

Fast Rendering of Relativistic Objects

christopher betts

*Department of Computer Science, Monash University, Wellington Rd., Clayton, Vic. 3168,
Australia (email: cbetts@bruce.cs.monash.edu.au)*

SUMMARY

This paper presents a number of simple techniques that allow very fast rendering of multiple independently moving objects under the conditions of special relativity, as well as a brief tutorial in the theory of special relativity for non-experts. In addition to geometric rendering, a simplified lighting model is presented, which allows the apparent colour changes of objects to be observed, along with changes in the apparent direction of illumination.

The geometric rendering can be summarized as a two-stage process. In the first, the Lorentz contraction is applied to each object in the observer's reference frame, and in the second the apparent visual position is found for every object. Once the visual geometry is found, the coloration of each object is handled, with a traditional graphics engine such as a z-buffer performing the final rendering.

Rendering speeds in excess of 30,000 polygons a second were obtained on a 133 MHz Silicon Graphics Indigo2 workstation. © 1998 John Wiley & Sons, Ltd.

key words: relativity; simulation; animation

INTRODUCTION

Relativity theory is closely linked to the behaviour of light, and owes much of its origin to consideration of how light propagates. Experiments with light, such as the early Michelson–Morley experiments, which failed to find the ‘ether’ in which light was thought to propagate, were early pointers towards relativity. Einstein’s ‘thought experiments’ on the behaviour of light, such as what his reflection in a mirror would look like were he travelling at the speed of light, and others, were the building blocks of his theory.

Given this early interest, it is surprising that over fifty years elapsed between the time that Einstein first published his seminal paper, and the time that some of the more interesting visual effects of special relativity were revealed. Although the Lorentz contraction was an integral part of special relativity from the beginning, it was not until 1959 that James Terrell pointed out¹ that this contraction could be quite difficult to observe visually, and in some cases would be seen as a simple rotation of the object. At the same time Roger Penrose was pointing out other effects, such as the fact that the outline of a relativistically moving sphere was always seen to be circular.² To discover why, we need to consider both special relativity, and the normal process of viewing.

A thumbnail sketch of special relativity

We must assume that the reader has some familiarity with special relativity: nevertheless, a quick review may not go astray.

Special relativity is based on two postulates:

1. The laws of nature appear the same to all observers, as long as they are not accelerating.
2. All observers will measure the speed of light in a vacuum to be the same (3×10^8 m/s).

From these postulates, we can derive a number of corollaries, among which are:

- (a) Moving objects, as measured by a stationary observer, appear to contract in their direction of motion (by a factor γ^{-1}).
- (b) Clocks on moving objects, as measured by a stationary observer, appear to slow down (by a factor γ).

Here $\gamma = 1/\sqrt{1 - v^2/C^2}$, v being the speed of the object relative to the observer, and C being the speed of light.

These equations tell us that, if we measure a moving object using synchronized, stationary clocks at known distances, the moving object will appear to have shrunk. In fact, if it is travelling very close to the speed of light, it will appear to have flattened a great deal—spheres become pancakes, a spaceship appears to become a sheet of tin foil. To someone moving at the same velocity as the moving object, however, everything appears completely normal. This relativistic shortening is called ‘Lorentz contraction’.

Not only are there these geometric distortions, but the rate at which time passes on the moving object appears to decrease. This means that if we measure a moving clock, again using an array of stationary clocks at known distances, it will appear to be running slow. This also means that lights on moving objects, after taking account of normal Doppler effects, will appear to be shifted towards the red end of the spectrum, as light sources on the moving object appear to oscillate more slowly, producing lower frequency photons (‘red-shifting’ their spectrum). This slowing of moving clocks is usually known as ‘time dilation’.

As a final note, it is important to remember that special relativity deals with objects with constant velocity, and does not deal with accelerating objects; such motion normally coming under the purview of the more complex general theory of relativity.

