Real-World Relativity: Image-Based Special Relativistic Visualization

Daniel Weiskopf, Daniel Kobras, and Hanns Ruder*

Institute for Astronomy and Astrophysics†
University of Tübingen

Abstract

This paper describes a novel rendering technique for special relativistic visualization. It is an image-based method which allows to render high speed flights through real-world scenes filmed by a standard camera. The relativistic effects on image generation are determined by the relativistic aberration of light, the Doppler effect, and the searchlight effect. These account for changes of apparent geometry, color, and brightness of the objects. It is shown how the relativistic effects can be taken into account by a modification of the plenoptic function. Therefore, all known image-based non-relativistic rendering methods can easily be extended to incorporate relativistic rendering. Our implementation allows interactive viewing of relativistic panoramas and the production of movies which show super-fast travel. Examples in the form of snapshots and film sequences are included.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture I.3.8 [Computer Graphics]: Applications—Special relativity J.2 [Physical Sciences and Engineering]: Physics—Theoretical astrophysics

Keywords: image-based rendering, plenoptic function, scientific visualization, special relativity

1 Introduction

Special relativity is widely regarded as a difficult and hardly comprehensible theory, mainly because the properties of space, time, and light in relativistic physics are totally different from those in classical, Newtonian physics. In many respects, they are contrary to human experience and everyday perception, which is based on low velocities.

In the real world, mankind is limited to very slow velocities compared to the speed of light. For example, the speed of light is a million times faster than the speed of an airplane and 40,000 times faster than the speed at which the Space Shuttle orbits the earth. Even in the long term, there is no hope of achieving velocities comparable to the speed of light. Therefore, visualization is the only means of directly exploring the realm of special relativity and may help to stimulate intuition and motivation of people interested in the theory.

There exist well-known rendering techniques for special relativity. However, these methods use a standard geometry-based representation of three-dimensional scenes and hence require time-consuming geometrical modeling and costly rendering. The crucial shortcoming of the geometry-based methods is missing photo-realism. In discussions with editors and producers from TV industry, the demand for relativistic flights through real-world scenes was strongly expressed. Accordingly, special relativistic visualization has not been widely used in professional TV and film productions for edutainment and education yet, although there exist numerous TV documentations about Einstein and his special theory of relativity.

In this paper, we propose a novel image-based approach to special relativistic rendering. This approach overcomes problems of geometry-based rendering and has the following important advantages: No three-dimensional geometric modeling is needed, rendering costs are negligible, and photo-realism is easily achieved.

The basic idea of the image-based approach to relativistic visualization is presented in Sect. 3. We show how all relativistic effects on image generation can be covered by a modification of the plenoptic function[1]. Therefore, the full three-dimensional information about the scene is not required for relativistic rendering. In this framework, only one additional step is appended to the normal non-relativistic rendering pipeline, which is otherwise left unchanged. Hence, the relativistic transformation can easily be incorporated in all known image-based rendering methods.

We present two implementations of image-based relativistic rendering. The first implementation is an interactive panorama viewer which creates snapshots of a panorama with the camera moving at arbitrary speed. The second implementation is a batch job-oriented tool for the production of relativistic movies playing in real-world scenes. It stitches and blends series of views taken by different cameras in order to generate a sequence of images for a relativistic flight.

2 Previous and Related Work

Remarkably, the issue of visual perception in special relativity was ignored for a long time, or wrong interpretations were given. Apart from a previously disregarded article by Lampa[17] in 1924 about the invisibility of the Lorentz contraction, it was only in 1959 that

* {weiskopf,kobras,ruder}@tat.physik.uni-tuebingen.de

† Institute for Astronomy and Astrophysics, Section Theoretical Astrophysics, University of Tübingen, Auf der Morgenstelle 10, D-72076 Tübingen, Germany

the first solutions to this problem were described by Penrose[23] and Terrell[32]. Later, more detailed descriptions of the geometrical appearance of fast moving objects were given by Weisskopf[35], Boas[2], Scott and Viner[29], and Scott and van Driel[28].

Hsiung and Dunn[13] were the first to use advanced visualization techniques for image shading of fast moving objects. They proposed an extension of normal three-dimensional ray tracing. This technique accounts for relativistic effects on the apparent geometry as seen by the observer. Hsiung et al.[14] investigated relativistic ray tracing in more detail and included the visualization of the Doppler effect.

Hsiung et al.[15] introduced the time-buffer for fast visualization of relativistic effects. The time-buffer technique resembles the normal z-buffer and can be mapped onto it. The time-buffer technique allows for relativistic polygon rendering, a scan-line method. It is based on the apparent shapes of objects as seen by a relativistic observer. Gekelman et al.[9] and Chang et al.[3] investigated the polygon rendering approach in detail and gave a comprehensive presentation.

Weiskopf[33] introduced texture-based relativistic rendering for visualizing the apparent geometry of fast moving objects. This approach performs the relativistic transformation on the image plane by texture mapping.

A lot of research has been conducted on the field of non-relativistic image-based rendering. *QuickTime VR*[4] is a well-known method for image-based rendering, which uses panorama pictures. More advanced techniques include plenoptic modeling[20], light fields[18], the lumigraph[12], view morphing[30], and hybrid geometry and image-based rendering[6].

3 Basic Idea

One basic feature of special relativity is the absence of a single universal frame of reference and of a universal time. Any inertial frame is equally valid to describe the physical world.

Often an egocentric point of view is adopted to derive the properties of relativistic rendering, i.e., the camera is at rest and the objects are moving. In this paper, we rather take an exocentric point of view. Here, the objects are considered to be at rest and the observer—the camera—is moving at high speed. In the appendix, the equivalence of both points of view is shown explicitly.

