

# Towards Transient Imaging at Interactive Rates with Single-Photon Detectors

David Lindell  
Stanford University  
lindell@stanford.edu

Matthew O’Toole  
Stanford University  
motoole@stanford.edu

Gordon Wetzstein  
Stanford University  
gordon.wetzstein@stanford.edu

## Abstract

Active imaging at the picosecond timescale reveals transient light transport effects otherwise not accessible by computer vision and image processing algorithms. For example, analyzing the time of flight of short laser pulses emitted into a scene and scattered back to a detector allows for depth imaging, which is crucial for autonomous driving and many other applications. Moreover, analyzing or removing global light transport effects from photographs becomes feasible.

While several transient imaging systems have recently been proposed using various imaging technologies, none is capable of acquiring transient images at interactive frame rates. In this paper, we present an imaging system that leverages single-photon avalanche diodes together with a pulsed picosecond laser to record transient images with up to 25 Hz at a low spatial resolution of  $64 \times 80$  pixels or 1 Hz at a moderate resolution of  $256 \times 250$  pixels. We show several transient video clips recorded with this system and demonstrate transient imaging applications, including direct-global light transport separation and enhanced depth imaging.

## 1. Introduction and Overview

The emergence of ultra-fast photodetector arrays brings time-resolved or *transient* imaging within reach for widespread application in remote sensing, robotic and machine vision, and computational photography. Whereas a conventional imager integrates photoelectrons over a certain exposure time and reports the measurements as a 2D array of intensity values, a transient image is a 3D volume that contains photon counts for short time bins, typically at the picosecond time scale, for each 2D pixel location. Thus, each pixel not only measures a single intensity value but a histogram that shows the temporal distribution of the time of arrival of the photoelectrons throughout the exposure time. A camera equipped with such technology has the potential to acquire precise depth maps of a scene [20, 15, 14, 22], see through scattering media [11, 23], disambiguate direct and indirect light transport [26], resolve

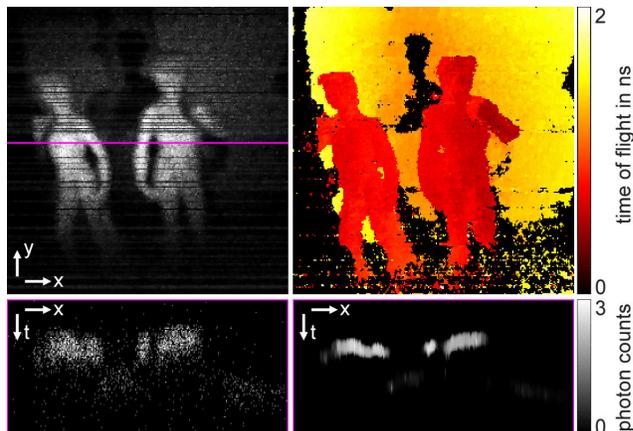


Figure 1. A transient image, i.e., one frame of a transient video, is a spatio-temporal volume that captures the time-of-flight of a short light pulse propagating through a scene. With the proposed imaging system, we can record such transient images at several frames per second. The raw data (lower left) is very noisy and only records about a few photon counts in each voxel of the transient image. We process this data (top left, bottom right) to mitigate noise and deblur the temporal dimension. The top right image shows a color-coded visualization of the light in flight.

motion [8] and shapes [24, 19] around corners, or operate without the need for focusing optics [27].

Although transient imaging has been approached using a number of different technologies, including optical holography [1], digital interferometry [9], streak cameras [25, 6], photon mixer devices (PMDs) [13, 10], and single-photon avalanche diodes (SPADs) [7, 17], all of these technologies have different drawbacks. Among the digital techniques, interferometry requires acquisition times on the order of hours, streak cameras are expensive and not compatible with most imaging systems used in practice, and the temporal resolution of PMDs is limited to the nanosecond timescale. SPAD arrays have become increasingly popular and widely available, for example in light detection and ranging (LiDAR) applications, but unfortunately SPAD sensors are currently available with only low pixel counts. To record a scene at a reasonably high resolution, some form

	Acquisition time	Temporal resolution	Spatial resolution
Holography [1]	12 s	800 ps	high
Streak camera [25]	>1 h	0.3-5.7 ps	1000 × 672
PMD sensor [10]	90 s	1000 s	160 × 120
Optical interferometry [9]	>1 h	0.033 ps	655 × 648
SPAD [17]	64 s	300/100 ps	320 × 256
SPAD (ours)	0.04 s 1 s	494/286 ps	64 × 80 256 × 250

Table 1. Comparison of spatial and temporal resolution as well as acquisition times of several transient imaging technologies. The temporal jitter of the SPAD systems is reported for raw measurements and after deconvolution. Our system achieves shorter acquisition times with a corresponding decrease in measurement SNR; we leverage post-processing to improve the quality of the final results. A detailed overview of transient imaging technologies and applications can be found in [12].

of temporal scanning has to be employed, which severely limits achievable frame rates. The fastest transient imaging system to date records at 1 frame per minute [17]. Please refer to Table 1 for a more detailed overview of these techniques along with reported spatial and temporal resolution as well as acquisition times.

