

Relativistic Effects for Time-Resolved Light Transport

Adrian Jarabo¹ Belen Masia^{1,2,3} Andreas Velten⁴ Christopher Barsi² Ramesh Raskar² Diego Gutierrez¹

¹Universidad de Zaragoza

²MIT Media Lab

³I3A Institute

⁴Morgridge Institute for Research

Abstract

We present a real-time framework which allows interactive visualization of relativistic effects for time-resolved light transport. We leverage data from two different sources: real-world data acquired with an effective exposure time of less than 2 picoseconds, using an ultrafast imaging technique termed femto-photography, and a transient renderer based on ray-tracing. We explore the effects of time dilation, light aberration, frequency shift and radiance accumulation by modifying existing models of these relativistic effects to take into account the time-resolved nature of light propagation. Unlike previous works, we do not impose limiting constraints in the visualization, allowing the virtual camera to explore freely a reconstructed 3D scene depicting dynamic illumination. Moreover, we consider not only linear motion, but also acceleration and rotation of the camera. We further introduce, for the first time, a pinhole camera model into our relativistic rendering framework, and account for subsequent changes in focal length and field of view as the camera moves through the scene.

1. Introduction

Analyzing and synthesizing light transport is a core research topic in computer graphics, computer vision and scientific imaging [GNJJ08]. One of the most common simplifications, rarely challenged, is the assumption that the speed of light is infinite. While this is a valid assumption in most cases, it is certainly not true: light travels extremely fast, but with finite speed. In this paper, we lift this assumption and explore the consequences of dealing with time-resolved data (finite speed of light), focusing on the effects predicted by special relativity when the camera moves at speeds comparable with the speed of light.

Transient imaging has recently emerged as a vibrant, exciting area of research. Being able to analyze light transport at picosecond scale has already helped gain a better understanding of the complexities of light propagation [VWJ*12, VWJ*13], to approximate the shape of hidden objects [VWG*12] or reflectance properties [NZV*11, NBVR14]. In this paper, we offer a novel contribution by visualizing relativistic effects of time-varying radiance. Beyond the pure scientific interest of advancing the field of relativistic visualization, our work has direct applications in education and even games (see for instance *OpenRelativity* from the MIT Game Lab [KTS13]). Additionally, it can also contribute to recent works which have set the ground for a time-resolved framework for light transport simulation [JMM*14].

Relativistic rendering is not new [HD89, CLC96, WBE*06]. However, our particular time-resolved framework implies by definition that radiance is not constant in the temporal domain. In particular, we need to address how time affects light aberration, the Doppler effect and the searchlight effect typically handled in relativistic rendering, and include time dilation effects due to Lorentz contraction.

Additionally, we modify the traditional camera model used in previous literature and introduce for the first time—to our knowledge—a pinhole camera model in a relativistic rendering framework. This implies the consideration of two reference frames in the camera, co-moving with the pinhole and with the sensor respectively, and thus a focal length (distance pinhole-sensor) which can vary in certain reference frames under certain conditions; this, in turn, has an effect on the field of view of the camera, affecting the recorded imagery. We describe the required formulation for this camera model both in the presence of constant velocity and accelerated motion. Our model of relativistic acceleration is based on the clock hypothesis; we assume instantaneously co-moving frames of reference, thus being able to apply special relativity [Dew63, Wei01]. Besides, we develop a model that accounts for relativistic rotation, allowing to freely rotate the camera.

We demonstrate our relativistic rendering framework with fly-throughs in a range of scenes. We rely on both captured data (using *femto-photography* [VWJ*13], which al-

lows capturing light propagation at effective exposure times of less than 2 ps per frame), and synthesized data (using the recent transient rendering framework by Jarabo et al. [JMG13, JMM*14]).

In summary, we have developed a rendering and visualization tool for transient light transport, capable of simulating relativistic effects, lifting some of the assumptions and constraints of previous works. This paper is an extension of our previous work on rendering of time-resolved data [JMV*13], in which we showed results of relativistic rendering of time-resolved data, including camera rotations. Here, we offer a more detailed formulation, significantly extend the discussion on previous works, and develop a novel framework to be able to handle a pinhole camera model with a finite focal length; this has an effect on image formation, since, as we will show, the focal length can suffer contractions and dilations in the sensor frame of reference. Our contributions can be summarized as follows:

- We describe a special relativity rendering framework which is able to handle time-resolved data
- We extend the traditional camera model used in the literature to a pinhole camera model, and describe the necessary transformations for rendering time-resolved scenes under this new model, both in the case of constant velocity and accelerated motion
- We propose an approximate solution for the case of relativistic *rotation*, including relativistic deformation of the sensor, so the camera can be moved freely in 3D space

2. Related Work

Time Resolved Light Transport A modified rendering equation can account for the finite speed of light and handle transient effects [Arv93, SSD08]. However, in previous works no practical rendering framework is derived from the proposed transient rendering framework. A fully functional time-resolved rendering system was recently presented by Jarabo and colleagues [JMG13, JMM*14]. In their work, they introduce the transient path integral framework, and devise new sampling strategies and density estimation techniques for efficient rendering. Wu et al. [WWB*13] perform a rigorous analysis on transient light transport in frequency domain. They derive an analytic expression that models the information transfer between dimensions, and show that this derivation can be exploited to build a bare-sensor imaging system.

Time-resolved imaging is further analyzed by Wu et al. [WVO*13] to separate direct and global illumination components in macroscopic table-top scenes. The authors analyze the time profile for each pixel and decompose it into direct, subsurface scattering and interreflection components. Kirmani et al. [KHDR09] utilized global information in time-of-flight images to infer geometries of hidden objects, not directly visible by the camera, by using the encoded

time-of-flight of diffuse indirect illumination. This work was further improved by Velten et al. [VWG*12]. Material BRDFs of distant patches were reconstructed [NZV*11] via light-bounce analysis from ultrafast image measurements.

Very recently, Velten et al. [VWJ*12, VWJ*13] developed femto-photography, a novel technique that allows ultra-fast (in the order of picoseconds) capture of transient light transport, at an effective frequency of a trillion frames per second. We explain this system in more detail in Section 3, since we rely on the data it provides to render some of the relativistic effects shown in this paper.

