C.P., Schuman, J.S., Stinson, W.G., Chang, W., Hee, M.R., Flotte, T., Gregory, K., Puliafito, C.A., Fujimoto, J.G., 1991. Optical coherence tomography. *Science* 254 (5035).
- Hullin, M.B., 2014. Computational imaging of light in flight. In: *SPIE/COS Photonics Asia*.
- Ihrke, I., Ziegler, G., Tevs, A., Theobalt, C., Magnor, M., Seidel, H.-P., 2007. Eikonal rendering: Efficient light transport in refractive objects. *ACM Trans. Graph.* 26 (3).
- Jamtsho, S., Lichti, D.D., 2010. Modeling scattering distortion in 3D range camera. *Intl. Arch. Photo. Remote Sens. Spat. Info. Sci.* 38 (5).
- Jarabo, A., 2012. Femto-photography: Visualizing Light in Motion. Universidad de Zaragoza.
- Jarabo, A., Marco, J., Muñoz, A., Buisan, R., Jarosz, W., Gutierrez, D., 2014. A framework for transient rendering. *ACM Trans. Graph.* 33 (6).
- Jarabo, A., Masia, B., Velten, A., Barsi, C., Raskar, R., Gutierrez, D., 2013. Rendering relativistic effects in transient imaging. In: *Congreso Español de Informática Gráfica*.

- Jarabo, A., Masia, B., Velten, A., Barsi, C., Raskar, R., Gutierrez, D., 2015. Relativistic effects for time-resolved light transport. *Comp. Graph. Forum* 34 (8).
- Jensen, H.W., 2001. Realistic image synthesis using photon mapping. AK Peters, ISBN: 1568811470.
- Jiménez, D., Pizarro, D., Mazo, M., Palazuelos, S., 2012. Modeling and correction of multipath interference in time of flight cameras. In: *CVPR*.
- Jiménez, D., Pizarro, D., Mazo, M., Palazuelos, S., 2014. Modeling and correction of multipath interference in time of flight cameras. *Image Vision Comput.* (ISSN: 0262-8856) 32 (1).
- Jongenelen, A.P., Carnegie, D.A., Payne, A.D., Dorrington, A.A., 2010. Maximizing precision over extended unambiguous range for ToF range imaging systems. In: *I2MTC*.
- Kadambi, A., Schiel, J., Raskar, R., 2016a. Macroscopic interferometry: Rethinking depth estimation with frequency-domain time-of-flight. In: *CVPR*.
- Kadambi, A., Taamazyan, V., Jayasuriya, S. and Raskar, R., 2015. Frequency domain ToF: encoding object depth in modulation frequency, arXiv preprint arXiv:1503.01804.
- Kadambi, A., Whyte, R., Bhandari, A., Streeter, L., Barsi, C., Dorrington, A., Raskar, R., 2013. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Trans. Graph.* 32 (6).
- Kadambi, A., Zhao, H., Shi, B., Raskar, R., 2016b. Occluded Imaging with Time-of-Flight Sensors. *ACM Trans. Graph.* (ISSN: 0730-0301) 35 (2).
- Kajiya, J.T., 1986. The rendering equation. In: *SIGGRAPH*.
- Kavli, T., Kirkhus, T., Thielemann, J.T., Jagielski, B., 2008. Modeling and compensating measurement errors caused by scattering in time-of-flight cameras. In: *Opt. Eng. Appl.*
- Keller, M., Kolb, A., 2009. Real-time simulation of time-of-flight sensors. *Sim. Model. Prac. Theory* 17 (5).
- Keller, M., Orthmann, J., Kolb, A., Peters, V., 2007. A simulation framework for time-of-flight sensors. In: *ISSCS*.
- Kirmani, A., Benedetti, A., Chou, P.A., 2013. Spumic: simultaneous phase unwrapping and multipath interference cancellation in time-of-flight cameras using spectral methods. In: *ICME*.
- Kirmani, A., Hutchison, T., Davis, J., Raskar, R., 2009. Looking around the corner using transient imaging. In: *ICCV*.
- Kirmani, A., Hutchison, T., Davis, J., Raskar, R., 2011a. Looking around the corner using ultrafast transient imaging. *IJCV* 95 (1).
- Kirmani, A., Velten, A., Hutchison, T., Lawson, M.E., Goyal, V.K., Raskar, M.B.R., 2011b. Reconstructing an image on a hidden plane using ultrafast imaging of diffuse reflections.