In this paper, we address the tradeoff between image resolution and frame rate of SPAD-based cameras from a systems level. To this end, we follow O’Toole et al. [17] and combine a linear array of 256 SPADs [5] with a mirror galvanometer and a pulsed picosecond laser. Our system is capable of recording transient videos with a spatial resolution of  $64 \times 80$  pixels at 25 Hz or with  $256 \times 250$  pixels at 1 Hz. The increased acquisition rate of our systems comes by trading off exposure time, resulting in a decrease in measurement SNR. Despite this tradeoff, we demonstrate that post-processing of the captured measurements greatly improves the quality of our results and so the additional computation can compensate for lower SNR. We also demonstrate several applications, including direct-global illumination separation and enhanced depth imaging. With this work, we take first steps towards making transient imaging practical for dynamic scenes (see example in Fig. 1).

## 2. System

In this section, we describe the hardware and software used for recording and processing transient video clips and we outline different spatio-angular resolution tradeoffs offered by our system.

### 2.1. Hardware

The sensing platform used in the system is a LinoSPAD [5] and contains a  $256 \times 1$  array of single-photon avalanche diodes. The time-to-digital converters (TDCs) are implemented on a Xilinx Spartan 6 field-programmable

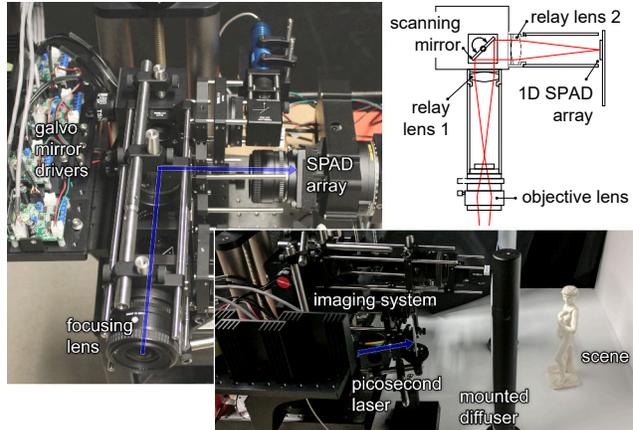


Figure 2. Prototype single-photon avalanche diode (SPAD) imaging system. The instrument uses a linear array of 256 SPADs together with a mirror galvanometer system to scan transient images at interactive rates. A laser is positioned next to the system to illuminate the scene from arbitrary perspectives.

gate array (FPGA). Each SPAD has a diameter of  $24 \mu\text{m}$ , a fill factor of 41%, and a photon detection efficiency of 20% at 450 nm. The dark count rate (i.e., number of falsely detected events) varies over the array with a median value of 2.5 kHz. The SPAD array is enclosed in a 3D-printed housing and mounted using standard Thorlabs opto-mechanical components. These components include a mechanical cage system, two optical relay lenses, and a mirror galvanometer that allows us to horizontally scan the vertical scanline imaged by the 1D SPAD array. A National Instruments data acquisition device (NI-DAQ USB-6343) acts as the master clock in the system, synchronizing the acquisition time of the SPADs, the mirror, and also the laser. The laser is an ALPHALAS PICOPOWER-LD-450-50 that has a wavelength of 450 nm, peak power of 450 mW, and a repetition rate configured to 25 MHz. Photographs of the imaging system and a schematic that illustrates the optical path are shown in Figure 2.

### 2.2. Calibration and System Tradeoffs

Calibration of the system involves several steps, including focusing the system, positioning the scanning mirrors, and characterizing the laser pulse shape and temporal jitter of the SPAD. To focus the system, we direct the pulsed laser source to a point within the field of view, and adjust the position of the objective lens and SPAD sensor to produce a sharp image of the laser point. The extreme horizontal positions of the scanning mirror are also configured to achieve a broad field of view while avoiding vignetting from the relay and objective lenses.

To calibrate the pulse shape of the laser and temporal jitter of the SPADs, we repeatedly fire laser pulses onto a

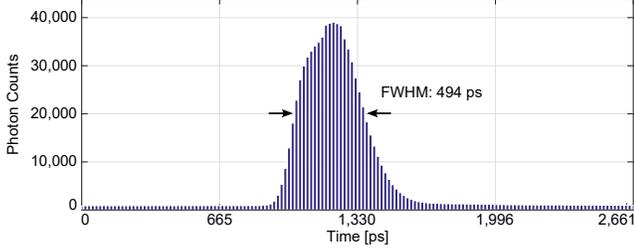


Figure 3. Temporal point spread function measured with the system. The full width at half maximum (FWHM) of the imaging system is approx. 494 ps, which combines the pulse width of the laser and the temporal jitter of the detector. After deconvolution, the FWHM of the PSF is approx. 286 ps.

planar surface and measure the response. As seen in Figure 3, the recorded temporal point spread function (PSF) is not symmetric and has a footprint of 494 ps, which corresponds to a path length of about 15 cm. Using a laser with a shorter pulse width or one with a longer wavelength can reduce the temporal spread of the PSF. However, the photon detection efficiency of the SPAD is highest for the wavelength used in this system. After deconvolution, the PSF has a footprint of approx. 286 ps or 9 cm.