The essence of all image-based rendering methods is the evaluation of the plenoptic function[1]. The full plenoptic function $P(x, y, z, t, \theta, \phi, \lambda)$ is the radiance of the light depending on the direction (θ, ϕ) in spherical coordinates, the spatial position (x, y, z), the time t, and the wavelength λ. The definition of wavelength-dependent radiance can be found, e.g., in [11, Chapt. 13]. Polarization is usually neglected.

We restrict ourselves to a static world, in which all objects and light sources are at rest relative to each other and relative to the objects' frame denoted S_{obj}. In S_{obj}, the plenoptic function can be determined by standard image-based rendering algorithms, since the finite speed of light can be neglected in this static situation.

First, consider the generation of a snapshot taken by a camera at rest in S_{obj}. The spatial position of the camera is (x, y, z) and the time is t. All the information needed for this snapshot is contained in the reduced three-parameter plenoptic function $\tilde{P}(\theta, \phi, \lambda)$, which is evaluated at the respective position and time.

Then, let us bring special relativity back into the game. Consider another observer that is moving relative to the objects. His or her rest frame is denoted S_{observer}. This observer is taking a snapshot at the same position and time as the first observer that is at rest in S_{obj}. What is the plenoptic function for this moving observer and how is it connected to the plenoptic function for the observer at rest?

In general, physical properties can be transformed from one frame of reference to another by the so-called Lorentz transformation. Here, all relevant physical properties are contained in the plenoptic function. Therefore, only the Lorentz transformation of the plenoptic function has to be known. This transformation is discussed in the next section.

Once the plenoptic function $\tilde{P}(\theta, \phi, \lambda)$ with respect to S_{obj} is transformed to $\tilde{P}'(\theta', \phi', \lambda')$ with respect to S_{observer}, the normal rendering process can generate the image seen by the fast moving camera because $\tilde{P}'(\theta', \phi', \lambda')$ is the plenoptic function at rest relative to this camera. (The primed quantities are with respect to S_{observer}.) In this way, all relativistic effects are isolated in the form of the Lorentz transformation of the plenoptic function. The locality property of this transformation allows us to generate relativistic images without knowledge of the depth, or three-dimensional, information about the surrounding scene. Due to the relativity principle the transformation of the plenoptic function can account for both a fast camera and rapidly moving objects.

4 Lorentz Transformation

In this section, the Lorentz transformation of the plenoptic function is described. Relevant for this transformation are the relativistic aberration of light, the Doppler effect, and the searchlight effect. For a detailed presentation of special relativity we refer to [21, 22, 27].

The relativistic aberration of light causes a rotation of the direction of light when one is changing from one inertial frame of reference to another. The aberration of light is sufficient to completely describe the apparent geometry seen by a fast moving camera. Figure 1 illustrates the aberration of light.

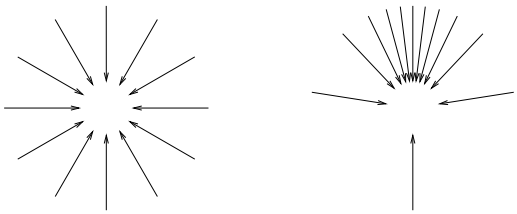

Figure 1: Relativistic aberration of light. The left image shows some of the light rays registered by an observer at rest. The right image shows the same light rays with the observer moving upwards at 90 percent of the speed of light.

The Doppler effect accounts for the transformation of wavelength from one inertial frame of reference to another and causes a change in color.

The searchlight effect is based on the transformation of wavelength-dependent radiance from one inertial frame of reference to another. The transformation of radiance increases the brightness of objects ahead when the observer is approaching these objects at high velocity.

Let us consider two inertial frames of reference, S and S', with S' moving with velocity v along the z axis of S. The usual Lorentz transformation along the z axis connects frames S and S'.

In reference frame S, consider a light ray with the direction (θ, ϕ) and the wavelength λ. In frame S', the light ray is described by the direction (θ', ϕ') and the wavelength λ'. These two representations are connected by the expressions for the relativistic aberration of light, cf. [22],

$$\cos\theta' = \frac{\cos\theta - \beta}{1 - \beta\cos\theta}, \tag{1}$$

$$\phi' = \phi, \tag{2}$$

and for the Doppler effect,

$$\lambda' = D\,\lambda. \tag{3}$$

The Doppler factor D is defined as

$$D = \frac{1}{\gamma(1 - \beta\cos\theta)} = \gamma(1 + \beta\cos\theta'), \tag{4}$$

where $\gamma = 1/\sqrt{1-\beta^2}$, $\beta = v/c$, and c is the speed of light.

Wavelength-dependent radiance L_λ is transformed from one frame of reference to another according to

$$L'_\lambda(\lambda', \theta', \phi') = D^{-5} L_\lambda(\lambda, \theta, \phi). \tag{5}$$

A derivation of this relation can be found in [34]. Please note that the subscript λ is only attached to indicate wavelength dependency of radiance; it is not a parameter. The proper wavelength parameter is included as a function parameter.

The relativistic aberration of light, the Doppler effect, and the searchlight effect can be combined to form the transformation of the plenoptic function from S to S':

$$\begin{aligned} \tilde{P}'(\theta', \phi', \lambda') &= D^{-5}\tilde{P}(\theta, \phi, \lambda) \\ &= D^{-5}\tilde{P}\left(\arccos\frac{\cos\theta' + \beta}{1 + \beta\cos\theta'}, \phi', \frac{\lambda'}{D}\right) \end{aligned} \tag{6}$$

By inverting Eqs. (1)–(3), the parameters θ, ϕ, and λ are substituted by terms containing θ', ϕ', and λ'.

Usually, the direction of motion is not identical to the z axis. Therefore, additional rotations of the coordinate system have to be considered before and after the aberration transformation. These rotations are identical to the standard rotations in three-dimensional Euclidean space. By including these rotations, we obtain the complete Lorentz transformation of the plenoptic function.