Viewing relativistic objects

These relativistic effects, Lorentz contraction and time dilation, are interesting enough in themselves. Curiously, though, they are difficult to observe: if viewed directly (say through a very powerful telescope, or by way of a radio beacon) a relativistically moving clock may actually be seen to be running fast—and rather

* There are a number of good introductory relativity texts available, such as those by Schwartz and McGuinness,³ H. Bondi,⁴ and Rindler.⁵ An excellent plain-language introduction (possibly out of print) is by Bertrand Russell.⁶ Brief explanations can also be found in many encyclopedias, e.g. *Everyman's Encyclopedia*.⁷

than being red-shifted, may actually be blue-shifted, with light emitted from the clock having a higher frequency than normal. In addition, as mentioned before, objects that have been Lorentz contracted may appear to be merely rotated, or they may appear skewed or sheared—but they do not appear to be shrunk.

The reasons for these effects, and for the difficulty in directly viewing the Lorentz contraction, is the finite speed of light. Since relativistic objects are moving at a significant fraction of the speed of light, Doppler effects become important. Just as a fire engine's siren increase in pitch when the vehicle is approaching, light from a rapidly approaching object increases in frequency, by more than time dilation causes it to decrease.

Difficulties also occur because of the lag, caused again by the finite speed of light, between an event occurring, and our observation of it. For example, consider what happens when we try to observe an approaching clock. Imagine that the clock suddenly appears one 'light hour' away from us, showing twelve noon, and approaching at 80 per cent of the speed of light. It will take one hour and fifteen minutes for the clock to reach us. However, it will take one hour for the light, showing us the initial appearance of the clock, to travel to our position. During the one hour and fifteen minute flight, the clock *should* (in a non-relativistic universe) advance to a quarter past one. Owing to time dilation, however, it will in fact only advance forty-five minutes—so when the clock reaches us, it will show a quarter to one.

So, at our first view of the clock (an hour after it appears) the clock will show twelve noon, and a bare fifteen minutes later the clock will reach us, showing a quarter to one. To our eyes, it will appear that the clock has advanced forty five minutes in a quarter of an hour, and be running fast by a factor of three. Also, the clock will appear brighter, since the rate at which it appears to be producing photons will likewise increase by three. And finally, all those photons will have their frequencies tripled.

Similar difficulties occur when we try to visually observe the shape of a rapidly moving object. Since light from the back of an object takes longer to reach us than light from the front, we see the back of an object at an earlier position, further back along the object's path. As a result, the object appears to be deformed by a non-linear shear, which, if the object is sufficiently distant, appears as a simple rotation of the object, with no contraction at all.

Animating relativistic scenes

Animating the geometry of scenes with relativistically moving objects is challenging because of the difficulty in determining quickly the positions of the (sometimes severely distorted) visual images, owing to the affects mentioned above. Once this is achieved, however, simple rendering of the scene can be done using a z-buffer or some similar fast-rendering technique.

Aside from the distorted shapes of relativistically moving objects, the other main visual effects in special relativity are the changes in apparent colour, intensity, and the direction of illumination that an object undergoes as it moves past the observer.* The Doppler effect and time dilation cause the object to appear blue-shifted or red-shifted, and the object's motion may cause its apparent luminosity to change by

*These effects are described in most relativity text books, e.g. those by Lawden⁸ and Rindler.⁵

several orders of magnitude. The direction from which the object appears to be illuminated may also change radically from that of other, non-moving objects, or objects moving in different directions.

PREVIOUS WORK IN THE FIELD

The visual geometry of objects has been dealt with mathematically by a number of authors. Terrell published a seminal article¹ in 1959 showing that in the limiting case of a distant object subtending a small angle, the Lorentz contraction is indistinguishable from rotation. Other authors have also shown mathematically a number of interesting effects, such as the distortion of the surface of a sphere,⁹ the appearances of cubes,¹⁰ and the actual three dimensional geometry of objects as perceived by the observer, either by some form of direct distance measurement¹¹ or by stereoscopic vision.¹²

There have been a number of graphical computer simulations of relativistic phenomena, dating back to the early seventies. Scott and van Dreil¹³ used computers in 1970 to calculate the appearance of starfields and spherical surface distortions. More recently, full relativistic ray tracers have been written by people such as Nguyen,¹⁴ and Hsiung¹⁵. These completely model a scene under the conditions of special relativity, but at great computational cost.