The system can also be configured to trade off the spatial resolution and frame rate. We acquire measurements in two different modes to demonstrate this capability: a high-resolution ( $256 \times 250$ ), low frame rate (1 Hz) mode, and a low-resolution ( $64 \times 80$ ), high frame rate (25 Hz) mode. The high-resolution mode captures measurements with each of the four sets of 64 pixels in succession, while only a single set of 64 pixels is used in the low-resolution mode. The exposure time for each scan position in the high-resolution mode is 900  $\mu\text{s}$  and 450  $\mu\text{s}$  in the low-resolution mode.

### 2.3. Data Processing

Following [21, 17], we model the photon flux  $r$  incident on a detector during a discrete time interval  $t$ , relative to an emitted pulse  $g$ , as

$$r[t] = (\tau * g)[t] + a[t], \quad (1)$$

where  $*$  is the 1D convolution operator,  $a[t]$  is the ambient photon flux and  $\tau$  is the temporal impulse response or transient image of the scene. The impulse response models photon travel time, albedo, and any other light transport effect.

Single photon avalanche diodes typically only timestamp up to one photon event for any emitted laser pulse. After a photon event is detected, the SPAD requires approximately 100 ns to reset. During this *dead time*, no photon events are recorded. The SPAD image formation can be modeled as a set of Bernoulli trials that are repeated by firing  $N$  laser pulses. The sampled photon events are aggreg-

gated in a histogram  $h$ , which follows a Poisson distribution (cf. [22, 17]), such that

$$h[t] \sim \mathcal{P}(N(\eta(r * f)[t] + d)), \quad (2)$$

where  $d$  is the dark count rate and  $f$  is the detector jitter (i.e., the uncertainty in the time stamping process).

The raw histograms can be processed to estimate the discrete latent transient image  $\tau \in \mathbb{R}^{n_x n_y n_t \times 1}$ . To this end, the temporal convolutions by the laser pulse shape  $g$  and the detector jitter  $f$  are combined in a convolution matrix  $\mathbf{A} \in \mathbb{R}^{n_x n_y n_t \times n_x n_y n_t}$ . Estimating  $\tau$  is thus a joint denoising and deconvolution problem, which can be formulated as the following objective function

$$\begin{aligned} \underset{\{\tau\}}{\text{minimize}} \quad & -\log(p(\mathbf{h}|\mathbf{A}\tau)) + \Gamma(\tau), \quad (3) \\ \text{subject to} \quad & 0 \leq \tau \end{aligned}$$

Here,  $p(\mathbf{h}|\mathbf{A}\tau)$  is the likelihood of observing measurements  $\mathbf{h}$  for a known transient image  $\tau$  and  $\Gamma(\tau)$  is an additional prior on the recovered signal. For all results shown in the following, we use an isotropic 3D total variation (TV) prior and solve Equation 3 using the alternating direction method of multipliers (ADMM) [4]. Using this method with 25 iterations per transient image, we can process a transient image within a few minutes. For more details on transient image formation and reconstruction with SPADs, we refer the interested reader to [17].

### 2.4. Results

We recorded transient video clips of six scenes with varying tradeoffs between spatial resolution and frame rate. Figure 4 shows several spatio-temporal slices of two of these scenes. Both examples are recorded with a frame rate of 1 frame per second for 5 seconds with a spatial resolution of  $256 \times 250$  pixels and 1536 temporal histogram bins. Thus, the total resolution of these transient video clips is  $5 \times 256 \times 250 \times 1536$ . Storing the raw data uncompressed with double floating point precision occupies 3.93 GB of memory per dataset. Due to the sparsity of the data, compressing the data into MATLAB's binary .mat file format only requires a few hundred MB.

The raw measurements of both scenes (Fig. 4, rows 1 and 3) show that the low photon count of the recorded data results in very noisy data. Denoising and deconvolution (Fig. 4, rows 2 and 4), mitigates noise artifacts and also significantly sharpens the temporal profiles of the transient images. The remaining scenes are shown in the supplemental video. An overview of all recorded scenes with their individual settings are shown in Figure 5 and Table 2.