With the notation from the previous section, frame S coincides with S_{obj} and frame S' with S_{observer}. Please note that the transformed plenoptic function depends only on the original plenoptic function, the observer's velocity and direction of motion, and the orientations of the two reference frames.

All the information registered by the observer's camera is carried from the surrounding environment to the observer's position via light rays—there is no direct interaction between the camera and the outside objects. The process of image generation is localized at the observer's position, where the actual interaction between the incoming photons and the detector (camera) takes place. In geometric optics, direction, wavelength, radiance, and polarization completely determine the incoming light. The relativistic transformation of the first three quantities is exactly and uniquely described by the aberration, and the Doppler and searchlight effects. Polarization is neglected because it is not registered by standard cameras.

The plenoptic function combines the information about direction, wavelength, and radiance. Accordingly, the Lorentz transformation of the plenoptic function is just the combination of the aberration, and the Doppler and searchlight effects. Bearing the locality of light detection in mind, it can be seen that the transformed plenoptic function provides all the information that is registered by the moving camera. Therefore, the method of this paper generates images which are in total compliance with the physics of special relativity.

5 Relativistic Rendering

Image-based relativistic rendering extends the standard non-relativistic techniques by a transformation of the plenoptic function according to the previous section. This extension is located at the end of the rendering pipeline, just before the final image is generated. All other parts of the rendering pipeline are unaffected.

In the following, some variations of relativistic rendering are described. In particular, we address the issue of missing data, since the wavelength dependency of the plenoptic function cannot be measured by standard cameras. In most cases, data for image-based rendering is acquired by cameras which are sensitive to only three RGB colors and not to the full power spectrum of the incoming light.

5.1 Completely Relativistic Rendering

If the wavelength-dependent plenoptic function $\tilde{P}(\theta, \phi, \lambda)$ is provided in the non-relativistic situation, the transformed plenoptic function $\tilde{P}'(\theta', \phi', \lambda')$ can be computed according to Sect. 4. It is important that $\tilde{P}(\theta, \phi, \lambda)$ is known for an extended range of wavelengths, so that $\tilde{P}'(\theta', \phi', \lambda')$ can be evaluated for wavelengths in the visible range after Doppler-shifting.

Each pixel on the image plane has corresponding spherical coordinates (θ', ϕ'), which are transformed to (θ, ϕ) in the objects' frame. Therefore, each pixel is associated with the wavelength-dependent radiance,

$$L'_\lambda(\lambda') = \tilde{P}'(\theta', \phi', \lambda').$$

For the final display on the screen, three tristimulus values such as RGB have to be calculated from this wavelength-dependent radiance. The RGB values (c_R, c_G, c_B) can be obtained by

$$c_i = \int L'_\lambda(\lambda')\,\bar{f}_i(\lambda')\,d\lambda', \qquad i = R, G, B,$$

where $\bar{f}_i$ are the respective color-matching functions for RGB, cf. [38].

5.2 Apparent Geometry

The relativistic effects on the apparent geometry can be visualized by using only a partial transformation of the plenoptic function. Here, solely the effects of the aberration of light are taken into account and the searchlight and Doppler effects are neglected, i.e., only the direction (θ, ϕ) of the incoming light is transformed and all other effects are ignored.

This visualization technique is useful when the full spectral information of the plenoptic function is not available, since this information is not needed for the visualization of apparent geometry. Nevertheless, even this restricted relativistic rendering provides some insight into the special theory of relativity and creates impressive visual effects, as shown in Sect. 7.

5.3 Reconstruction of the Power Spectrum

In most cases, data for image-based rendering does not comprise the full power spectrum, but only three RGB values. The power spectrum has to be reconstructed from RGB values in order to include the relativistic effects on geometry and illumination. Unfortunately, this reconstruction is not unique because infinitely many spectra map to one RGB triplet. This phenomenon is called metamerism, cf. [38].

However, a possible spectrum can always be determined and metamerism gives a lot of freedom of doing so. A straightforward approach models the three RGB values by the line spectrum consisting of the corresponding primaries[10]. Sun et al.[31] propose the representation by Gaussian functions with adapted width. Another approach uses Fourier functions[10].

We find the dominant wavelength model[7] useful because it provides a smooth change of color and brightness for a wide range of Doppler factors. The corresponding spectral power distribution consists of a spike at the dominant wavelength and of a uniform distribution, i.e., white light. The luminance and excitation purity determine the levels of the two parts of the spectrum. The parameters for the dominant wavelength model can be computed from RGB values according to [7]. The relativistic situation requires only one slight extension of the original model. Here, the uniform part of the spectrum is not restricted to the range of visible wavelengths, but comprises a larger interval. In this way, the spectrum is still present after Doppler-shifting.

With the reconstructed wavelength-dependent plenoptic function, the algorithm from Sect. 5.1 can generate the fully relativistic image.

5.4 Rendering of a Film Sequence

So far, the generation of just a single snapshot has been described. But how can a film sequence with a fast camera be produced?

In principle, it works the same way as in the non-relativistic situation. The path of the camera is discretized into a finite set of positions. For every element of this set the plenoptic function is evaluated. Therefore, the plenoptic function has to be known at these positions. Then, the relativistic transformation is computed and the corresponding image is generated. Finally, a list of snapshots which represent the film sequence is obtained.

For the film to be physically sound, not only the generation of each single snapshot has to be correct, but also the path of the camera itself. As long as the camera is moving uniformly—at constant speed and with a constant direction of motion—the camera is trivially placed at equidistant positions. However, even an accelerated camera can be described by special relativity. In [26] it is shown how the trajectory of an accelerating observer can be computed. Therefore, the positions and velocities of the camera for each snapshot can be calculated, and image-based relativistic rendering can be performed. This method is valid because the generation of a single image is only determined by the position and velocity of the viewer and by the standard camera parameters, but not by the "history" of the trajectory or the acceleration of the observer.