Relativistic Rendering Here we discuss the most relevant work on relativistic rendering. For a wider survey, we refer to [WBE*06], where the different proposed techniques for both general and special relativistic rendering are discussed, including their application as educational tools. A number of early works by Hsiung et al. [HD89] were the first to tackle rendering of relativistic effects in the field of computer graphics and visualization, applying the Lorentz transform to render scenes in which the observer moves at relativistic speeds relative to the scenes. A set of works followed, extending this work to three dimensional motion, considering the Doppler shift and time dilation, and improving rendering times [HTCD90, HT90, HTW90, HTC*90]. Later, Chang et al. [CLC96] developed a more comprehensive framework accounting for both geometric and radiance transformations, including light aberration (Lorentz transform), the Doppler effect, and the searchlight effect. However, their formulation modeled the searchlight and Doppler effects incorrectly; these were corrected by Weiskopf et al. [WKR99], who later further described the rendering pipeline in [Wei00a]. Weiskopf was also the first to formalize a complete rendering pipeline which could deal with special relativity and considered not only moving observers but also dynamic scenes, including acceleration, and moving light sources [Wei01]. This pipeline, with an application to a virtual reality environment, is described in [Wei00b], while special relativity renderings of accelerated extended objects are considered in [WBE*06, KRWZ02][†]. Acceleration is further dealt with [WKR00b], correcting earlier work by Betts [Bet98]. Related work simulates relativistic effects in synthetic scenes modeled with image-based techniques, by applying the relativistic transformations directly on the light field [WKR00a]. Finally, visualization approaches and games have been created with a didactic goal, aiming at helping students in the understanding of relativity. Doat et al. [DPV11] developed a billiard game, with a relativistic rendering engine and a non-Newtonian physics engine to deal with the collisions, to help students visualize the effects of special relativity. The game *A Slower Speed of Light*, also notable among these, uses the open-source toolkit *OpenRel-*

[†] English version available at <http://www.spacetime.travel/tompkins/tompkins.html>




Figure 1: Scene setups and selected frames for the four scenes used throughout the paper. The scenes on the left (cube and tomato) depict real data captured using femto-photography [VWJ*13], while the scenes on the right (bunny and dragons) have been synthesized using time-resolved rendering [JMM*14].

activity which works with the *Unity* engine and can also simulate special relativity effects [KTS13]. In this paper we develop a relativistic rendering framework including a pinhole camera model, with a finite focal length; this has an effect on image formation, since, as we will show, this focal length can suffer contractions and dilations in the sensor frame of reference. Further, we show results of relativistic rendering of time-resolved data acquired from the real world.

3. Time-resolved Data

In this paper we leverage two different sources of time-resolved light transport data: real data from the femto-photography technique of Velten et al. [VWJ*13], and synthesized data using the transient rendering framework of Jarabo et al. [JMM*14].

Femto-photography [VWJ*13] is a novel imaging technique with an effective exposure time of 1.85 picoseconds, which allows to image the *propagation* of light as it interacts with objects in a scene. A *Ti:Sapphire* femto-second laser pulse is repeatedly shot against a diffuser, which reflects it into the scene. Photons enter the camera through a horizontal slit, and are converted into electrons and then deflected vertically by a time-varying voltage. This allows to effectively code the time of arrival of each photon as a spatial coordinate in the sensor ‡. Figure 1 (left) shows the setups and selected frames for the two captured scenes used in this work: *cube* and *tomato*.

On the other hand, Jarabo et al. [JMM*14] extend the path integral framework to formally describe light transport in transient state, and propose several techniques to efficiently compute transient images. In particular, the authors propose a novel density estimation technique to reuse sampled paths to reconstruct time-resolved radiance, as well as new sampling strategies according to the distribution of radiance *along time* in participating media. Figure 1 (right)

shows selected frames for the synthesized *bunny* and *dragons* scenes: both of them show a direct spherical wavefront illuminating the scene, creating secondary wavefronts over time due to indirect illumination.

4. Relativistic Rendering

Time-resolved data allows us to explore light transport like never before, no longer being constrained by the assumption that light speed is infinite. However, current capture systems only allow to visualize light propagation from the capture point of view. Assuming that the geometry in the scene is known (which can be easily acquired with a digitizer arm or from time-of-flight data), we can synthesize new viewpoints and animations of the scene by taking an image-based rendering approach, using x-y textures from the x-y-t data cube (see [VWJ*13] and Section 3) and projecting them onto the geometry. This allows us to visualize real-world events from new angles.

4.1. Time-warping

Visualizing light transport events at pico-second scale yields counter-intuitive results, as observed by Velten et al. [VWJ*13]. Due to the finite speed of light, events are not captured in the sensor as they occur, which leads to unexpected apparent distortions in the propagation of light. From this observation, it follows that different temporal frames of reference must be employed: one for the world O (when the events occur), and one for the camera sensor O' (when the events are actually captured).

As a consequence, sensor data acquired by the femto-photography technique appears *warped* in the temporal domain, and must be *time-unwarped* to take into account the finite speed of light. So for each frame in the synthesized animations, we access the original warped data and apply the following transformation [VWJ*13]:

$$t'_{ij} = t_{ij} + \frac{z_{ij}}{c/\eta} \quad (1)$$

‡ Videos and data from scenes captured with this setup can be found online at: <http://giga.cps.unizar.es/~ajarabo/pubs/femtoSIG2013/>

where t'_{ij} and t_{ij} are sensor (camera) and world times respectively, z_{ij} is the depth from each point (i, j) to the new camera position, c is the speed of light and η the index of refraction of the medium. Note how a naive approach based on simply sticking the textures from the first frame to the geometry through the animation would produce wrong results; the distance from each geometry point to the center of projection of the camera varies for each frame, and thus a *different* transformation must be applied each time to the original, *warped* x-y-t data (see Figure 2).