### 3. Applications

In this section, we outline applications of transient imaging, including separation of direct and global illumination

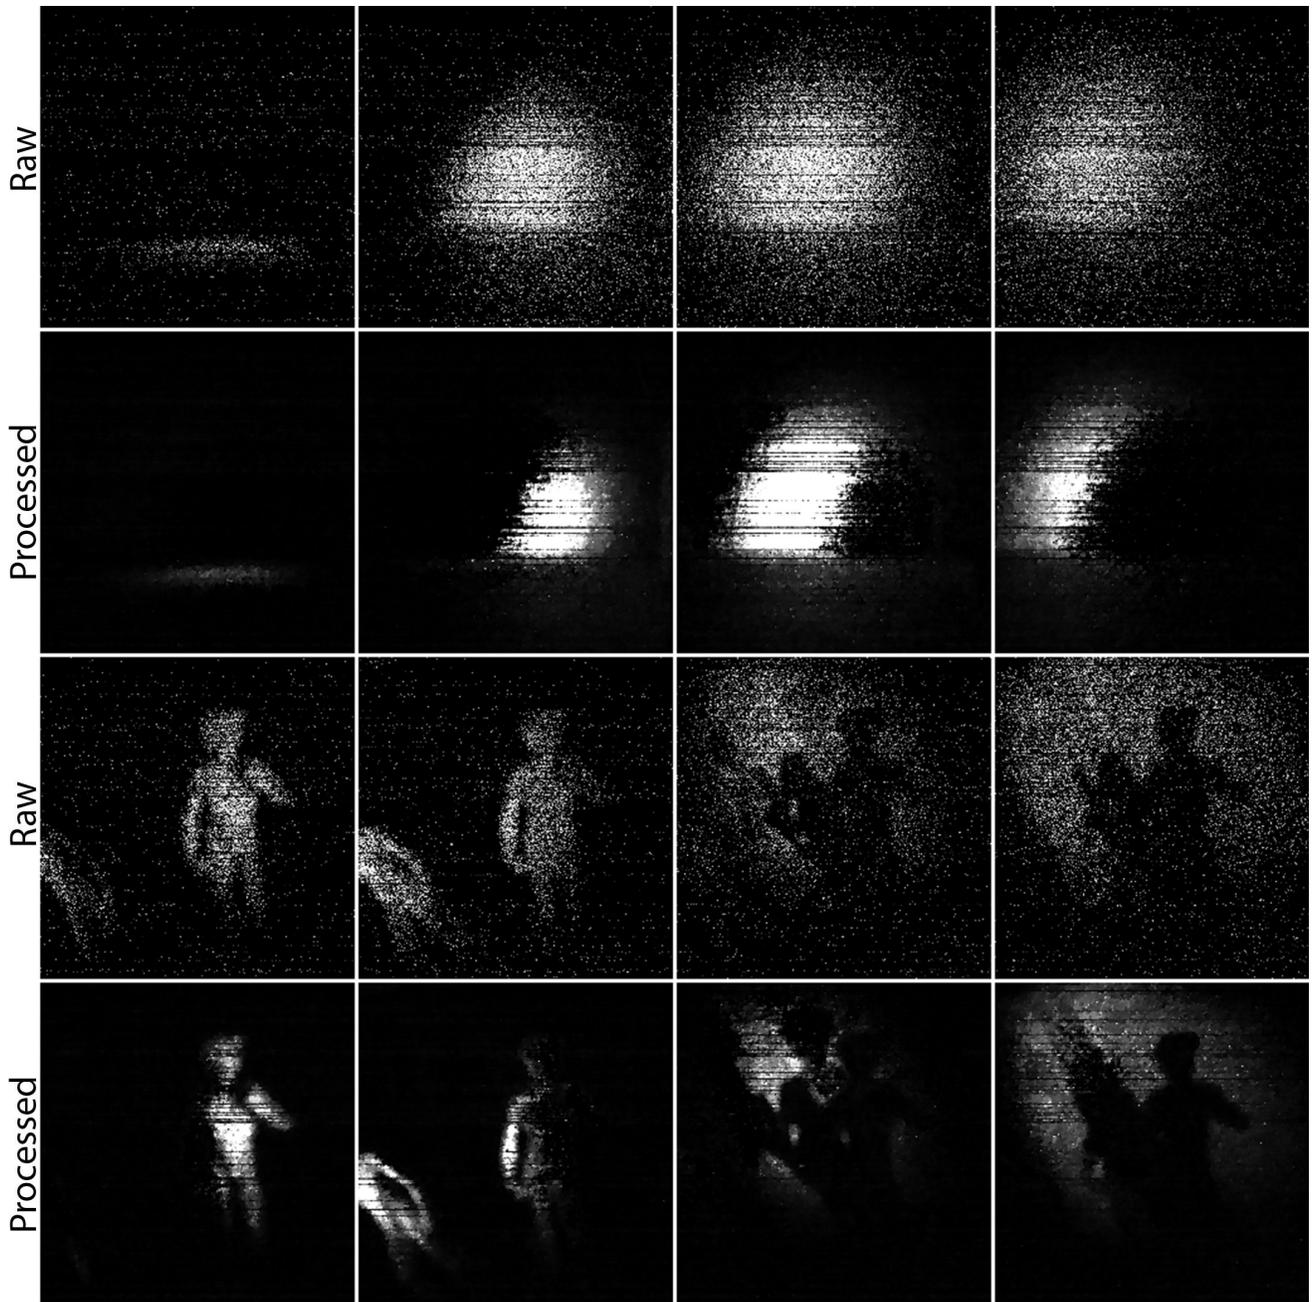


Figure 4. Two example scenes captured with our imaging system. The “Wall” scene (upper two rows) contains a simple planar floor and background. This scene intuitively visualizes the diffused laser pulse as it propagates through the scene. The first row shows several sub-frames of raw measurements and the second row shows the same frames after data processing, i.e. denoising and deconvolution. The bottom rows show raw and processed frames of a transient video of the “David” scene, including a statue and its reflection in a mirror. Line artifacts are due to SPAD pixels with a high dark count (number of false detections).

components as well as enhanced depth imaging. We capture these results in the high-resolution mode of the system for easier visualization.

### 3.1. Direct-Global Light Transport Separation

We demonstrate direct-global separation of light transport in the “Diffuser” scene for rapidly acquired transient images. In this scene, we place a planar, retroreflective ob-

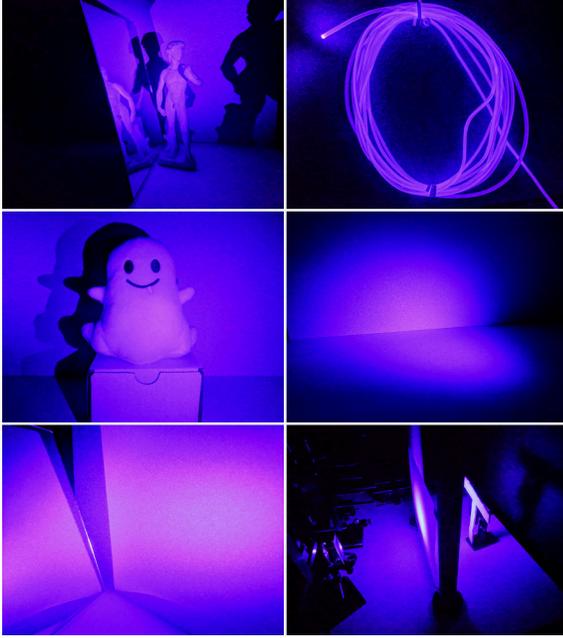


Figure 5. Photographs of recorded scenes under blue laser illumination: “David”, “Fiber”, “Ghost”, “Wall”, “Mirror”, “Diffuser”.

	Spatial Resolution	Exp.	FPS	Frames	Avg. Counts per Frame
David	$256 \times 250$	$900 \mu s$	1	5	$1.33 \times 10^6$