Our everyday experience is based on length scales in the range of meters, time scales in the range of seconds, and velocity scales in the range of meters per second, i.e., the velocities we are used to are approximately eight orders of magnitude smaller than the speed of light. Length l, time t, and velocity v are related by $v = dl/dt$. Therefore, one has to change the length, time, or velocity scales in order to notice relativistic effects. For example, the time scales could be reduced to the orders of 10^{-8} seconds. We can think of playing a respective recording at super slow-motion, so that we are able to watch processes which involve time spans of only 10^{-8} seconds. Another option is to artificially reduce the speed of light in vacuum, for example, to walking speed. An instructive illustration of reduced speed of light can be found in Mr Tompkins' world by Gamow[8][1]. In the third approach, length scales are increased to the range of light seconds.

The change of scales is the reason why we can only support static scenes. The real-world camera image is recorded using the true values for the length, time, and velocity scales. In particular, the true speed of light is effectively infinite in all practical situations. The relativistic simulation of a dynamic scene would use images which are instantaneously transported from the object to the camera, instead of the correct, retarded images which take into account the reduced speed of light.

[1] Please note that the illustrations in Mr Tompkins do not show visual perception within special relativity, but only the measurements of Lorentz-contracted lengths.

5.5 Magnification and Anti-Aliasing

The aberration of light does not conserve the element of solid angle. In fact, the infinitesimal solid angle is transformed according to

$$\begin{aligned} \frac{d\Omega'}{d\Omega} &= \frac{\sin\theta'}{\sin\theta}\frac{d\theta'}{d\theta}\frac{d\phi'}{d\phi} = \frac{d(\cos\theta')}{d(\cos\theta)}\frac{d\phi'}{d\phi} \\ &= \frac{d(\cos\theta')}{d(\cos\theta)} = D^2, \end{aligned} \tag{7}$$

with the use of Eqs. (1), (2), and (4).

Therefore, the transformation of the plenoptic function causes a magnification opposite to the direction of motion, whereas objects ahead are scaled down. The demand for a higher resolution towards the back has to be taken into account when the original data for the plenoptic function is acquired. In the rendering process, the sampled data is accessed by bilinear interpolation.

The image contraction for the front view might cause aliasing effects, especially for extremely high velocities. These effects can be reduced by standard supersampling and postfiltering on the image plane.

Since the sampled plenoptic function can be stored in the form of a two-dimensional texture for the coordinates θ and ϕ, antialiasing can alternatively be based on texture filtering techniques. Texture mapping can be considered as the process of calculating the projection of a screen pixel onto the texture image—which is called *footprint*—and computing an average value which best approximates the correct pixel color. There exist a couple of filtering methods, the most prominent of which is MIPmapping[36]. This standard technique supports only a quadratic footprint. Hence, it is not very well suited for our application. The mapping by the aberration equation can generate prolate, anisotropic, and distorted footprints because it substantially changes the angle θ, whereas it leaves the angle ϕ invariant.

Therefore, techniques which support more complex footprints are required. Summed-area tables[5] (SAT), for example, allow prolate rectangular footprints. We have adopted the idea of rectangular axes-aligned footprints for the relativistic situation. In contrast to SAT, we leave out the computation of prefiltered data tables, since each texture is used only once in the rendering process. Filtering provides fair image-quality, even for velocities as high as $\beta = 0.99$ and for images with high spatial frequencies. Respective examples are shown in Sect. 7. The main advantage of a rectangular footprint over more complex footprints is faster computation and rendering.

Fast footprint MIPmapping[16] is based on quadrilateral footprints and makes use of precomputed MIPmaps and weighting tables. Quadrilateral footprints are an improved approximation compared to rectangular axes-aligned footprints. They support anisotropic, rotated, and distorted footprints. Despite the associated complexity, fast footprint MIPmapping should be able to achieve good rendering performance. Its relativistic adaption will be implemented in future work.

6 Implementation

We have implemented the relativistic panorama viewer *Imagine* (IMAge-based special relativistic rendering enGINE), which can read panoramas in the *LivePicture* format[19]. This format is similar to *QuickTime VR*, but uses a spherical projection instead of a cylindrical projection. Therefore, a complete 4π sterad view is supported.

The interactive viewer is written in C++ and is based on standard OpenGL 1.1[37] and *QGLViewer*[24]. The virtual camera is surrounded by a sphere onto which the panorama texture is mapped.

Figure 2: Digital video camera mounted on fork arm.

Texture mapping hardware is used to achieve high rendering performance. The relativistic effects on the apparent geometry are implemented by transforming the texture coordinates according to the relativistic aberration of light. The non-interactive part of the viewer uses software rendering to implement completely relativistic visualization by reconstructing the spectrum according to Sect. 5.3. Texture filtering as described in the previous section is not implemented yet.

Another implementation is *Off-Terdingen*, which is an off-screen, batch job-oriented relativistic movie renderer. It is able to produce movies of relativistic flights through real-world scenes. It is a C++-based software renderer which stitches and blends series of views taken by different cameras in order to generate a sequence of images for a relativistic flight. The parameters and orientations of the cameras are supplied manually. *Off-Terdingen* provides anti-aliasing by means of texture filtering, as described in Sect. 5.5. Additionally, standard supersampling on the image plane can be applied. The data of the original images is accessed by bilinear interpolation.

Adequate data acquisition for the non-relativistic panorama is an issue, since relativistic image-based rendering demands for higher quality of the initial data than standard panorama techniques. First, the resolution of original images has to suffice the magnification by the aberration formula, Eq. (7), when the observer looks into the backward direction. Secondly, a complete 4π sterad panorama should be recorded. Most commercially available panorama systems, however, are based on cylindrical projection, e.g., *QuickTime VR*.

