Understanding Video at 30 Billion Frames Per Second with *Transient Rendering*

Adam Smith James Skorupski James Davis
{ amsmith, jskorups, davis } @ cs.ucsc.edu

Ultra-Fast Video

•Imagine this: *a photon of light moves one centimeter during each frame*
•Ultra-fast videos can *time-resolve* the distinct paths light takes
•Scenes can be probed with ultra-short flashes

Transient Rendering

An ultra-fast video is nonsense without a way to relate it to a scene. Transient rendering provides expected observations by generalizing the traditional rendering equation to account for light propagation delays.

Inputs
•Surface geometry
•Scattering properties (reflectance/transmittance kernels)
•Initial flux functions (known emission functions)
Process
•Repeat and sum for all surfaces:
  •Propagate last radiant flux (including distance-based time delays)
  •Scatter last irradiant flux (no delay)
Output
•Instantaneous total flux field: $F([x,y,z],[\vartheta,\varphi,\rho],t)$
•Restrictions as time-pixels: $I(t) = F(v,\omega,t)$

Ultra-Fast Video Understanding

Understanding ultra-fast videos will require explaining *emergent* patterns in space and time in terms of *localized* elements of a scene. With transient rendering we can follow an analysis-by-synthesis approach.

Rationale

•Global transient light transport is the operator composition propagation and scattering, both are linear, time-invariant processes – invertible processes!

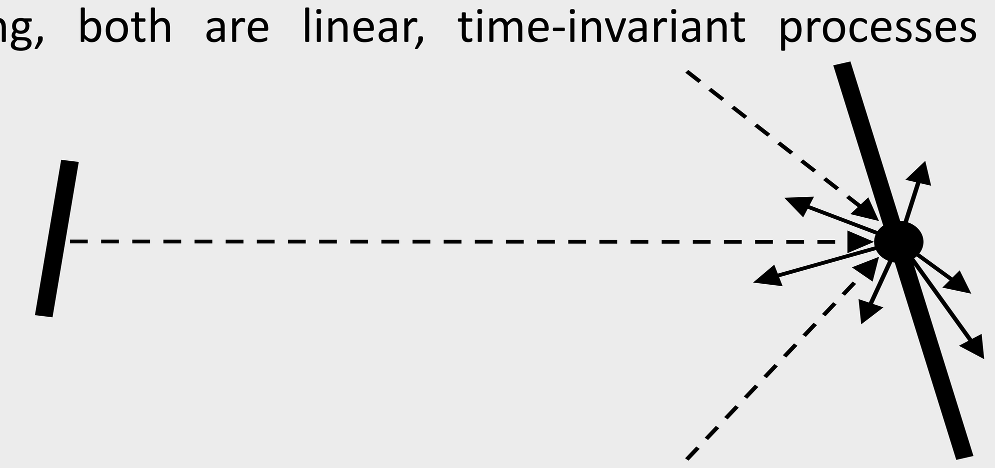

Hollywood Fantasy

courtesy Warner Bros.