Nguyen has also written a faster program that directly calculates the appearance of a scene, every element of which is moving at the same velocity, to an observer.¹⁴ This method turns out to be quite similar mathematically to the method used in this paper, although it had not been designed for real time animation of multiple objects. Another fast program, using a z-buffer-like construct called a 'time-buffer' or 't-buffer' has also been written by Hsiung.¹⁶ In the t-buffer, all the events visible to the observer were loaded into the buffer, where more recent events replaced (occulted) older events. The t-buffer is then displayed on screen, each pixel displaying the most recent event.

The basic equations governing the behaviour of light, and light emission, under relativistic conditions are well known, and available in most standard relativity textbooks.^{5,8} Modelling this behaviour, and translating this to a computer RGB display is not straightforward though. This paper presents a method for quickly rendering lighting effects by greatly restricting the properties of the illuminants; however, far more general and accurate, though slower, methods have also been created.^{14,17}

This simulation works by simplifying the mathematics of the lighting, pre-calculating a large amount of scene information, and taking full advantage of an existing z-buffer, present in most graphics workstations, to make real-time animation possible. In addition, by breaking the simulation up into a two-stage process, it becomes possible to completely simulate any special-relativistic scene that can be described in terms of a number of independently moving objects at no extra computational cost.

MODELLING THE VISUAL GEOMETRY

The apparent shape of a relativistically moving object can be modelled by considering two simple effects: the Lorentz contraction, and the non-linear shear distortion caused by the finite speed of light. It should be emphasized that this second distortion has

nothing at all to do with relativity, and is simply due to the observer receiving delayed images.

The Lorentz contraction

Objects moving relativistically are (in the observer's frame of reference) Lorentz contracted in the direction of motion (as illustrated in Figure 1). Provided that the object's velocity is constant, this contraction need only be calculated once, at the beginning of the simulation. This is easily done by scaling each point of the object by γ^{-1} in the direction of motion where as usual:

$$\gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}} \quad (1)$$

Objects are Lorentz contracted at the start of the simulation, and then placed in their desired locations in the observer's reference frame. This is equivalent to locating the origin of the object's reference frame at a chosen point in the observer's reference frame.

Modelling the finite speed of light

If we know the position of a point on an object in the observer's frame, and the velocity of that point, it is possible to find the apparent position of the point, as it appears to the observer.

To determine the apparent position of a moving particle, we backtrack along the particle's path until we find a point at which the physical distance from the observer is equal to the temporal distance times the speed of light. In the language of relativity, we are finding the intersection of the particle's world line with the observer's past light cone (see Figure 2).

As a simple example, if a point was two light-seconds away from an observer, and travelling directly towards the observer at half the speed of light, the observer would see the point as being four light-seconds away. The observer would thus be seeing a four-second old view of the particle.

