

# The Dawn of Computational Light Transport

IT IS EASY TO FORGET, when casually observing our surroundings, that the speed of light is finite. Light travels so quickly that even though it may scatter, refract or bounce many times on nearby surfaces before reaching our eyes, these events are spaced just trillionths of a second apart—too fast for any conventional camera to resolve, and certainly too fast for our own visual system to perceive as anything but instantaneous.

But are these individual light transport events really beyond the realm of direct visual observation? What would the world look like if we had a chance to observe it with a trillion-frame-per-second video camera? And what insights might one gain from such observations?

The following paper by Velten *et al.* represents an audacious attempt to answer these questions for the very first time for general objects. Much like a microscope can zoom to a tiny area on a specimen, the authors describe an imaging system that can zoom to a tiny interval of time—just a nanosecond across—and record a 512-frame video that spans it. Their online videos show wavefronts of light travelling across a variety of natural and man-made objects and are simply stunning (<http://web.media.mit.edu/~raskar/trillionfps/>, <http://giga.cps.unizar.es/~ajarabo/pubs/femtoSIG2013/>).

The authors' approach was to take a page out of Doc Edgerton's playbook. In the 1920s Edgerton, inventor of the electronic flash and famous for his photos of speeding bullets frozen in time, used a stroboscope to inspect rapidly-rotating motors. By flashing a bright light very briefly in sync with the motor's rotation, he could make the motor look like it was standing still (and therefore easy to photograph). Here the authors use the same technique at much,

much shorter timescales to make light itself appear still. How short? The flash lasts one twenty-thousandth of a nanosecond; a frame of video is exposed for about forty times as long in sync with the flash; and this is repeated half a billion times at intervals of about thirteen nanoseconds in order to collect enough light. Simple!

The devil, of course, is in the details. Such brief flashes can only come from a femtosecond laser. Also, cameras with picosecond-scale frame rates do not actually exist. The closest sensor the authors could find is a streak camera—a 1D photodetector used primarily in chemistry and physics to take x-versus-time light measurements at picosecond resolutions. To capture 2D video, the sensor is panned across the field of view and makes repeated acquisitions, building the video row by row. Additional processing then ensures the final video is consistent with the arrow of time.

As with any "first" attempt to solve a problem, it is relatively easy to criticize. The equipment costs hundreds of thousands of dollars; it takes hours to record one video; the scene must remain perfectly still; and its surface geometry must be captured laboriously with a 3D digitizer.

These issues, however, are beside the point. The paper's principal contribution is the highest-order bit: that light transport through an everyday scene *can* be observed directly by ultra-fast imaging. With that settled, cheaper and faster solutions will surely follow (and they have—see <http://www.cs.ubc.ca/labs/imager/tr/2013/TransientPMD/>).

On a broader level, the paper marks the maturation of *computational light transport*—an emerging family of techniques that use lights, sensors and computation to reveal and interpret the flow of light in our everyday world. The topic has been

brewing for well over two decades in the research communities of computer vision and computer graphics but required the confluence of several factors to finally come into its own. This includes new sensors; programmable light sources; powerful computing engines; new algorithmic tools for data-driven analysis of light transport; and growing interdisciplinary ties to optics.

Will these techniques ultimately prove useful? There are several reasons to be optimistic. As a person living in Canada, I can say that there is a great need for imaging techniques that work well in the presence of complex light transport, and scattering in particular (fog, snow, *etc.*) Scattering is also *the* barrier to imaging deep under the skin and through other biological tissues and causes major accuracy problems in modern depth cameras and 3D scanners. While I don't expect vision-based driving on snowy roads and through raging snowstorms anytime soon, transport-robust 3D imaging techniques are already here (see DOIs 10.1145/2735702 and 10.1145/2766897), never-seen-before abilities such as looking around corners are a reality (see DOI 10.1038/ncomms1747), and there is great potential for synergies with scientific imaging (see DOIs 10.1038/NPHYS3373 and 10.1145/2766928).

But perhaps the biggest reason for optimism is this: ultimately, computer vision algorithms can only be as powerful as the sensor data they take as input. Given the tremendous strides the field has made recently in tasks such as object recognition and navigation from conventional images, one can only speculate how much more could be accomplished from light transport data. So, read the paper and consider the possibilities.