Figure 2: Time unwarping between camera time and world time for synthesized new views of a cube scene. Top row, left: Scene rendered from a novel view keeping the unwarped camera time from the first frame (the small inset shows the original viewpoint). Right: The same view, warping data according to the new camera position. Notice the large changes in light propagation, in particular the wavefronts on the floor not visible in the previous image. Bottom row: Isochrones visualization of the cube-scene for a given virtual camera (color encodes time); from left to right: original x-y-t volume in the time-frame of the capturing camera, unwarped x-y-t data in world time frame, and re-warped data for the new virtual camera.

4.2. Camera Model

We model our camera as a pinhole camera, with a focal distance d from the sensor to the pinhole aperture, and a field of view (FOV) $\alpha = 2 \arctan(w/d)$, with w the half width of the sensor (Figure 3, left). We assume the camera has no lenses, the pinhole being what ensures correct focusing on the sensor. This model effectively splits the camera into two frames of reference: one frame O' which will be co-moving with the sensor, and one for the pinhole aperture O'' . Note that considering the camera as a finite, instead of a point-like, object, has important effects on the recorded imagery due to relativistic contractions and dilations of the focal length, which will need to be taken into account. We detail how this is done in the following sections.




Figure 3: Left: Pinhole camera model, with P marking the position of the pinhole, w the half-width of the sensor, and d the focal length. Right: World lines for the sensor and the pinhole, moving with constant speed $\beta = 0.2$. The green dotted axes depict a local frame of reference of the sensor; these axes show how, in the frame of the sensor, the focal length d' is larger than when measured in the world frame, d . Note that the t axis actually corresponds to ct , but only t is stated in this and subsequent figures for clarity.

4.3. Relativistic Effects

In addition to the time-warping of data, which occurs in both static and dynamic cameras, macroscopic camera movement at pico-second time scales, like the one synthesized in Figure 2, would give rise to relativistic effects. This requires a relativistic framework to correctly represent and visualize light traveling through the 3D scene.

According to special relativity, light aberration, the Doppler effect, and the searchlight effect need to be taken into account when simulating motion at fast speeds. *Light aberration* accounts for the apparent geometry deformation caused by two space-time events measured in two reference frames moving at relativistic speeds with respect to each other. The *Doppler effect* produces a wavelength shift given by the Doppler factor. Last, the *searchlight effect* increases or decreases radiance, according to whether the observer is approaching or moving away from a scene.

We build our relativistic visualization framework on the derivations by Weiskopf et al. [WKR99]. We consider two inertial frames, O and O' , where O' (the sensor) is moving with velocity $v = \beta c$ with respect to O , with $\beta \in [0.. \pm 1)$. L represents wavelength-dependent radiance[§] measured in O , defined by direction (θ, ϕ) and wavelength λ . θ and ϕ are defined with respect to the direction of motion, with $\theta \in [0.. \pi]$ being the angle between the view and the motion vector, and $\phi \in [0.. 2\pi]$ the angular coordinate of the view vector around the motion vector on its tangent plane. The corresponding

[§] For legibility we use the term "radiance" to describe "wavelength-dependent radiance" through the text.

primed variables (θ', ϕ') and λ' define radiance L' measured in O' . To obtain the modified radiance L' given L and the speed of the sensor, we need to apply the following equation [WKR99]:

$$L'(\theta', \phi', \lambda') = D^{-5} L\left(\arccos \frac{\cos \theta' + \beta}{1 + \beta \cos \theta'}, \phi', \frac{\lambda'}{D}\right), \quad (2)$$

where $D = \gamma(1 + \beta \cos \theta')$ and $\gamma = 1/\sqrt{1 - \beta^2}$ are the Doppler and Lorentz factors, respectively. This equation accounts for all three factors: light aberration, the Doppler effect, and the searchlight effect. However, it cannot model explicitly the effect of special relativity on non-constant radiance. In the following paragraphs we explain each effect separately, and discuss the modifications needed to handle time-resolved radiance.

Time dilation: Breaking the assumption of constant radiance means that we cannot ignore the effect of time dilation [Ein61]. Time dilation establishes that, given an observer moving with respect to a stationary (world) frame O , if Δt_0 is the elapsed time in the rest frame of the moving observer, the time elapsed for an observer in the stationary frame Δt , and the time elapsed for the moving observer are related by $\Delta t = \gamma \Delta t_0$. Since $\gamma \geq 1$, this effectively means that a clock carried by the moving observer will tick more slowly than a clock carried by the stationary observer. Consequently, if we take as the proper time the time in the frame O' in which the camera is at rest, the relationship between time intervals in the camera frame O' and the world frame O is $\Delta t = \gamma \Delta t'$. Thus, we need to keep track of both world t and camera time t' , since they differ depending on the speed of motion.

In addition to time dilation, Lorentz contraction shrinks the space in the direction of motion. This means that the time warping modeled with Equation 1, which accounts for the time it takes light to travel to the camera, has to be computed in the moving frame O' . Establishing the proper length for z in the world frame, the transformation to obtain z' is given by:

$$z' = z \sqrt{\sin^2 \theta + \frac{\cos^2 \theta}{\gamma^2}}. \quad (3)$$

Camera deformation: In addition to the effects on propagation delays, the Lorentz contraction also affects the focal length d of our pinhole camera (see Section 4.2). We define the proper length of the camera in its own rest frame O' , thus being d' the proper focal length. Therefore, the focal length in the world frame O is $d = d'/\gamma$, as can be seen in the space-time diagram in Figure 3 (right). This effectively produces a reduction in the field of view of the camera α' , such that:

$$\alpha' = 2 \arctan\left(\frac{w}{\gamma d}\right) = 2 \arctan\left(\frac{\tan(\alpha/2)}{\gamma}\right). \quad (4)$$

Light aberration: An easy example to understand light

aberration is to visualize how we see rain drops when traveling on a speeding train. When the train is not moving, raindrops fall vertically; but as the train picks up speed, raindrop trajectories become increasingly diagonal as a function of the train's speed. This is because the speed of the train is comparable with the speed of raindrops. A similar phenomenon occurs with light if moving at relativistic speeds. However, as opposed to rain drops, relativistic light aberration cannot be modeled with classical physics aberration; the Lorentz transformation needs to be applied on top of it.