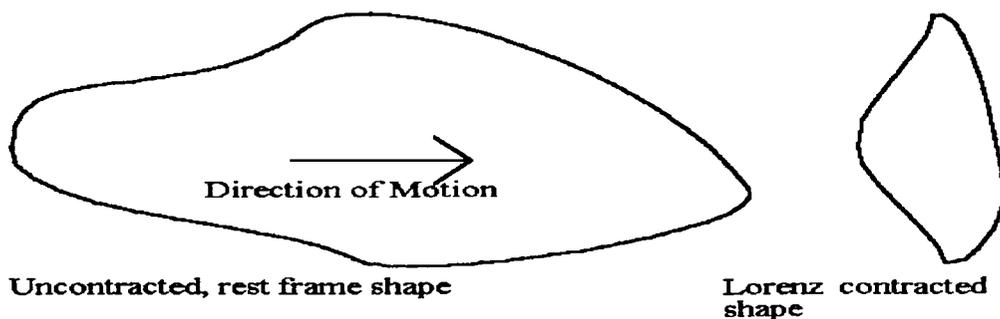


Figure 1. Lorentz contraction

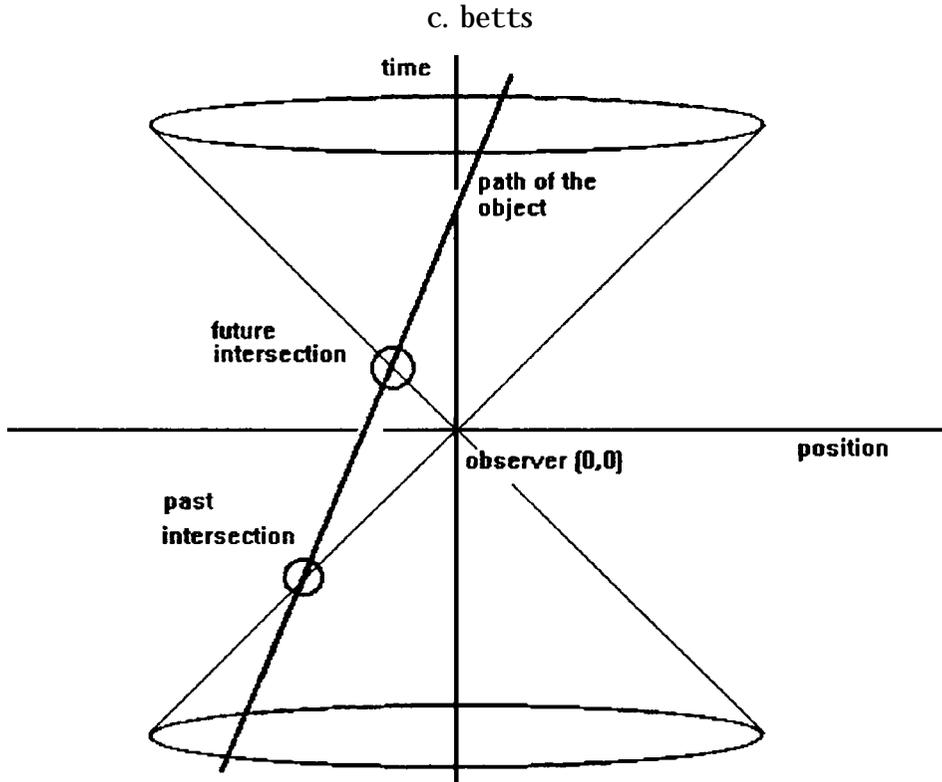


Figure 2. Space-time diagram of object and observer's light cone

Expressing this mathematically, we have the path of the object as

$$\mathbf{P}(t) = \mathbf{P}(0) + \mathbf{v}t \quad (2)$$

where $\mathbf{P}(t)$ is the position of the point at time t , $\mathbf{P}(0)$ is the position of the point at $t=0$, \mathbf{v} is the velocity of the point and t is time (all quantities refer to measurements in the observer's frame).

For a point event to be visible from the observer's position, the distance to the point must be equal to the elapsed time multiplied by the speed of light. In normalized co-ordinates, where the speed of light is 1, we have:

$$|\mathbf{x}| = |\Delta t| \quad (3)$$

where \mathbf{x} is the distance to the point event from the observer and Δt is the time from the event to the present.

In the language of relativistic physics, this is the same as saying that the interval between the observer and the observed event is null. The combination of equations (2) and (3) gives us

$$(\mathbf{P}(0) + \mathbf{v} \Delta t)^2 = \Delta t^2 \quad (4)$$

Explicitly including the vector components, and rearranging, we obtain a quadratic equation in t :

$$\Sigma P_i^2 + 2\Sigma P_i v_i \Delta t + (\Sigma v_i^2 - 1)\Delta t^2 = 0 \quad (5)$$

where P_i are the vector components of $\mathbf{P}(0)$, the position of the point in the observer's frame at $t=0$, and v_i are the vector components of \mathbf{v} , the velocity of the point in the observer's frame.

This equation has two solutions, corresponding to past and future intersections of the observer's light cone and the object's world line. As the observer cannot see into the future, we are only interested in past solutions, so we find

$$\Delta t = \frac{-\Sigma P_i v_i + \sqrt{\Sigma P_i^2 v_i^2 - (\Sigma v_i^2 - 1)(\Sigma P_i^2)}}{(\Sigma v_i^2 - 1)} \quad (6)$$

Using this equation, our knowledge of the point's actual position at $t=0$, and its velocity, we can find the point's apparent position to the observer (at $t=0$) as simply

$$\mathbf{x}_{\text{apparent}} = \mathbf{P}(0) + \mathbf{v}\Delta t \quad (7)$$

(compare with equation 2). This is easily expanded for all t to

$$\mathbf{x}_{\text{apparent}} = \mathbf{P}(t) + \mathbf{v}\Delta t \quad (8)$$

where Δt is calculated from equation (6), taking P_i as the current position of the point $\mathbf{P}(t)$.