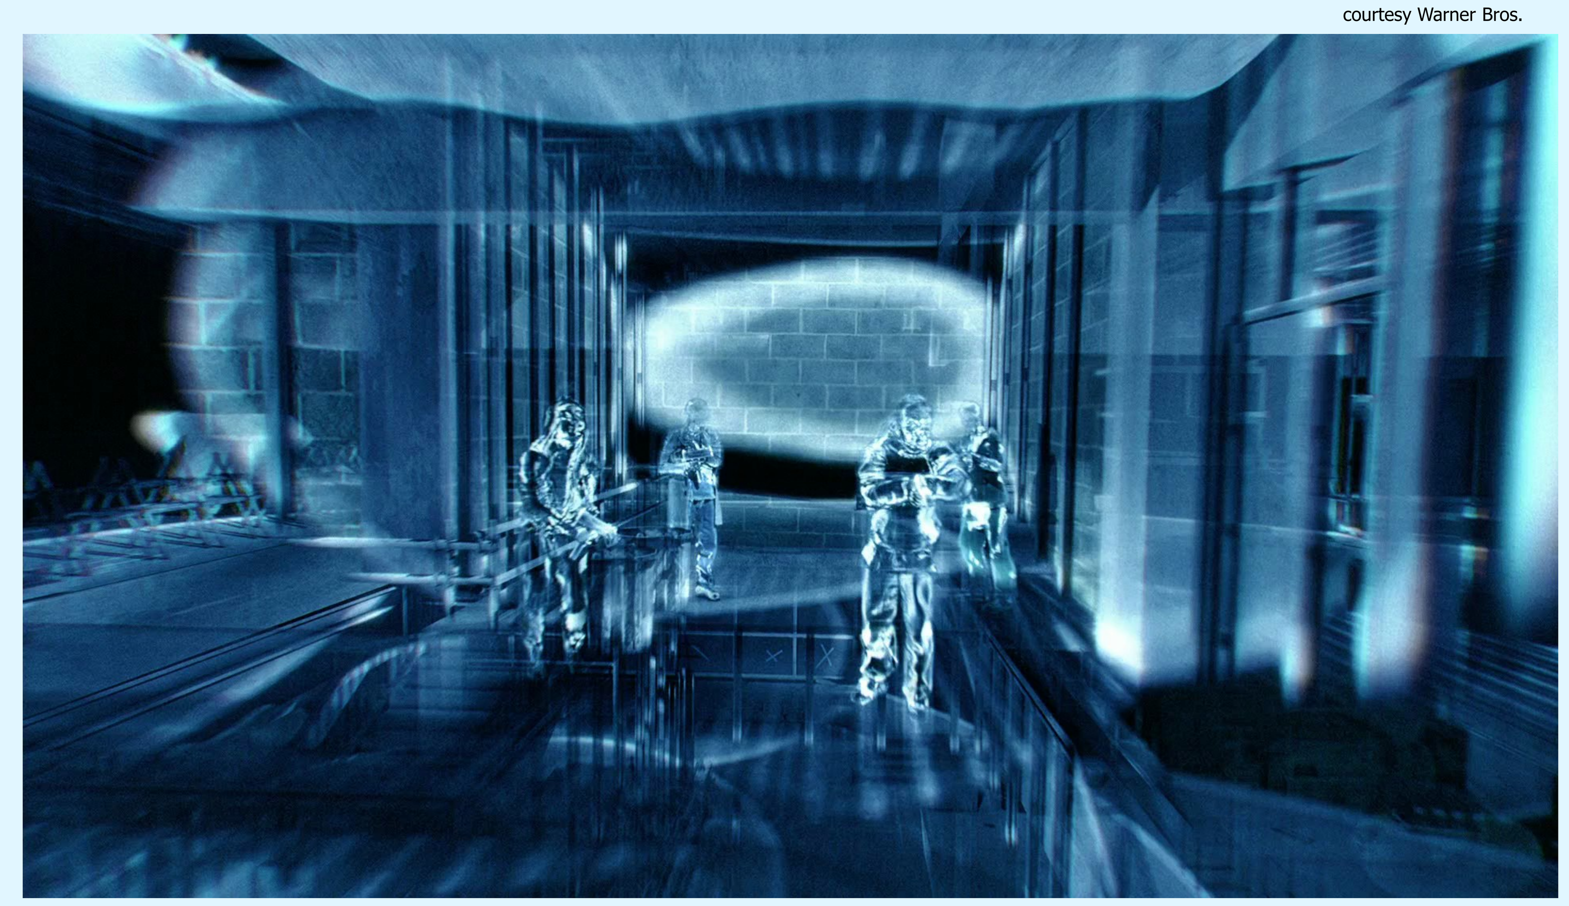

In *The Dark Knight*, Batman rigs a Gotham City's cell phones to use their speakers and microphones as a distributed sonar network, building a real-time 3D map of the city. Pshh – unrealistic! The only way Batman can avoid the rampant diffraction of sound waves is to use *light pulses* and *light sensors*. Let's fight crime with a new theoretical framework. :)

Baby Steps

Synthesis
a) a simple 1D world with an eye, light source, and two surfaces (one translucent and one opaque)
b) transient rendering
c) a transient photometric response function with distinct peaks

Analysis
Inter-surface distances can be read off from delays between distinct peaks, but, for geometry beyond point elements, integration quickly smears the TPRF in time requiring robust signal processing methods to recover geometry.