Therefore, we built a camera system which can automatically film a 4π sterad field of view. A standard digital video camera is mounted on a fork arm which was originally designed for a telescope. Figure 2 shows the fork arm with camera. The fork arm is controlled by a mobile computer. Due to the specific geometry of the fork arm the camera can be placed in a way that avoids parallax artifacts when the camera is turned in different directions. The camera system is DV-based. Images are transferred to the mobile computer via IEEE 1394 (Firewire) link. The calibrated image data is stitched by *Off-Terdingen* to render spherical panoramas or relativistic views.

Figure 3: Non-relativistic view.

Figure 4: Relativistic visualization of apparent geometry with $\beta = 0.99$.

7 Results

Figures 3–5 and Color Plates 10–12 provide examples of image-based relativistic rendering. These images were produced by *Off-Terdingen*.

Figures 3–5 show a long corridor. Figure 3 provides the non-relativistic view of the scene. Figure 4 illustrates the effects on apparent geometry when the viewer is rushing into the scene with $\beta = 0.99$. A dominant effect is the increased apparent field of view—the objects seem to move away. Furthermore, straight lines which are perpendicular to the direction of motion become distorted to hyperbolae.

Figure 5 shows completely relativistic rendering with $\beta = 0.3$. Here, the power spectrum is reconstructed by using the dominant wavelength model. Changes in brightness due to the searchlight effect are noticeable. Color changes due to the Doppler effect are present, but cannot be reproduced on the gray-scale image. The searchlight effect heavily brightens the image, so the overall intensity has to be reduced to one half of that in Figs. 3 and 4 in order to avoid extreme clamping of the final gray-scale values.

The pictures on the first page and in Color Plate 10 show the apparent geometry for a snapshot of Yosemite Valley at $\beta = 0.95$. They exhibit the same effects as in Fig. 4, for example the distortion of straight lines to hyperbolae.

Color Plate 11 compares non-relativistic view, apparent geometry, and completely relativistic visualization, analogous to Figs. 3–5. In Color Plate 11(c), the color shift due to the Doppler effect is reproduced and a noticeable blueshift is shown.

Color Plate 12 compares filtering and supersampling techniques. This example shows the visualization of apparent geometry at $\beta = 0.99$. Image 12(a) is rendered without filtering and super-

Figure 5: Completely relativistic rendering with $\beta = 0.3$. The overall intensity is reduced to one half of that in Figs. 3 and 4 in order to avoid extreme clamping of the final gray-scale values.

sampling. Aliasing effects are noticeable, especially on the ceiling. Image 12(b) illustrates filtering with rectangular footprint, as described in Sect. 5.5. Aliasing artifacts are attenuated by texture filtering. In image 12(c), filtering and 2×2 supersampling are combined, yielding a better image quality than mere filtering. This indicates that more sophisticated footprints might improve filtering quality.

The accompanying video presents further examples of relativistic visualization. Here, only the relativistic effects on apparent geometry are taken into account.

The first part of the video shows a relativistic trip across a bridge. This movie was produced with the use of *Off-Terdingen* as well. The recordings of three cameras were stitched together to form the final movie. The film was shown on TV as part of a broadcast on Einstein and special relativity[25]. The first sequence shows the non-relativistic situation. The second sequence presents the relativistic case. The observer accelerates from non-relativistic speed to a maximal velocity of $\beta = 0.9$. In the third sequence, the relativistic and the non-relativistic views are compared. The second part of the video illustrates the interactive relativistic panorama viewer *Imagine*. It was recorded from computer screen during a simulation on an SGI Onyx2 with InfiniteReality graphics board.

8 Conclusion and Future Work

In this paper an image-based approach to special relativistic rendering has been introduced. This approach closes the gap between the well-known non-relativistic image-based techniques and relativistic visualization. We have shown how all relativistic effects on image generation can be covered by a transformation of the plenoptic function. Therefore, only slight modifications of existing rendering methods are required to incorporate the physically correct rendering of super-fast objects.

Photo-realistic images of rapidly moving real-world objects can be generated with great ease. Therefore, image-based special relativistic rendering is a powerful tool to generate movies and snapshots for both edutainment and educational purposes. We have specifically addressed aliasing problems caused by relativistic rendering. We have described supersampling and texture filtering methods to overcome these problems and render high-quality images.

In future work, we will improve the techniques for data acquisition. In particular, we will build a robot-based camera system which can automatically record a sequence of panoramas in order to generate a relativistic film with a large field of view. Resolution and sampling rate will be adapted to the different scalings due to aberration in order to achieve high-quality final pictures. The stitching software will be extended to automatically correct color and brightness variations and small misalignments of the original images. Furthermore, advanced texture filtering techniques will be investigated and implemented.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) and is part of the project D4 within the Sonderforschungsbereich 382. We thank WDR and the Quarks&Co team for their cooperation and the permission to use their material in the accompanying video.

References

[1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. In M. Landy and J. A. Movshon, editors, *Computational Models of Visual Processing*, pages 3–20, Cambridge, 1991. MIT Press.

[2] M. L. Boas. Apparent shape of large objects at relativistic speeds. *American Journal of Physics*, 29(5):283–286, May 1961.

[3] M.-C. Chang, F. Lai, and W.-C. Chen. Image shading taking into account relativistic effects. *ACM Transactions on Graphics*, 15(4):265–300, Oct. 1996.

[4] S. E. Chen. QuickTime VR – An image-based approach to virtual environment navigation. In *SIGGRAPH 95 Conference Proceedings*, pages 29–38, Aug. 1995.