Light aberration is computed by transforming θ' and ϕ' with the following equations, which provide the geometric transformation between two space-time events measured in two reference frames which move at relativistic speeds with respect to each other:

$$\cos \theta' = \frac{\cos \theta - \beta}{1 - \beta \cos \theta} \quad (5)$$

$$\phi' = \phi. \quad (6)$$

The end result is that light rays are tilted towards the direction of motion as velocity increases, causing light rays reaching the sensor from behind the camera to become visible. Finally, as β approaches 1, and thus $v \approx c$, most incoming light rays are compressed towards the motion direction; this makes the scene collapse into a single point as the camera moves towards it (note that this produces the wrong impression that the camera is moving *away* from the scene). The first two rows in Figure 4 show the effects of light aberration with increasing velocity as the sensor moves at relativistic speeds, towards and away from the scene respectively.

Doppler effect: The Doppler effect is better known for sound, and it is not a phenomenon restricted to relativistic velocities. In our case, the Doppler effect alters the observed frequency of the captured events in the world when seen by a fast-moving camera, which produces a wavelength shift, as defined by the Doppler factor D :

$$\lambda' = D\lambda. \quad (7)$$

The overall result is a color shift as a function of the velocity of the sensor relative to the scene. Somewhat less known, the Doppler effect also creates a perceived speed-up (or down, depending on the direction of camera motion) of the captured events. This means that the *original frame rate* of the time-varying radiance f in world frame is Doppler shifted, making the perceived frame rate f' in camera frame become $f' = f/D$. Figure 4 (third row) shows an example of the Doppler effect.

Searchlight effect: Due to the searchlight effect, photons from several instants are captured at the same time differential, in part as a cause of the Doppler shift on the camera's perceived frame rate. This results in increased (if the observer is approaching the scene) or decreased (if the observer is moving away) brightness (see Figure 4, bottom row):

$$L'(\theta', \phi', \lambda') = D^{-5} L(\theta, \phi, \lambda). \quad (8)$$




Figure 4: Relativistic effects shown separately for the cube scene. First rows: Distortion due to light aberration as the camera moves towards and away from the scene at different velocities, with $\beta = \{0, 0.3, 0.6, 0.9, 0.99\}$. We assume a laser wavelength of 670 nm for visualization purposes. Note how in both cases light aberration produces counter-intuitive results as the camera appears to be moving in the opposite direction. Third row: Doppler effect, showing the shift in color as a consequence of the frequency shift of light reaching the sensor, with $\beta = \{0, 0.15, 0.25, 0.35, 0.50, 0.55\}$. Fourth row: Searchlight effect, resulting in an apparent increase in brightness as the speed of the approaching camera increases, with $\beta = \{0, 0.2, 0.3, 0.4, 0.5\}$ (simulated laser at 508 nm). All images have been tone-mapped to avoid saturation.

Intuitively, continuing with our previous rain analogy, it is similar to what occurs in a vehicle driving in the rain: the front windshield will accumulate more water than the rear windshield. Equation 8 has an implicit dependence on time; for our time-varying streak-data, we need to make it explicit, yielding:

$$L'(\theta', \phi', \lambda', t') = D^{-5} L(\theta, \phi, \lambda, t). \quad (9)$$

In our case, since in the end we need to obtain the radiance accumulated over the duration of each frame of the movie, we are interested in computing the incoming radiance impinging the camera during a certain camera time interval $\mathcal{T}' = 1/f'$, where f' is the frame rate in the camera frame. As a consequence of the Doppler shift in the frame rate mentioned in the previous paragraph ($f' = f/D$), this camera time per frame \mathcal{T}' corresponds to a time interval $\mathcal{T} = \mathcal{T}'/D$ in the world frame O . To do this, we need to integrate the wavelength-dependent radiance of Equation 9 over that time

interval, resulting in the following:

$$\int_{\mathcal{T}'} L'(\theta', \phi', \lambda', t') dt' = \int_{\frac{\mathcal{T}'}{D}} D^{-5} L(\theta, \phi, \lambda, t) dt'. \quad (10)$$

When incorporating the previously described effects into Equation 10, we obtain the final expression for incoming radiance in the frame of the camera for each camera time frame (which has a duration \mathcal{T}'):

$$\int_{\mathcal{T}'} L'(\theta', \phi', \lambda', t') dt' = D^{-5} \int_{\frac{\mathcal{T}'}{D}} L\left(\arccos \frac{\cos \theta' + \beta}{1 + \beta \cos \theta'}, \phi', \frac{\lambda'}{D}, \frac{t'}{\gamma} - \frac{z'}{c/\eta}\right) dt'. \quad (11)$$

Note that the D^{-5} factor used is valid only for a velocity vector v parallel to the sensor normal (see [WKR99] for a de-




Figure 5: Relativistic phenomena for the cube (top, captured data) and bunny (bottom, simulated data) scenes including light aberration, Doppler effect and the searchlight effect, as the camera approaches the scene at increasing velocities $v = \beta c$, with β increasing from 0 to 0.77 (cube), and from 0.2 to 0.9 (bunny). In both scenes, we compare between fixing the focal distance d' in the sensor frame O' (left), and fixing it in the world frame O , and therefore being transformed in the moving frame (right). Note that we transform RGB radiance into luminance.

tailed derivation); for other motion configurations that factor should be re-derived.

Finally, Figure 5 shows the result of combining all these relativistic effects, both for the *cube* scene (data captured with femto-photography techniques) and the *bunny* scene (simulated data by rendering) respectively, both fixing the camera focal length in the world frame (i.e. varying d' with speed, as shown in Figure 3, right), and fixing it in the sensor frame. The laser wavelength is set at 670 nm for visualization purposes. We refer the reader to the supplementary videos to see the full animations.