By making this transformation for every point making up the object, we find the apparent visual position of the object in the observer's coordinate frame at any time t . In addition we note that since more recent visible events appear closer than older visible events, we can use any rendering system where closer points occult more distant points. In the program described below, a z-buffer was used.

MODELLING THE COLOUR AND LIGHTING OF THE OBJECT

The changes in apparent colour and illumination that occur for a rapidly moving object can be rather laborious to model, since to be completely accurate a large portion of the spectrum of both the object and the illuminant must be considered. A number of simplifying assumptions can be made, however, which result in a reasonably swift and straightforward method of dealing with these effects, while sacrificing a more rigorous colour and reflectance model.

Illuminants

The easiest way to handle illuminants would be to avoid them altogether, and make all objects self-luminous. This has the advantage of simplicity, and removes the necessity of modelling in detail the spectrum of the illuminant. Unfortunately, it

also removes the opportunity to model the interesting effect of the changing direction of apparent illumination for a rapidly moving object.

However, if we make some artificial assumptions about our illuminant's spectrum, we can make the objects in effect be self-luminous, but retain the shading effects of the illuminant, along with the apparent change in the direction of illumination. From the point of view of our computer model, we simply keep the illuminant's intensity constant through the simulation. Some readers may be content with this, and may wish to pass over the next section, which describes in more detail our illuminant's spectrum.

A mathematically convenient spectral model for illuminants

Modelling our illuminants becomes easier if they illuminate the relativistically moving objects at a constant intensity, regardless of the object's speed relative to the illuminant. Although it might appear unlikely that any real illuminant would have this property, it turns out that a class of physical illuminants approximate it.

Let us consider an illuminant with a power output distribution of the inverse fourth-power of wavelength ($P \approx 1/\lambda^4$). This is the Rayleigh–Jeans distribution of classical thermodynamics, and is an approximation of the actual spectral distribution of classical thermodynamics, and is an approximation of the actual spectral distribution of black-body radiators at long wavelengths. Since photon energy is proportional to the inverse square of wavelength ($1/\lambda^2$), this leaves us with a photon flux also proportional to the inverse square of wavelength:

$$\text{Rayleigh–Jeans distribution: No. of photons} \propto 1/\lambda^2 \quad (9)$$

Considering the range of frequencies the observer can see (their 'visual spectrum'), we note two effects for light illuminating an object moving relative to such an illuminant. First, when the spectrum is red or blue shifted (by non-relativistic Doppler shifting), the range of the original spectrum that is shifted to the visual region varies directly with wavelength ($\propto \lambda$). For example, if the object is moving towards the illuminant at half the speed of light, the frequency of the incoming light with regard to the object will be doubled, so the span of wavelengths of the original spectrum that are shifted into the visible region will be twice the span of wavelengths visible were the object stationary. Since the span increases, so does the photon flux, so we can say that the photon flux is proportional to wavelength:

$$\text{Non-relativistic Doppler shifting: No. of photons} \propto \lambda \quad (10)$$

Secondly, when the object is moving relative to the illuminant, it receives more or less light depending on whether it moves towards or away from the light source. An object moving towards an illuminant at a fractional speed of light α (or away from it, if α is negative) will receive $(1 + \alpha)$ as many photons per unit time as otherwise. The wavelength of the received light will also change by Doppler shifting by the same factor $(1 + \alpha)$. Since both wavelength and photon flux are changing at the same rate, we can express the change in photon flux due to relative motion between object and illuminant as being directly proportional to the change in wavelength.

Intensity change due to motion between object and illuminant: No. of photons $\propto \lambda$
(11)

The net effect of the photon flux, varying as the inverse square of wavelength, the non-relativistic Doppler shift, which increases visible photons in proportion to the wavelength viewed, and the increase of visible photons in proportion to wavelength due to the object's motion (equations (9), (10) and (11)), is that the object will appear to an observer to be illuminated at a constant intensity by an illuminant with a Rayleigh–Jeans distribution.