[5] F. C. Crow. Summed-area tables for texture mapping. In *SIGGRAPH 84 Conference Proceedings*, pages 207–212, July 1984.

[6] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry and image-based approach. In *SIGGRAPH 96 Conference Proceedings*, pages 11–20, Aug. 1996.

[7] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, Reading, Massachusetts, 1990.

[8] G. Gamow. *Mr Tompkins in Paperback*. University Press, Cambridge, 1993.

[9] W. Gekelman, J. Maggs, and L. Xu. Real-time relativity. *Computers in Physics*, 5(4):372–385, July/Aug. 1991.

[10] A. S. Glassner. How to derive a spectrum from an RGB triplet. *IEEE Computer Graphics and Applications*, 9(4):95–99, July 1989.

[11] A. S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, 1995.

[12] S. J. Gortler, R. Grzeszczuk, and R. S. M. F. Cohen. The lumigraph. In *SIGGRAPH 96 Conference Proceedings*, pages 43–54, Aug. 1996.

[13] P.-K. Hsiung and R. H. P. Dunn. Visualizing relativistic effects in spacetime. In *Proceedings of Supercomputing '89 Conference*, pages 597–606, 1989.

[14] P.-K. Hsiung, R. H. Thibadeau, C. B. Cox, R. H. P. Dunn, M. Wu, and P. A. Olbrich. Wide-band relativistic doppler effect visualization. In *Proceedings of the Visualization 90 Conference*, pages 83–92, Oct. 1990.

[15] P.-K. Hsiung, R. H. Thibadeau, and M. Wu. T-buffer: Fast visualization of relativistic effects in spacetime. *Computer Graphics*, 24(2):83–88, Mar. 1990.

[16] T. Hüttner and W. Straßer. Fast footprint MIPmapping. In *Proceedings of the Eurographics/SIGGRAPH workshop on graphics hardware 1999*, pages 35–44, 1999.

[17] A. Lampa. Wie erscheint nach der Relativitätstheorie ein bewegter Stab einem ruhenden Beobachter? *Zeitschrift für Physik*, 27:138–148, 1924. In German.

[18] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH 96 Conference Proceedings*, pages 31–42, Aug. 1996.

[19] *Live Picture*. Web site: http://www.livepicture.com.

[20] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In *SIGGRAPH 95 Conference Proceedings*, pages 39–46, Aug. 1995.

[21] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. Freeman, New York, 1973.

[22] C. Møller. *The Theory of Relativity*. Clarendon Press, Oxford, second edition, 1972.

[23] R. Penrose. The apparent shape of a relativistically moving sphere. *Proceedings of the Cambridge Philosophical Society*, 55:137–139, 1959.

[24] *QGLViewer*. Web site: http://www.qglviewer.de.

[25] Quarks & Co. Die Relativitätstheorie—Einfach erklärt. Broadcast on 12 November 1999 by German TV station WDR.

[26] R. T. Rau, D. Weiskopf, and H. Ruder. Special relativity in virtual reality. In H.-C. Hege and K. Polthier, editors, *Mathematical Visualization*, pages 269–279. Springer Verlag, Heidelberg, 1998.

[27] W. Rindler. *Introduction to Special Relativity*. Clarendon Press, Oxford, second edition, 1991.

[28] G. D. Scott and H. J. van Driel. Geometrical appearances at relativistic speeds. *American Journal of Physics*, 38(8):971–977, Aug. 1970.

[29] G. D. Scott and R. R. Viner. The geomentrical appearence of large objects moving at relativistic speeds. *American Journal of Physics*, 33(7):534–536, July 1965.

[30] S. M. Seitz and C. R. Dyer. View morphing. In *SIGGRAPH 96 Conference Proceedings*, pages 21–30, Aug. 1996.

[31] Y. Sun, F. D. Fracchia, T. W. Calvert, and M. S. Drew. Deriving spectra from colors and rendering light interference. *IEEE Computer Graphics and Applications*, 19(4):61–67, July/Aug. 1999.

[32] J. Terrell. Invisibility of the Lorentz contraction. *Physical Review*, 116(4):1041–1045, Nov. 1959.

[33] D. Weiskopf. A texture mapping approach for the visualization of special relativity. In *IEEE Visualization 1999 Late Breaking Hot Topics Proceedings*, pages 41–44, 1999.

[34] D. Weiskopf, U. Kraus, and H. Ruder. Searchlight and doppler effects in the visualization of special relativity: A corrected derivation of the transformation of radiance. *ACM Transactions on Graphics*, 18(3):278–292, July 1999.

[35] V. F. Weisskopf. The visual appearance of rapidly moving objects. *Physics Today*, 13(9):24–27, 1960.

[36] L. Williams. Pyramidal parametrics. In *SIGGRAPH 83 Conference Proceedings*, pages 1–11, July 1983.

[37] M. Woo, J. Neider, T. Davis, and OpenGL Architecture Review Board. *OpenGL programming guide: the official guide to learning OpenGL, version 1.1*. Addison-Wesley, Reading, MA, USA, 1997.

[38] G. Wyszecki and W. S. Stiles. *Color Science*. John Wiley & Sons, New York, 2 edition, 1982.

A Equivalence of Exocentric and Egocentric View

One basic feature of special relativity is the absence of a single universal frame of reference and of a universal time. Any inertial frame is equally valid to describe the physical world. Therefore, an egocentric point of view (the camera is at rest and the objects are moving) and an exocentric point of view (the objects are at rest and the camera is moving) are totally equivalent.

Nevertheless, we would like to explicitly show how both points of view can be matched. Only the issues related to the geometrical appearance are discussed. The Doppler and searchlight effects are neglected because they are usually presented in a way equivalent to the exocentric point of view and thus need no further presentation.