4.4. Relativistic Acceleration

Until now, we have assumed that the camera accelerates instantaneously, and therefore no effect is produced by such acceleration. Similar to previous works [Wei01, Dew63, Sty07], our framework for relativistic acceleration assumes that the clock hypothesis is valid. This hypothesis is implicit in the original formulation of special relativity, and states that the rate of a clock is only dependent on its *instantaneous* velocity. Therefore, an accelerating frame suffers the same relativistic effects as in its instantaneous co-moving inertial frame (i.e. a frame uniformly moving with the same instantaneous speed) [Wei01], modeled using Equation 11. It is important to note that the clock hypothesis is not universally accepted [Bri70, FG10]; breaking the assumption of its validity makes the formulation of relativistic acceleration significantly more complex.

With this assumption, we model the acceleration a in the

world frame O as [Ter59]:

$$a' = \frac{dv'}{dt} = \frac{d}{dt} \gamma v = \gamma^3 a, \quad (12)$$

where a' is the constant proper acceleration in the moving frame; following the standard notation in physics literature, in the following we call it g . By integration we can get the velocity in the world frame as a function of time t [Sty07]:

$$v(t) = \frac{g(t-t_0)}{\sqrt{1+(g(t-t_0)/c)^2}}, \quad (13)$$

where t_0 is the time at which acceleration starts ($v(t_0) = 0$). Integrating $v(t)$ allows us to obtain the position of the accelerating object in the world frame, $x(t)$, as:

$$x(t) = \frac{c^2}{g} \sqrt{1 + \left(\frac{g(t-t_0)}{c}\right)^2} - \frac{c^2}{g} + x_0, \quad (14)$$

where x_0 is the initial position of the object. This means that due to Lorentz contraction, the acceleration in the world frame is dependent on the velocity of the camera, which results in a hyperbolic motion with asymptotic speed c in the world frame O , as seen in Figure 8. Both v and x are defined as a function of time t in the world frame O , while we are capturing the scene from the sensor frame O' , and therefore accounting for the proper time t' . These two times t and t' are related as:

$$t = \frac{c}{g} \sinh\left(\frac{gt'}{c}\right). \quad (15)$$




Figure 6: Relativistic phenomena for the cube scene due to an accelerating camera with focal length in world frame $d = 36$ mm. The camera starts at rest, and accelerates up to $v = 0.9c$ towards the scene, with acceleration $g = 0.01c s^{-1}$ in the camera frame O' . Note the slight differences in the field of view with respect to Figure 5.




Figure 7: Comparison of taking into account the effect of relativistic acceleration on the camera (right), against the effects of using constant speed (left). The camera has a focal length $d = 36$ mm, and proper acceleration $g = 0.01c s^{-1}$. Velocities shown here are $v = \beta c = [0.0..0.99c]$. Note that the simultaneity between the sensor and pinhole frames is broken in the scenario with acceleration, and therefore they have different speeds at proper time t' . As a consequence, the focal length d' in the moving frame O' is transformed differently in the presence of camera acceleration than when moving at constant speeds, an effect which becomes more visible at increasingly higher speeds.




Figure 8: (a) World lines for the sensor and the pinhole in an accelerated scenario ($g = 0.002c s^{-1}$), and local frames of reference of the sensor at different time instants and thus with different velocities (dotted axes). (b-d) Space-time diagrams for the sensor and the pinhole in the three local frames of reference drawn in (a). Note how the focal length (distance between sensor and pinhole) measured in the sensor frame increases as the velocity of the sensor increases due to acceleration.

Note that Equations 13 and 14 are only valid for constant proper acceleration g ; for time varying acceleration we would need to solve an initial value problem [Wei00b].

Acceleration has an important effect on our extended camera model, since an accelerating body changes its frame of reference at each instant; this, according to Equation 4, would make the focal distance d' increase as the object accelerates. However, when considering two different objects with the same acceleration and initial speed with respect to a

world frame, we need to take into account that simultaneity is broken: two simultaneous instants in the world frame O might not be so in accelerating frame O' . The consequence of this is that the distance between the two objects (sensor and pinhole in our case) will be constant when measured in the world frame (Figure 8(a)), but will keep changing when measured in a frame local to either of the objects (as shown in Figure 8(b-d) for three different velocities of the sensor with respect to the world frame, corresponding to three different time instants along the world line of the sensor, de-

picted in Figure 8(a)). Note that in our case, the distance between the two objects corresponds to the focal length (d in world frame O), which, when measured in the sensor frame O' (d'), is changing as the sensor (and pinhole) accelerates. This affects the field of view of the camera, and thus needs to be taken into account for the rendering framework.

We model this varying distance d' using the formulae derived by Styer [Sty07]:

$$d' = \frac{c^2}{g} \left(1 - \sqrt{\frac{1 - (v/c)^2 + (gd/c^2)^2 (v/c)^2}{1 - (v/c)^2}} \right) + \gamma d, \quad (16)$$

where v is the speed of the sensor in the world frame (Equation 13). Note that once both the sensor and pinhole stop accelerating, and therefore both are at constant speed βc , then $d' = \gamma d$, as explained in Section 4.3[¶]. Equation 16 shows that the distance d' varies non-linearly with the distance d in the world frame; in fact, the larger the distance, the more notable are the effects due to an accelerating camera.

Given d' , we can obtain the transformed field of view of the camera α' as $\alpha' = 2 \arctan(w/d')$ (for constant speed $g = 0$ it corresponds to Equation 4). Note that in order to maintain a constant field of view α' during acceleration, different accelerations would be needed in the pinhole and the sensor, as derived by Franklin [Fra10].

Figure 6 shows a sequence of an accelerating camera in the *cube* scene. We can see that, as opposed to the case in which a constant camera speed is used, the reduced field of view due to the Lorentz transform reduces the effect of light aberration observed with a camera at constant speed. This effect is more explicit in Figure 7, which compares the effect of constant speed and acceleration.

4.5. Relativistic Rotation

Providing free navigation of a scene depicting time-resolved light transport implies that the viewers should be allowed to rotate the camera. However, there is no universally accepted theory of relativistic rotation [RR04].

In the absence of a universal theory of relativistic rotation, we propose an approximation for the particular case of very small rotation angles per frame. This assumption, which is feasible in a scenario in which the camera moves at relativistic speeds, allows us to effectively treat each differential area of the sensor as undergoing linear translation only in the direction of the sensor's normal, with a different velocity for each sensor differential, according to the linear velocity field Ψ , and with a zero-crossing at the axis of rotation.