	$64 \times 80$	$450 \mu s$	25	125	$4.84 \times 10^4$
Wall	$256 \times 250$	$900 \mu s$	1	5	$1.41 \times 10^6$
	$64 \times 80$	$450 \mu s$	25	125	$7.42 \times 10^4$
Fiber	$256 \times 250$	$900 \mu s$	1	5	$1.61 \times 10^6$
	$64 \times 80$	$450 \mu s$	25	125	$5.81 \times 10^4$
Diffuser	$256 \times 250$	$900 \mu s$	1	1	$1.69 \times 10^6$
Ghost	$256 \times 250$	$900 \mu s$	1	1	$1.09 \times 10^6$
Mirror	$256 \times 250$	$900 \mu s$	1	1	$1.19 \times 10^6$

Table 2. Overview of experiments. Most of the scenes are captured for multiple frames with varying tradeoffs between spatial resolution and frame rate. All scenes have 1536 temporal histogram bins for each transient frame. The exposure is listed for each scanline of the SPAD pixels. The average photon count lists the total number of detected events divided by the number of frames.

ject with the shape of a “T” character behind a diffuser and direct a diffused laser pulse at the diffuser (see Fig. 5, bottom right). To decompose the captured measurements into direct and global components, we apply a method similar to that of Wu et al. [26]. As summarized in Algorithm 1, the onset of the direct component is determined by finding the time at which the gradients of the temporal histograms first exceed a provided threshold. The time at which the direct component ends is calculated by adding twice the full width at half maximum (FWHM) of the system PSF to the onset time. All direct light is assumed to be recorded before