A.1 Lorentz Transformation

Events in spacetime are described by four-vectors. A four-vector x^μ consists of one temporal and three spatial coordinates,

$$x^\mu = (x^0, x^1, x^2, x^3) = (ct, x, y, z),$$

where c is the speed of light and $\mu \in \{0, 1, 2, 3\}$.

For an observer moving with velocity $v = \beta c$ along the positive z axis, the four-vector for the same event can be calculated according to the respective Lorentz transformation. The transformed four-vector is

$$x^{\mu\prime} = (\gamma(ct - \beta z), x, y, \gamma(z - \beta ct)), \tag{8}$$

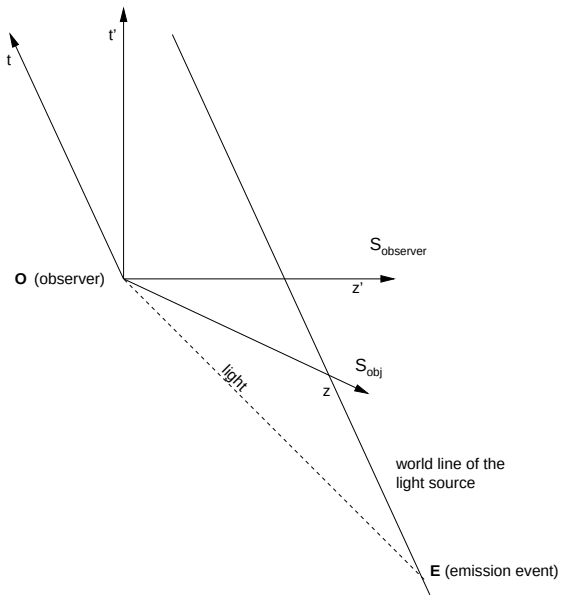

Figure 6: Minkowski diagram for egocentric view.

with $\gamma = 1/\sqrt{1 - \beta^2}$. The inverse Lorentz transformation originates from the Lorentz transformation by substituting β by $-\beta$.

The following considerations are based on relations between the observer's rest frame S_{observer} and the objects' rest frame S_{obj}. Without loss of generality let the origins of the two frame coincide at the event of image production and the observer be moving with β along the positive z axis of S_{obj}. Then, the two frames of reference are related to each other by the above Lorentz transformation. The primed terms are with respect to S_{observer}, the unprimed terms are with respect to S_{obj}.

A.2 Direction of Incoming Light

First, it will be shown that the direction of the incoming light is identical for both points of view. The light emitted by a single point-light source is considered. Figures 6 and 7 show the respective Minkowski diagrams. A Minkowski diagram is a spacetime diagram without the spatial coordinates x and y.

In S_{obj}, the event of light emission is

$$x_E^\mu = (-\sqrt{x_E^2 + y_E^2 + z_E^2}, x_E, y_E, z_E),$$

if the light source is located at the spatial position (x_E, y_E, z_E). The component x_E^0 reflects the time of flight from the emission event to the absorption event at the camera. Alternatively, the emission event can be expressed in spherical coordinates,

$$x_E^\mu = (-r_E, r_E \cos\phi_E \sin\theta_E, r_E \sin\phi_E \sin\theta_E, r_E \cos\theta_E), \tag{9}$$

with $r_E = \sqrt{x_E^2 + y_E^2 + z_E^2}$.

In S_{observer}, the emission event is obtained by the Lorentz transformation,

$$x_E^{\mu\prime} = (-\gamma(r_E + \beta z_E), x_E, y_E, \gamma(z_E + \beta r_E)).$$

The comparison to the analog of Eq. (9) in the observer's frame of

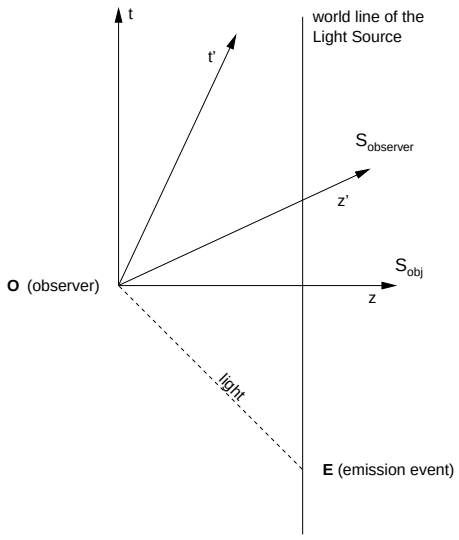

Figure 7: Minkowski diagram for exocentric view.

reference yields the transformed angles:

$$\cos\theta' = \frac{\cos\theta - \beta}{1 - \beta\cos\theta}, \quad (10)$$

$$\phi' = \phi. \quad (11)$$

Therefore, the direction of the incoming light is identical for the egocentric and the exocentric point of view which is based on the aberration equations (1) and (2).

A.3 Visibility

The next concern is the issue of visibility. If one object is hidden by another object in one frame of reference, is it as well hidden in all other frames? Are we allowed to compute visibility in S_{obj} and then use the result in S_{observer}?

Light travels along straight lines in four-dimensional spacetime with respect to every frame. Therefore, the order of emission events along a light ray is independent of the chosen frame of reference and so is the visibility property.

The explicit calculation is as follows. Let us consider two emission events E_1 and E_2. In S_{obj}, let E_2 be hidden by E_1. The respective coordinates are related by

$$x^\mu_{E_2} = a\, x^\mu_{E_1}, \quad (12)$$

with a constant $a > 1$. With the use of the Lorentz transformation (8), the coordinates of the emission events in S_{observer} follow:

$$x^\mu_{E_2}{}' = a\, x^\mu_{E_1}{}'.$$

Combined with the fact that the aberration formulae (10) and (11) are invertible, the invariance of visibility under Lorentz transformations is proven. Since the inverse Lorentz transformation is just a Lorentz transformation with opposite direction of motion, the invariance of invisibility is valid for the transformation from S_{observer} to S_{obj} as well.