[¶] When this takes place, there will be a transition period in which, measured in the sensor frame, the sensor will stop accelerating before the pinhole; we have omitted this there for clarity but refer the interested reader to the thorough description in [Sty07].

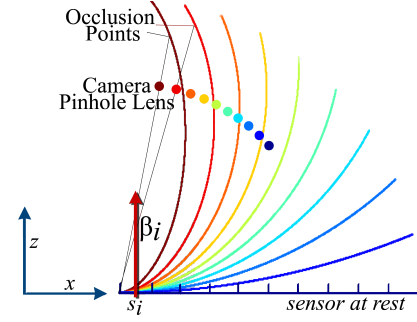


Figure 9: Distortion of the sensor (curves) and the camera pinhole lens (dots) due to relativistic rotation, as seen for the leftmost sensor differential area s_i . The velocity vector is perpendicular to such area (marked with a red arrow), with speeds ranging from $\beta_i = 0$ (dark blue, sensor at rest, no distortion), to $\beta_i = 0.9$ (dark red). Note how aberration in the sensor increases with higher speeds, even occluding the line at sight for very high values.

To simulate the rotation of the camera we therefore first divide the sensor S in different areas $s \in S$. Our approximation effectively turns each of them into a different translational frame, with linear velocity ψ_s . Then, for each s we render the scene applying the novel relativistic transformations introduced in this section, with a different β_s for each s (trivially obtained from an input β measured at the edge of the sensor). This causes the incoming radiance to be deformed differently depending on the position of the sensor where it is imaged. Figure 10 shows an example, where the sensor is rotating clockwise.

This transformation is enough if we are assuming that for each surface area s the camera lens moves with $\beta = \beta_i$, and that each sensor area s_i is independent of the rest. The center of the sensor thus has velocity $\beta = 0$, with each sensor area moving at different speeds. As a consequence, for an individual sensor area s_i , the sensor areas s_j suffer light aberration, each with different speed $\Delta\beta_{i \rightarrow j}$. This speed is the composition of β_i and β_j , and it is obtained using the relativistic law of composition of velocities:

$$\Delta\beta_{i \rightarrow j} = \frac{\beta_i + \beta_j}{1 + \beta_i \beta_j}. \quad (17)$$

This difference in speed makes both the camera lens and the sensor undergo light aberration, and therefore their projection is distorted following Equations 5 and 6. As a consequence, for each s_i with speed β_i , each ray from s_i to the camera lens is distorted, effectively changing the field of view at each sensor area. Additionally, since the rest of the sensor is aberrated with spatially-varying $\Delta\beta_{i \rightarrow j}$, it is even possible that the sensor plane itself occludes the incoming light, as shown in Figure 9. Figure 11 shows the same scene depicted in Figure 10, but including the aberration effects in the camera lens and sensor.




Figure 10: Assuming that the rotation angle θ can be neglected between frames, we model the rotation as a continuous linear velocity field on the sensor Ψ , so each differential area is assigned a different velocity ψ_s (left). This causes that depending on the position on the sensor, different relativistic transformations are applied on the scene. The rest of the frames show the effects of a clockwise rotation of the sensor; with $\beta = \{0, 0.4, 0.8, 0.99\}$ (measured at the edge of the sensor), assuming that no relativistic effect affects the camera. Figure 11 shows the effect when this last assumption is broken.




Figure 11: Relativistic effects of a clockwise rotation of the sensor, with $\beta = \{0, 0.25, 0.5, 0.75, 0.99\}$ (measured at the edge of the sensor), including relativistic aberrations in the camera. The lower half of the sequence shows the effect of rotation when taking into account the aberrations suffered by the sensor as seen from s_i (Figure 9), including the change in the FoV and the self-occlusion, while the upper half shows the rotation at the same speed ignoring its effect on the sensor. Note that at high-speeds, the edges of the sensor are occluded by the sensor itself due to sensor aberration, as shown in Figure 9 (marked here as a red rectangle). The image appearing to the right of this sensor occlusion comes from the scene, visible due to the high relativistic aberration caused by the rotation at high speeds.

4.6. Discussion

In order to include the relativistic effects described above, there are several assumptions that need to be made in order to keep physical plausibility. First, our work is restricted to special relativity, so we assume that no gravitational force is present in the scene; otherwise, general relativity would need to be considered. Additionally, we model relativistic acceleration based on the *clock hypothesis*.

We also need to assume that all components of the camera have infinite elasticity, to support the stress produced by relativistic acceleration and rotation, as well as a lensless pinhole camera. Including a lens would require adding lens transformations due to relativistic contractions (Sections 4.3 and 4.4), which would effectively change the plane of focus of the camera, and introduce additional time-varying delays. Our model does take into account the finite nature of the camera, including two frames of reference (pinhole and sensor), and a focal length, which can change under certain conditions. While in certain scenarios this effect may be negligible, it is important for time-varying radiance. Current transient data can reach a temporal resolution of picoseconds; at this timescale, a pinhole camera model makes a significant difference (Figure 7). To include rotation, we need to define a sensor with differential width, to avoid problems related with radius contraction as discussed in the Ehrenfest paradox [RR04]. Furthermore, for rotation we assume that the

pinhole aperture is almost at the middle of the sensor, so it does not move when rotating. Finally, the formulae exposed here are constrained to camera motion parallel to the sensor normal; for other motion directions the equations modeling the Doppler and searchlight effect need to be modified (see [WKR99] for the full mathematical derivation); in particular, the Doppler factor D should be re-derived accounting for the new configuration.

5. Implementation

Our implementation allows for real-time visualization of relativistic effects, both from real and simulated data. It is implemented in OpenGL as a stand-alone application, taking as input the reconstructed geometry of the scene, as well as the time-resolved data. The system is based on classic image-based rendering (IBR) techniques, where the shading of the surface is modeled by the images projected over the surface. In our case, we use x-y images from the x-y-t data cube to shade the geometry. The cube is stored as a 3D texture on the GPU in *world time* coordinates. This allows us to apply time-warping to adapt it to the new viewpoint in rendering time, by simply applying the transformation defined in Equation 1 (see Section 4.3).