- A., Kirmani, H., Jeelani, V., Montazerhodjat, V.K., Goyal, 2012. Diffuse imaging: Creating optical images with unfocused time-resolved illumination and sensing. *IEEE Signal Process. Lett.* 19 (1).
- Kirmani, A., Venkatraman, D., Shin, D., Colaço, A., Wong, F.N., Shapiro, J.H., Goyal, V.K., 2014. First-photon imaging. *Science* 343 (6166).
- Klein, J., Peters, C., Martín, J., Laurenzis, M., Hullin, M.B., 2016. Tracking objects outside the line of sight using 2d intensity images. *Sci. Rep.* 6.
- Kolb, A., Barth, E., Koch, R., Larsen, R., 2010. Time-of-Flight sensors in computer graphics. *Comput. Graph. Forum* 29 (1).
- Křivánek, J., Georgiev, I., Kaplanyan, A., Canada, J., 2013. Recent advances in light transport simulation: theory and practice. In: *ACM SIGGRAPH 2013 Courses*.
- Lambers, M., Hoberg, S., Kolb, A., 2015. Simulation of time-of-flight sensors for evaluation of chip layout variants. *IEEE Sens. J.* 15 (7).
- Lange, R., Seitz, P., 2001. Solid-state time-of-flight range camera. *IEEE J. Quant. Elec.* 37 (3).
- Lange, R., Seitz, P., Biber, A., Lauthermann, S.C., 2000. Demodulation pixels in CCD and CMOS technologies for time-of-flight ranging. In: *Elec. Imaging*.
- Lau, A.K., Tang, A.H., Xu, J., Wei, X., Wong, K.K., Tsia, K.K., 2016. Optical time stretch for high-speed and high-throughput imaging from single-cell to tissue-wide scales. *IEEE J. Sel. Top. Quantum Electron.* 22 (4).
- Laurenzis, M., Bacher, E., 2011. Image coding for three-dimensional range-gated imaging. *Appl. Opt.* 50 (21).
- Laurenzis, M., Christnacher, F., Monnin, D., 2007. Long-range three-dimensional active imaging with superresolution depth mapping. *Opt. Lett.* 32 (21).
- Laurenzis, M., Velten, A., 2014. Nonline-of-sight laser gated viewing of scattered photons. *Opt. Eng.* 53 (2).
- Laurenzis, M., Woiselle, A., 2014. Laser gated-viewing advanced range imaging methods using compressed sensing and coding of range-gates. *Opt. Eng.* 53 (5).
- Lee, S., Shim, H., 2015. Skewed stereo time-of-flight camera for translucent object imaging. *Image Vis. Comput.* 43.
- Li, L., Wu, L., Wang, X., Dang, E., 2012. Gated viewing laser imaging with compressive sensing. *App. Opt.* 51 (14).
- Lima, I.T., Kalra, A., Sherif, S.S., 2011. Improved importance sampling for monte carlo simulation of time-domain optical coherence tomography. *Bio. Opt. Exp.* 2 (5).
- Lin, J., Liu, Y., Hullin, M.B., Dai, Q., 2014. Fourier analysis on transient imaging with a multifrequency time-of-flight camera. In: *CVPR*.
- Lin, J., Liu, Y., Suo, J., Da, Q., 2016. Frequency-Domain transient imaging. *IEEE Trans. Pattern Anal. Mach. Intell.* <http://dx.doi.org/10.1109/TPAMI.2016.2560814>.
- Lindner, M., Schiller, I., Kolb, A., Koch, R., 2010. Time-of-Flight sensor calibration for accurate range sensing. *Comput. Vis. Image Underst.* 114 (12).
- Marco, J., 2013. Transient Light Transport in Participating Media. Universidad de Zaragoza.
- Meister, S., Nair, R., Jähne, B., Kondermann, D., 2013a. Photon mapping based simulation of multi-path reflection artifacts in time-of-flight sensors.
- Meister, S., Nair, R., Kondermann, D., 2013b. Simulation of time-of-flight sensors using global illumination. In: *VMV*.
- Moravec, H.P., 1981. 3D graphics and the wave theory. *SIGGRAPH Comput. Graph.* 15 (3).
- Musbach, A., Meyer, G.W., Reitich, F., Oh, S.H., 2013. Full wave modeling of light propagation and reflection. *Comput. Graph. Forum* 32 (6).
- Naik, N., Barsi, C., Velten, A., Raskar, R., 2014. Estimating wide-angle, spatially varying reflectance using time-resolved inversion of backscattered light. *JOSA A* 31 (5).