After all these abstract calculations, how can it be explained that we are able to look "around" relativistic objects and see their back? Figures 8 and 9 clarify the situation for the example of a moving cube. In the egocentric view, Fig. 8, the camera is at rest and the cube is moving with $v = 0.7c$ to the left. The cube is Lorentz-contracted along the direction of motion. Here, the back of the cube becomes visible because the cube outruns the light rays emitted from its back. In the exocentric view, Fig. 9, the cube is at rest and the camera is moving with $v = 0.7c$ to the right. Here, the back of the cube becomes visible because the observer is behind the cube when she or he is registering the incoming light. The rightmost image in Fig. 9 also illustrates the direction of the light ray based on aberration in order to allow a direct comparison to the egocentric view.

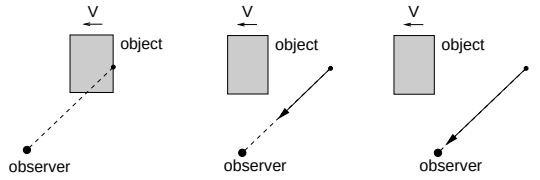

Figure 8: Egocentric view, three snapshots for $\beta = 0.7$.

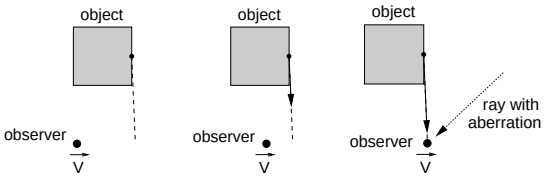

Figure 9: Exocentric view, three snapshots for $\beta = 0.7$.

A.4 Apparent Rotation

The apparent rotation of fast moving objects is closely related to the explanation in the previous section. In the egocentric view, an object seems to be rotated because light emitted from the normally invisible back of the object is outrun by the object and can thus reach the observer. In the exocentric view, the observer is already behind the object and can thus see its back. However, from the observer's point of view the object seems to still be ahead because of the aberration of the incoming light. Seeing the back side of an object is interpreted as an apparent rotation of the object.

A.5 Summary

We have shown that the egocentric and exocentric view are equivalent with respect to the apparent geometry in special relativistic rendering. The main difficulty in matching both views is to transform all physical components of the system, in particular the position of the observer and the objects.

Usually, the egocentric view is regarded more natural and hence is a widely used model of explanation. In fact, we deem the exocentric view to be more appropriate for the following reasons. First, the exocentric view allows for the transformation of all relevant information about the light field in one operation, as described in Sect. 4. Secondly, an accelerated motion of the observer can be incorporated in the exocentric view without any modification, cf. [26]. Thirdly, the exocentric view better reflects the physical reality. There is no direct interaction between the observer and remote objects. All the information about the surrounding environment is carried to the observer via photons. The generation of a snapshot is based on a local interaction between the photons and the detector (camera). Therefore, it is closer to physics to transform the photons which have reached the observer than to transform emission events far away from the camera.

Figure 10: Flight through Yosemite Valley with $\beta = 0.95$.

(a) (b) (c)

Figure 11: Image (a) shows the non-relativistic view, image (b) the apparent geometry for $\beta = 0.9$, and image (c) completely relativistic rendering with $\beta = 0.2$. The overall intensity in (c) is reduced to 10 percent of that in (a) and (b) to avoid extreme clamping of the final RGB values.

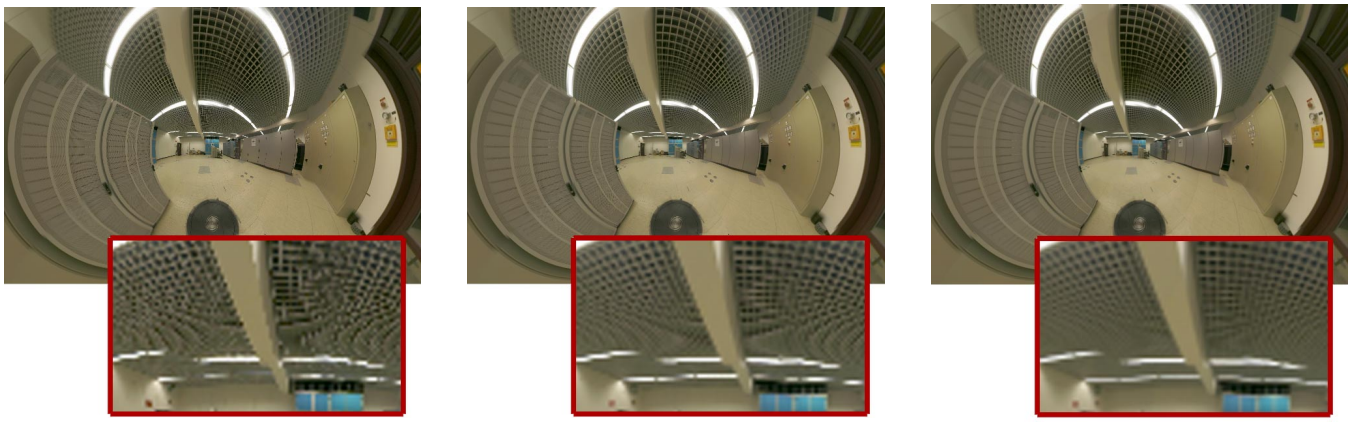

(a) (b) (c)

Figure 12: Comparison of filtering and supersampling techniques for the visualization of apparent geometry at $\beta = 0.99$. Image (a) is rendered without filtering and supersampling, image (b) illustrates filtering with rectangular footprint, and image (c) documents the combination of filtering and 2×2 supersampling.