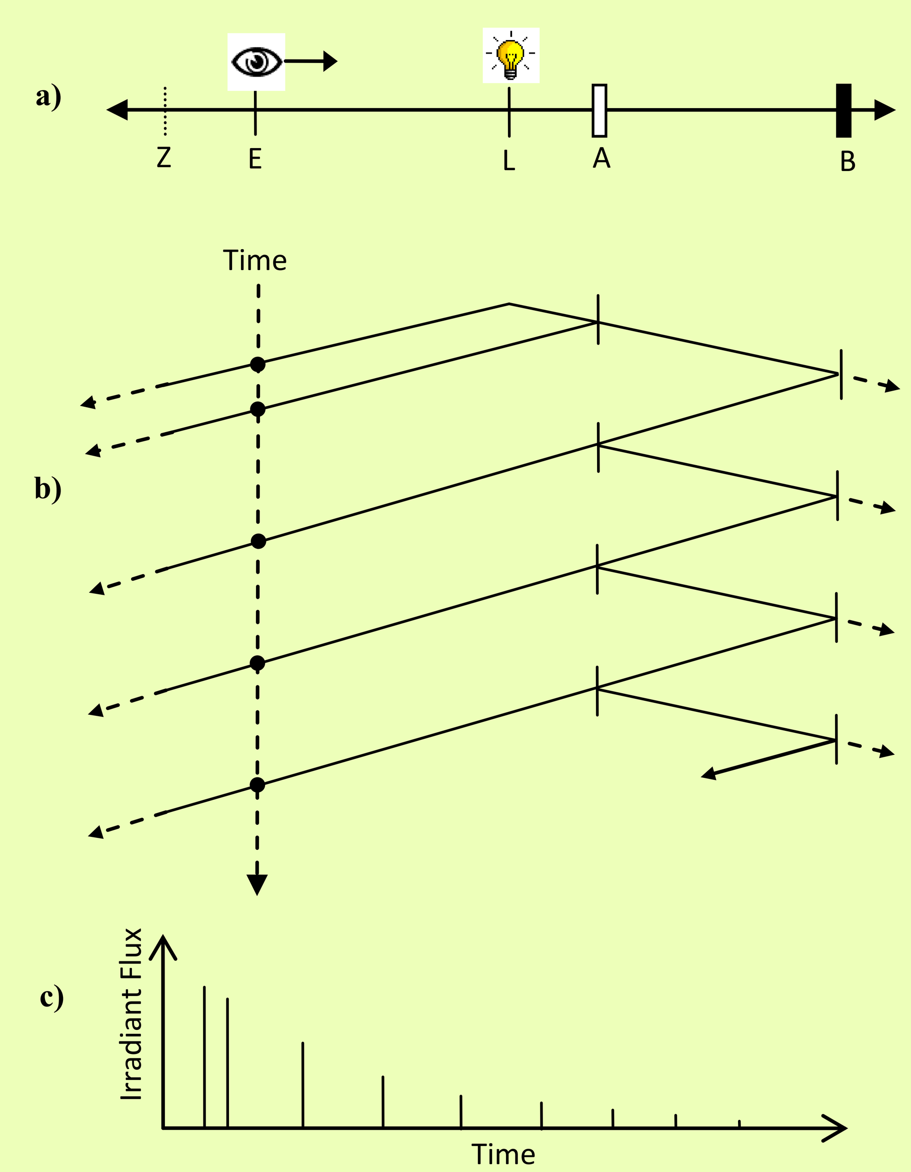

Potential Applications

Hidden Surface Recovery
•Recover geometric details from "around the corner" by observing delays for light that bounced there and then bounced back

Localizing Specular Objects
•Analyze light paths to infer the position of mirror-smooth objects

Advanced Single-viewpoint Surface Modeling
•Fit complex scattering models to known surfaces without moving the camera or assuming uniformity (BSDF, sub-surface scattering, etc.)

Fundamental Challenges

•Ultra-short pulse sources are difficult to form
•Very little energy is available for capture per-pixel-per-frame
•Transient rendering is incompatible with many existing global-illumination techniques (path-tracing is ok, radiosity and photon mapping are out)
•General recovery might require repeated layers of deconvolution

How do I use this?

Reasoning over a small set of distinct paths (as in traditional LIDAR) won't scale to complex scenes. Instead, fitting parameters to data with transient rendering is a yet another non-linear optimization problem with a rich background model to provide informative gradients.

How realistic is this?

•A probing pulse can be created using femtosecond laser
•Commodity hardware setups can sample at 5Ghz
•Initial "around the corner" tests yield <1cm error in positional estimates

See also
Kirmani, A., Hutchinson, T., Davis, J., Raskar, R. *Looking Around the Corner Using Transient Imaging*. ICCV, 2009.
Smith, A., Skorupski, J., Davis, J. *Transient Rendering*. Technical Report, UCSC-SOE-08-26, 2008.
Ramamoorthi, R., Hanrahan, P. *A Signal-Processing Framework for Inverse Rendering*. SIGGRAPH '06, 2006.