Due to light aberration the geometry viewed from the camera is distorted. This distortion causes straight lines to

become curved, so the geometry has to be re-tessellated. Image-space warping, which has been used in many scenarios (e.g. [WKR00a, TDR*12, MWA*13]) is not a viable alternative due to the large extent of the deformations. Our implementation performs the re-tessellation off-line on the CPU, but it is straightforward to tessellate it on the GPU on the fly. During render time, each vertex should be then transformed according to Equation 2. Another suitable alternative to the tessellation could be the use of 4D local ray tracing, such as the GPU-based system by Müller et al. [MGW10].

The Doppler effect is introduced by modifying the wavelength of the outgoing illumination from the surfaces. To avoid the complexity of a full-fledged spectral renderer, we assume light with energy in only one wavelength of the spectrum. To display radiance we use a simple *wavelength-to-RGB* conversion encoded as a 1D texture. Note that wavelengths out of the visible spectrum are displayed with grayscale values in this paper.

Finally, when modeling the searchlight effect, we avoid the straightforward approach to access all frames in the x-y-t data cube, bounded by the duration \mathcal{T} , and integrate them. This would require several accesses to the 3D texture, which would hinder interactivity. Instead, we pre-integrate radiance in the temporal domain, and use anisotropic mipmapping to access the pre-integrated values, using \mathcal{T} to select the mipmap level in the time dimension.

6. Conclusions and Future Work

We have described a method to visualize light transport from a new perspective, no longer constrained by the assumption of infinite speed of light, and including relativistic effects with time-resolved light transport data (both real and simulated). Additionally, we have developed an interactive image-based rendering application for free navigation and exploration. We are able to account for both constant and accelerated motion, and have developed a method to approximate the effects of camera rotation, for which a definite solution does not exist in the physics literature. We have also introduced a pinhole camera model in a relativistic rendering framework, taking into consideration the variation of the focal length with camera motion at relativistic speeds, which in turn results in a varying field of view. We hope this will spur future research and help to better understand the complex behavior of time-resolved interactions between light and matter.

Of course there is plenty of future work ahead: Several assumptions discussed in Section 4.6 should be relaxed to generalize our model to other types of camera. Also, in order to provide a full physically-based solution mass (and therefore gravitational forces) must be included by considering general relativity. Our current implementation assumes Lambertian surfaces, so the viewing angle with respect to the normal has no influence in the result. This assumption can be relaxed by using more sophisticated IBR techniques, e.g. [BG01].

Acknowledgements

We thank the reviewers for their insightful comments and suggestions, Paz Hernando for helping with the figures, and Adolfo Muñoz and Ivan Bailera for proofreading. This research has been funded by the European Commission, Seventh Framework Programme (project VERVE), the Spanish Ministry of Science and Technology (project LIGHTSLICE) and the BBVA Foundation. Diego Gutierrez is additionally supported by a Google Faculty Research Award. Belen Masia was additionally funded by an NVIDIA Graduate Fellowship.

References

- [Arv93] ARVO J.: Transfer equations in global illumination. In *Global Illumination, SIGGRAPH'93 Course Notes* (1993). 2
- [Bet98] BETTS C.: Fast rendering of relativistic objects. *Journal of Visualization and Computer Animation* 9, 1 (1998), 17–31. 2
- [BG01] BOIVIN S., GAGALOWICZ A.: Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proceedings of SIGGRAPH '01* (2001). 11
- [Bri70] BRILLOUIN L.: *Relativity reexamined*. Academic Press, 1970. 7
- [CLC96] CHANG M.-C., LAI F., CHEN W.-C.: Image shading taking into account relativistic effects. *ACM Trans. Graph.* 15, 4 (1996). 1, 2
- [Dew63] DEWAN E.: Stress effects due to Lorentz contraction. *American Journal of Physics* 31 (1963). 1, 7
- [DPV11] DOAT T., PARIZOT E., VÉZIEN J.-M.: Novative rendering and physics engines to apprehend special relativity. In *Proceedings of the Eurographics Conference on Virtual Environments* (2011). 2
- [Ein61] EINSTEIN A.: *Relativity: the special and the general theory*. Crown Publishers, 1961. 5
- [FG10] FRIEDMAN Y., GOFMAN Y.: A new relativistic kinematics of accelerated systems. *Physica Scripta* 82, 1 (2010). 7
- [Fra10] FRANKLIN J.: Lorentz contraction, Bell's spaceships and rigid body motion in special relativity. *European Journal of Physics* 31, 2 (2010). 9
- [GNJ08] GUTIERREZ D., NARASIMHAN S., JENSEN H., JAROSZ W.: Scattering. In *ACM SIGGRAPH Asia'08 Courses* (2008). 1
- [HD89] HSIUNG P.-K., DUNN R. H. P.: Visualizing relativistic effects in spacetime. In *Proceedings of the Supercomputing Conference* (1989). 1, 2
- [HT90] HSIUNG P.-K., THIBADEAU R. H.: Spacetime visualization of relativistic effects. In *ACM Computer Science Conference'90* (February 1990). 2
- [HTC*90] HSIUNG P.-K., THIBADEAU R. H., COX C. B., DUNN R. H. P., WU M., OLBRICH P. A.: Wide-band relativistic Doppler effect visualization. In *Proceedings of the 1st Conference on Visualization '90* (1990), pp. 83–92. 2
- [HTCD90] HSIUNG P.-K., THIBADEAU R. H., COXT C. B., DUNN R. H. P.: Time dilation visualization in relativity. In *Proceedings of Supercomputing'90* (1990), pp. 835–844. 2
- [HTW90] HSIUNG P.-K., THIBADEAU R. H., WU M.: T-buffer: Fast visualization of relativistic effects in space-time. *SIGGRAPH Comput. Graph.* 24, 2 (1990). 2