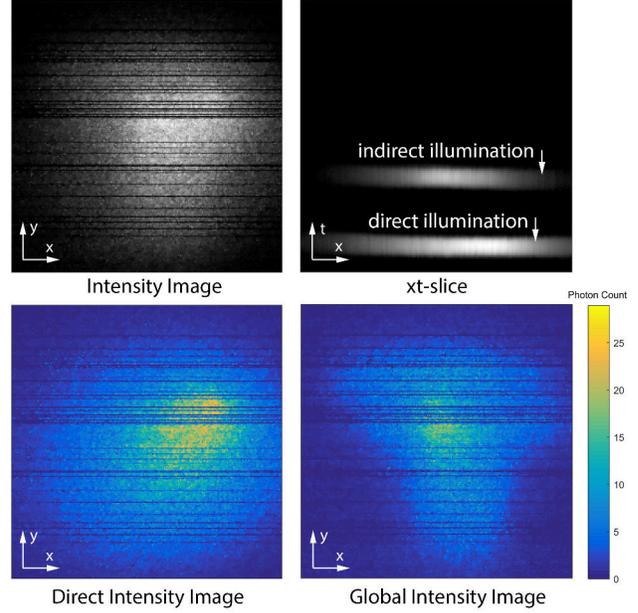


Figure 6. The “Diffuser” scene, consisting of a retroreflective “T” cutout placed behind a diffuser sheet. The intensity image captures all light from the scene including direct light from the diffuser sheet and indirect light from the “T” cutout. By separating the direct and global illumination components, the shape of the “T” can be identified in an intensity image of the global component.

the end of the direct pulse, and indirect light is assumed to occur after. This method works well in scenes where each imaged point in the scene receives a direct component, as the example shown in Figure 6.

---

**Algorithm 1** Algorithm for direct-global separation.

---

- 1: compute normalized gradient  $h'[t]/\max_{\{t\}} h'[t]$  of measurements  $h[t]$
  - 2: find onset of direct as  $t_1 = \min_{\{t\}} s.t. h'[t] > \alpha, \alpha = 0.1$
  - 3: compute end of direct as  $t_2 = t_1 \cdot 2 \cdot FWHM$
  - 4: let the direct component be  $h[t]$  for  $t_1 \leq t \leq t_2$
  - 5: let the global component be  $h[t]$  for  $t_2 < t$
- 

### 3.2. Enhanced Depth Imaging

Arguably the most important application of transient imaging is depth estimation. Long-range depth imaging is crucial in autonomous driving, in remote sensing, and in many other applications. The depth of a scene point is usually estimated by measuring the time of flight of an emitted laser pulse that scatters off an object and directly returns to the detector. Depth estimation algorithms aim to detect the peak in the recorded temporal histogram that corresponds to the direct illumination. In the absence of global illumination, usually there is only one peak in each histogram.

Indirectly reflected light, however, makes it more challenging to disambiguate which peak in the histogram resulted from direct light. In most cases, the brightest peak corresponds to the direct illumination, but problems arise when this assumption is violated.

One of most intuitive depth estimation algorithms is log-matched filtering [3]. A log-matched filter applies template matching by correlating the measured histograms with the laser pulse shape, followed by a search for the highest peak in the resulting signal. This operation is performed in the log domain to account for the Poisson-distributed noise in the measurements.

Figure 7 shows a scene that is challenging for the log-matched filter. An object is placed in the scene and receives direct illumination first, followed by a secondary pulse that arrives at a later time but that is brighter than the first one (Fig. 7, top right). We create this effect by passing light on a direct path from the illumination source through a wide-angle diffuser and also along a delayed path through a smaller-angle diffuser onto the scene. Similar situations may appear when the scene contains reflective, refractive, or specular objects that create caustics or otherwise focus the indirect light onto certain scene parts. In this case, the log-matched filter falsely detects the second pulse as the scene depth (Fig. 7, bottom left). Using the intuitive direct-global light transport separation algorithm outlined in the previous subsection, we can separate these illumination components in the transient image, apply the log-matched filter to the direct-only component, and thereby significantly improve the robustness of the depth estimation (Fig. 7, bottom right).

#### 4. Discussion

In summary, we present a system that allows for transient images to be recorded at up to 25 frames per second. The proposed imaging system uses a picosecond laser and a linear array of  $256 \times 1$  single-photon avalanche diodes. We use synchronized mirror galvanometers to scan a transient image and synchronize all system components. In addition to showing multiple recorded scenes which exhibit diverse direct and global light transport effects, we explore applications of such an imaging system, including direct-global light transport separation and enhanced depth imaging.