- Naik, N., Kadambi, A., Rhemann, C., Izadi, S., Raskar, R., Bing Kang, S., 2015. A light transport model for mitigating multipath interference in time-of-flight sensors. In: *IEEE Computer Vision and Pattern Recognition*.
- Naik, N., Zhao, S., Velten, A., Raskar, R., Bala, K., 2011. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graph.* 30.
- Nair, R., Meister, S., Lambers, M., Balda, M., Hofmann, H., Kolb, A., Kondermann, D., Jähne, B., 2013. Ground truth for evaluating time of flight imaging. In: *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications*.
- Nakagawa, K., Iwasaki, A., Oishi, Y., Horisaki, R., Tsukamoto, A., Nakamura, A., Hirohara, K., Liao, H., Ushida, T., Goda, K., et al., 2014. Sequentially timed all-optical mapping photography (stamp). *Nature Photonics* 8 (9).
- Nayar, S.K., Krishnan, G., Grossberg, M.D., Raskar, R., 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.* 25 (3).
- Ng, R., Ramamoorthi, R., Hanrahan, P., 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22 (3).
- O'Toole, M., Heide, F., Xiao, L., Hullin, M.B., Heidrich, W., Kutulakos, K.N., 2014. Temporal frequency probing for 5D Transient Analysis of Global Light Transport. *ACM Trans. Graph.* 33 (4).
- O'Toole, M., Raskar, R., Kutulakos, K.N., 2012. Primal-dual coding to probe light transport. *ACM Trans. Graph.* 31 (4).
- Pandharkar, R., 2011. Hidden Object Doppler: Estimating Motion, Size and Material Properties of Moving Non-Line-of-Sight Objects in Cluttered Environments. Massachusetts Institute of Technology (Ph.D. thesis).
- 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: *CVPR*.
- Periyasamy, V., Pramanik, M., 2016. Importance sampling-based monte carlo simulation of time-domain optical coherence tomography with embedded objects. *App. Opt.* 55 (11).
- Peters, C., Klein, J., Hullin, M.B., Klein, R., 2015. Solving trigonometric moment problems for fast transient imaging. *ACM Trans. Graph.* 34 (6).
- Pitts, P., Benedetti, A., Slaney, M., Chou, P., 2014. Time of flight tracer.
- Qiao, H., Lin, J., Liu, Y., Hullin, M.B., Dai, Q., 2015. Resolving transient time profile in ToF imaging via log-sum sparse regularization. *Opt. Lett.* 40 (6).
- Raskar, R., Davis, J., 2008. 5d Time-Light Transport Matrix: What Can We Reason About Scene Properties? Technical Report MIT. <http://hdl.handle.net/1721.1/67888>.
- Raviv, D., Barsi, C., Naik, N., Feigin, M., Raskar, R., 2014. Pose estimation using time-resolved inversion of diffuse light. *Opt. Express* 22 (17).
- Remondino, F., Stoppa, D., 2013. TOF Range-Imaging Cameras. Springer, ISBN: 978-3642275234.
- Satat, G., Heshmat, B., Barsi, C., Raviv, D., Chen, O., Bawendi, M.G., Raskar, R., 2015a. Locating and classifying fluorescent tags behind turbid layers using time-resolved inversion. *Nature Comm.* 6.
- Satat, G., Heshmat, B., Naik, N., Redo-Sanchez, A., Raskar, R., 2016. Advances in ultrafast optics and imaging applications. In: *SPIE Defense+ Security*.
- Satat, G., Raviv, D., Heshmat, B., Raskar, R., 2015b. Imaging through thick turbid medium using time-resolved measurement. In: *Imaging and Applied Optics* 2015.
- Schäfer, H., Lenzen, F., Garbe, C.S., 2014. Model based scattering correction in time-of-flight cameras. *Opt. Express* 22 (24).
- Schwarte, R., Xu, Z., Heinel, H.-G., Olk, J., Klein, R., Buxbaum, B., Fischer, H., Schulte, J., 1997. New electro-optical mixing and correlating sensor: facilities and applications of the photonic mixer device (PMD). In: *Lasers Opt. Manu. III*.
- Shim, H., Lee, S., 2016. Recovering translucent objects using a single time-of-flight depth camera. *IEEE Trans. Circuits Syst. Video Technol.* 26 (5).