- [JMG13] JARABO A., MASIA B., GUTIERREZ D.: *Transient Rendering and Relativistic Visualization*. Tech. Rep. TR-01-2013, Universidad de Zaragoza, April 2013. 2
- [JMM*14] JARABO A., MARCO J., MUÑOZ A., BUISAN R., JAROSZ W., GUTIERREZ D.: A framework for transient rendering. *ACM Trans. Graph.* 33, 6 (2014). 1, 2, 3
- [JMV*13] JARABO A., MASIA B., VELTEN A., BARSİ C., RASKAR R., GUTIERREZ D.: Rendering relativistic effects in transient imaging. In *Proceedings of the Spanish Conference on Computer Graphics (CEIG'13)* (2013). 2
- [KHDR09] KIRMANI A., HUTCHISON T., DAVIS J., RASKAR R.: Looking around the corner using transient imaging. In *Proceedings of ICCV'09* (2009). 2
- [KRWZ02] KRAUS U., RUDER H., WEISKOPF D., ZAHN C.: Was Einstein noch nicht sehen konnte - Visualisierung relativistischer Effekte. *Physik Journal* 7/8 (2002). 2
- [KTS13] KORTEMEYER G., TAN P., SCHIRRA S.: A slower speed of light: Developing intuition about special relativity with games. In *Proceedings of the International Conference on the Foundations of Digital Games (FDG'13)* (2013). 1, 3
- [MGW10] MÜLLER T., GROTTTEL S., WEISKOPF D.: Special relativistic visualization by local ray tracing. *IEEE Trans. on Visualization and Computer Graphics* 16, 6 (2010). 11
- [MWA*13] MASIA B., WETZSTEIN G., ALIAGA C., RASKAR R., GUTIERREZ D.: Display adaptive 3D content remapping. *Computers & Graphics* (2013). 11
- [NBVR14] NAIK N., BARSİ C., VELTEN A., RASKAR R.: Estimating wide-angle, spatially varying reflectance using time-resolved inversion of backscattered light. *JOSA A* 31, 5 (2014). 1
- [NZV*11] NAIK N., ZHAO S., VELTEN A., RASKAR R., BALAK.: Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graph.* 30 (2011). 1, 2
- [RR04] RIZZI G., RUGGIERO M. L.: *Relativity in Rotating Frames*. Kluber Academic, 2004. 9, 10
- [SSD08] SMITH A., SKORUPSKI J., DAVIS J.: *Transient Rendering*. Tech. Rep. UCSC-SOE-08-26, School of Engineering, University of California, Santa Cruz, February 2008. 2
- [Sty07] STYER D. F.: How do two moving clocks fall out of sync? a tale of trucks, threads, and twins. *American Journal of Physics* 75, 9 (2007). 7, 9
- [TDR*12] TEMPLIN K., DIDYK P., RITSCHEL T., MYSZKOWSKI K., SEIDEL H.-P.: Highlight microdisparity for improved gloss depiction. *ACM Trans. Graph.* 31, 4 (2012). 11
- [Ter59] TERRELL J.: Invisibility of the Lorentz contraction. *Phys. Rev.* 116 (1959). 7
- [VWG*12] VELTEN A., WILLWACHER T., GUPTA O., VEERARAGHAVAN A., BAWENDI M. G., RASKAR R.: Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3 (2012). 1, 2
- [VWJ*12] VELTEN A., WU D., JARABO A., MASIA B., BARSİ C., LAWSON E., JOSHI C., GUTIERREZ D., BAWENDI M. G., RASKAR R.: Relativistic ultrafast rendering using time-of-flight imaging. In *ACM SIGGRAPH 2012 Talks* (2012). 1, 2
- [VWJ*13] VELTEN A., WU D., JARABO A., MASIA B., BARSİ C., JOSHI C., LAWSON E., BAWENDI M., GUTIERREZ D., RASKAR R.: Femto-photography: Capturing and visualizing the propagation of light. *ACM Trans. Graph.* 32, 4 (2013). 1, 2, 3
- [WBE*06] WEISKOPF D., BORCHERS M., ERTL T., FALK M., FECHTIG O., FRANK R., GRAVE F., KING A., KRAUS U., MÜLLER T., NOLLERT H.-P., RICA MENDEZ I., RUDER H., SCHAFHITZEL T., SCHAR S., ZAHN C., ZATLOUKAL M.: Explanatory and illustrative visualization of special and general relativity. *IEEE Trans. on Visualization and Computer Graphics* 12, 4 (2006). 1, 2
- [Wei00a] WEISKOPF D.: Fast visualization of special relativistic effects on geometry and illumination. In *Data Visualization 2000*. 2000, pp. 219–228. 2
- [Wei00b] WEISKOPF D.: An immersive virtual environment for special relativity. In *Proceedings WSCG'00* (2000). 2, 8
- [Wei01] WEISKOPF D.: *Visualization of Four-Dimensional Spacetimes*. PhD thesis, Eberhard-Karls-Universität Tübingen, 2001. 1, 2, 7
- [WKR99] WEISKOPF D., KRAUS U., RUDER H.: Searchlight and Doppler effects in the visualization of special relativity: A corrected derivation of the transformation of radiance. *ACM Trans. Graph.* 18, 3 (1999). 2, 4, 5, 6, 10
- [WKR00a] WEISKOPF D., KOBRAS D., RUDER H.: Real-world relativity: Image-based special relativistic visualization. In *Proceedings of IEEE Visualization'00* (2000). 2, 11
- [WKR00b] WEISKOPF D., KRAUS U., RUDER H.: Illumination and acceleration in the visualization of special relativity: a comment on fast rendering of relativistic objects. *Journal of Visualization and Computer Animation* 11, 4 (2000). 2
- [WVO*13] WU D., VELTEN A., O'TOOLE M., MASIA B., AGRAWAL A., DAI Q., RASKAR R.: Decomposing global light transport using time of flight imaging. *International Journal of Computer Vision* 105, 3 (2013). 2
- [WWB*13] WU D., WETZSTEIN G., BARSİ C., WILLWACHER T., DAI Q., RASKAR R.: Ultra-fast lensless computational imaging through 5D frequency analysis of time-resolved light transport. *International Journal of Computer Vision* (2013). 2