Our work builds on [17] but improves acquisition speed by up to two orders of magnitude. Yet, this improvement comes at the cost of reduced photon counts, which make it more challenging to reconstruct high-quality transient images. For example, our captured transient image of the “David” scene over a 1 second acquisition records an average of 21 photon counts per spatial position compared to 1200 counts in the similar “David” scene of [17] (with a 64 second exposure). In spite of these challenges, we demonstrate robust reconstructions and believe that our work is a significant step towards transient imaging at interactive

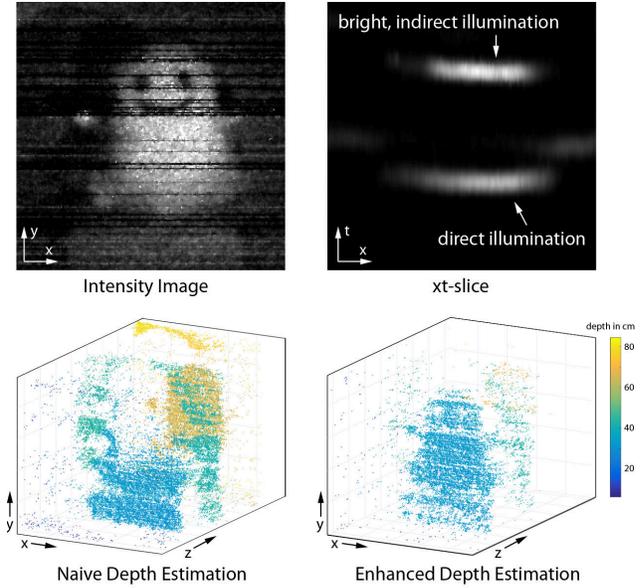


Figure 7. The “Ghost” scene showing a stuffed animal in front of a wall. The ghost is directly illuminated by a diffused laser pulse and also by bright indirect illumination that arrives at a later time than the direct light (top right). A simple log-matched filter falsely detects the indirect illumination as the pulse corresponding to the scene depth (lower left), thereby severely distorting the depth map. By separating direct and global illumination, we can apply the log-matched filter to the direct-only light transport to significantly improve the depth estimation (lower right).

frame rates.

Even though our system can acquire data at interactive rates, data processing (i.e., denoising and deconvolution) still requires a few minutes of compute time per transient imaging on a desktop computer with unoptimized MATLAB code. Hardware acceleration using a graphics processing unit (GPU), a field-programmable gate array (FPGA), or an application-specific integrated circuit (ASIC) would likely overcome this computational bottleneck. To make a specific application of transient imaging, such as depth estimation, practical, the algorithm may have to be implemented on a hardware accelerator anyway, which would require the algorithm to be reasonably simple. While more advanced algorithms than those presented in Section 3 may be easily devised, the discussed approaches may be simple enough to be practical.

Data-driven approaches to transient image reconstruction may also be useful in capturing difficult-to-model effects through learning. While such an approach may also lead to quicker reconstruction times compared to iterative optimization methods, a challenge is acquiring a well-constructed dataset, whether captured or simulated.

A limitation of the proposed approach to direct-global light transport separation is that it fails when a scene point

does not observe direct illumination, but only indirectly scattered light. Furthermore, the ability to capture subtle effects, such as subsurface scattering, is limited by the temporal resolution of the system (approx. 494 ps).

The primary challenge of any approach attempting to record transient images at fast acquisition rates is the low photon count. Even though it seems that higher-resolution, possibly two-dimensional SPAD arrays would naturally allow for transient videography, such a “flash LIDAR” approach also requires the laser to be diffused over the entire scene. This is also done for most of the scenes scanned with our scanned 1D SPAD array. Unfortunately, diffusing the laser illumination spreads its limited power over a large area and significantly reduces the number of photon events recorded at each sample point. Working at longer ranges makes this even more challenging due to the square distance intensity falloff of the light source. Using a higher-power laser could mitigate this limitation, but may cause safety concerns. These, in turn, could be addressed by operating in the infrared wavelength regime (e.g., 1550 nm), where it is eye-safe to use significantly higher-power lasers. Silicon SPADs may not be available for these wavelengths, however, because these are only sensitive up to the near-infrared regime where eye safety may be a concern. More advanced optical coding schemes [18, 16, 2] provide another alternative to optimize the light efficiency of transient imaging systems, but we leave such efforts for future work.

## Acknowledgements

D.B.L. is supported by a Stanford Graduate Fellowship in Science and Engineering. M.O. is supported by the Government of Canada through the Banting Postdoctoral Fellowships program. G.W. is supported by a National Science Foundation CAREER award (IIS 1553333), a Terman Faculty Fellowship, the DARPA REVEAL program, a Sloan Fellowship, and by the KAUST Office of Sponsored Research through the Visual Computing Center CCF grant.