- Shrestha, S., Heide, F., Heidrich, W., Wetzstein, G., 2016. Computational imaging with multi-camera time-of-flight systems. *ACM Trans. Graph.* 35 (4).
- Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Taylor & Francis.
- Smith, A., Skorupski, J., Davis, J., 2008. Transient Rendering. Technical Report UCSC-SOE-08-26 School of Engineering. University of California, Santa Cruz.
- Su, S., Heide, F., Swanson, R., Klein, J., Callenberg, C., Hullin, M., Heidrich, W., 2016. Material classification using raw time-of-flight measurements. In: *CVPR*.

- Tadano, R., Pediredla, A.K., Mitra, K. and Veeraraghavan, A. (2015) Spatial Phase-Sweep: Increasing temporal resolution of transient imaging using a light source array. arXiv preprint [arXiv:1512.06539](https://arxiv.org/abs/1512.06539).
- Tanaka, K., Mukaigawa, Y., Kubo, H., Matsushita, Y., Yagi, Y., 2016. Recovering transparent shape from time-of-flight distortion. In: CVPR.
- Tsakatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck, J.-L., Tsakalides, P., 2012. Active range imaging via random gating. In: SPIE Sec. Def.
- Tsakatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck, J.L., Tsakalides, P., 2013. Compressed gated range sensing. In: SPIE Opt. Eng. Appl.
- Tsakatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck, J.-L., Tsakalides, P., 2015. Multireturn compressed gated range imaging. *Opt. Eng.* 54 (3).
- Tsai, C.-Y., Veeraraghavan, A., Sankaranarayanan, A.C., 2016. Shape and reflectance from two-bounce light transients. In: ICCP.
- Veach, E., 1997. Robust Monte Carlo methods for light transport simulation (Ph.D. thesis).
- Velten, A., Willwacher, T., Gupta, O., Veeraraghavan, A., Bawendi, M.G., Raskar, R., 2012a. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Comm.* (3).
- Velten, A., Wu, D., Jarabo, A., Masia, B., Barsi, C., Joshi, C., Lawson, E., Bawendi, M., Gutierrez, D., Raskar, R., 2013. Femto-Photography: capturing and visualizing the propagation of light. *ACM Trans. Graph.* 32 (4).
- Velten, A., Wu, D., Jarabo, A., Masia, B., Barsi, C., Lawson, E., Joshi, C., Gutierrez, D., Bawendi, M.G., Raskar, R., 2012b. Relativistic ultrafast rendering using time-of-flight imaging. In: ACM SIGGRAPH 2012 Talks.
- Velten, A., Wu, D., Masia, B., Jarabo, A., Barsi, C., Joshi, C., Lawson, E., Bawendi, M., Gutierrez, D., Raskar, R., 2016. Transient imaging of macroscopic scenes at picosecond resolution. *Commun. ACM* 34 (9).
- Whyte, R., Streeter, L., Cree, M.J., Dorrington, A.A., 2015. Resolving multiple propagation paths in time of flight range cameras using direct and global separation methods. *Opt. Eng.* 54 (11).
- Wu, D., O'Toole, M., Velten, A., Agrawal, A., Raskar, R., 2012a. Decomposing global light transport using time of flight imaging. In: CVPR.
- Wu, D., Velten, A., O'Toole, M., Masia, B., Agrawal, A., Dai, Q., Raskar, R., 2014. Decomposing global light transport using time of flight imaging. *IJCV* 107 (2).
- Wu, J., Wang, Y., Perelman, L., Itzkan, I., Dasari, R.R., Feld, M.S., 1995. Time-resolved multichannel imaging of fluorescent objects embedded in turbid media. *Opt. Lett.* 20 (5).
- Wu, D., Wetzstein, G., Barsi, C., Willwacher, T., O'Toole, M., Naik, N., Dai, Q., Kutulakos, K., Raskar, R., 2012b. Frequency analysis of transient light transport with applications in bare sensor imaging. In: ECCV.
- Xiao, L., Heide, F., O'Toole, M., Kolb, A., Hullin, M.B., Kutulakos, K., Heidrich, W., 2015. Defocus deblurring and superresolution for time-of-flight depth cameras. In: CVPR.
- Zhang, X., Yan, H., 2011. Three-dimensional active imaging with maximum depth range. *Appl. Opt.* 50 (12).
- Zhang, X., Yan, H., Lv, J., 2012. Multireturn three-dimensional active imaging based on compressive sensing. *Opt. Lett.* 37 (23).