## References

- [1] N. Abramson. Light-in-flight recording by holography. *Opt. Lett.*, 3(4):121–123, Oct 1978. 1, 2
- [2] S. Achar, J. R. Bartels, W. L. R. Whittaker, K. N. Kutulakos, and S. G. Narasimhan. Epipolar time-of-flight imaging. *ACM Trans. Graph. (SIGGRAPH)*, 36(4):37:1–37:8, 2017. 7
- [3] I. Bar-David. Communication under the Poisson regime. *IEEE Trans. Information Theory*, 15(1):31–37, 1969. 6
- [4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011. 3
- [5] S. Burri, H. Homulle, C. Bruschini, and E. Charbon. LinoSPAD: a time-resolved 256x1 CMOS SPAD line sensor system featuring 64 FPGA-based TDC channels running at up to 8.5 giga-events per second. *Proc. SPIE*, 9899:98990D, 2016. 2
- [6] L. Gao, J. Liang, C. Li, and L. V. Wang. Single-shot compressed ultrafast photography at one hundred billion frames per second. *Nature*, 516:7477, 2014. 1
- [7] G. Gariepy, N. Krstaji, R. Henderson, C. Li, R. R. Thomson, G. S. Buller, B. Heshmat, R. Raskar, J. Leach, and D. Faccio. Single-photon sensitive light-in-flight imaging. *Nature Communications*, 6(6021), Jan 2015. 1
- [8] G. Gariepy, F. Tonolini, R. Henderson, J. Leach, and D. Faccio. Detection and tracking of moving objects hidden from view. *Nature Photonics*, 10, Jan 2016. 1
- [9] I. Gkioulekas, A. Levin, F. Durand, and T. Zickler. Micron-scale light transport decomposition using interferometry. *ACM Trans. Graph. (SIGGRAPH)*, 34(4):37:1–37:14, 2015. 1, 2
- [10] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich. Low-budget transient imaging using photonic mixer devices. *ACM Trans. Graph. (SIGGRAPH)*, 32(4):45:1–45:10, 2013. 1, 2
- [11] F. Heide, L. Xiao, A. Kolb, M. B. Hullin, and W. Heidrich. Imaging in scattering media using correlation image sensors and sparse convolutional coding. *Opt. Express*, 22:26338–26350, 2014. 1
- [12] A. Jarabo, B. Masia, J. Marco, and D. Gutierrez. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics*, 1, 2017. 2
- [13] A. Kadambi, R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Trans. Graph. (SIGGRAPH Asia)*, 32(6):167, 2013. 1
- [14] A. Kirmani, D. Venkatraman, D. Shin, A. Colaco, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal. First-photon imaging. *Science*, 343(6166):58–61, 2014. 1
- [15] A. McCarthy, N. J. Krichel, N. R. Gemmell, X. Ren, M. G. Tanner, S. N. Dorenbos, V. Zwiller, R. H. Hadfield, and G. S. Buller. Kilometer-range, high resolution depth imaging via 1560 nm wavelength single-photon detection. *Opt. Express*, 21(7):8904–8915, 2013. 1
- [16] M. O’Toole, S. Achar, S. G. Narasimhan, and K. N. Kutulakos. Homogeneous codes for energy-efficient illumination and imaging. *ACM Trans. Graph. (SIGGRAPH)*, 34(4):35:1–35:13, 2015. 7
- [17] M. O’Toole, F. Heide, D. Lindell, K. Zang, S. Diamond, and G. Wetzstein. Reconstructing transient images from single-photon sensors. *CVPR*, 2017. 1, 2, 3, 6
- [18] M. O’Toole, F. Heide, L. Xiao, M. B. Hullin, W. Heidrich, and K. N. Kutulakos. Temporal frequency probing for 5D transient analysis of global light transport. *ACM Trans. Graph. (SIGGRAPH)*, 33(4):87:1–87:11, 2014. 7
- [19] M. O’Toole, D. Lindell, and G. Wetzstein. Confocal non-line-of-sight imaging based on the light cone transform. *Nature*, 2018. 1
- [20] B. Schwarz. LIDAR: Mapping the world in 3D. *Nat. Photonics*, 4(7):1749–4885, 2010. 1
- [21] D. Shin. Computational imaging with small numbers of photons. *Ph.D. Thesis*, 2016. 3

- [22] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. C. Wong, and J. H. Shapiro. Photon-efficient imaging with a single-photon camera. *Nature Communications*, 7(12046):1–7, 2016. [1](#), [3](#)
- [23] S. Shrestha, F. Heide, W. Heidrich, and G. Wetzstein. Computational imaging with multi-camera time-of-flight systems. *ACM Trans. Graph. (SIGGRAPH)*, 35(4):33:1–33:11, 2016. [1](#)
- [24] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3, 2012. [1](#)
- [25] A. Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar. Femto-photography: Capturing and visualizing the propagation of light. *ACM Trans. Graph. (SIGGRAPH)*, 32(4):44:1–44:8, 2013. [1](#), [2](#)
- [26] D. Wu, A. Velten, M. O’Toole, B. Masia, A. Agrawal, Q. Dai, and R. Raskar. Decomposing global light transport using time of flight imaging. *IJCV*, 107(2):123–138, 2014. [1](#), [5](#)
- [27] D. Wu, G. Wetzstein, C. Barsi, T. Willwacher, M. O’Toole, N. Naik, Q. Dai, K. Kutulakos, and R. Raskar. Frequency analysis of transient light transport with applications in bare sensor imaging. *ECCV*, 2012. [1](#)