Total photon flux (in R–J distribution) for a moving object: No. of photons = constant
(12)

The Rayleigh–Jeans distribution is physically impossible, however, as it has an infinite energy. However, it is approximated at long wavelengths by black body radiators, so it is not completely unrealistic. In any event, it is a useful simplification for our simulation.

The direction of the illuminant

Since our objects are usually moving at an appreciable fraction of the speed of light, they will often appear to be illuminated in a different direction to that of the illuminant as seen by the observer, as illustrated in Figure 3.^{18*} This effect is similar to the fall of rain on a car windscreen: when the car is stationary, the direction of the rain may be from behind the car and miss the windscreen altogether, but once the car starts moving, the car 'catches' the rain falling ahead of it, and to the driver

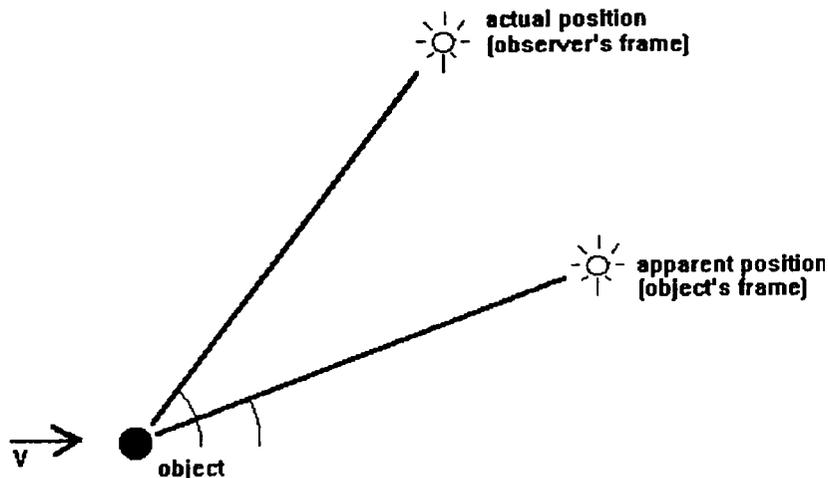


Figure 3. Apparent change in direction of the illuminant

*Explanations of these effects can be found in most special relativity text books, e.g. that of Rindler, (References, pp 39–48).

it appears that the rain is now falling from the front. For simplicity, we assume an infinitely distant illuminant. To obtain the apparent direction of the illuminant, from the object (i.e. in the object's reference frame) we add velocities to find

$$\alpha' = \cos^{-1} \left(\frac{\cos(\alpha) + \frac{v}{c}}{1 + \cos(\alpha) \frac{v}{c}} \right) \quad (13)$$

where α is the angle between the illuminant and the object's direction of motion, relative to the observer, α' is the same angle, relative to the object, v is the object's velocity and c is the speed of light.

Having found the direction of the illuminant in the object's reference frame, we can shade the object at the start of the simulation (before performing the Lorentz contraction) using our shading model of choice. For simplicity, a basic 'cosine law diffuse reflection' model was used, in which the lighting intensity at any point of the object is equal to the dot product of the illumination vector and the surface normal. This model is applied to compute the colour and intensity (which is equivalent from our point of view to a self-luminous radiance) at every point on the object.

Relativistic Doppler shifting

The relativistic Doppler shifting of light produced by (or reflected from) an object is easily calculated using standard equations.⁵ Using the illuminant model outlined earlier, there is no difference between light produced by the object and light reflected from our illuminants, so we need only concern ourselves with a single Doppler shift of light coming from the object to the observer. A slightly more elaborate model with a different illuminant would obviously need to explicitly model the Doppler shift between illuminant and object; however, as long as the illuminant was still considered to be an infinite distance away, this would only need to be done once, at the beginning of the simulation, and would not affect the speed of animation.

The relativistic Doppler shift is caused by two factors. First, there is a normal Doppler shift caused by the source moving (analogous, for example, to the audible Doppler shift produced by the siren of fire-engines). Secondly, there is a relativistic frequency shift caused by time dilation, which causes the emitter, as seen by the observer, to slow down (analogous to someone slowing down the fire-engine's siren by reducing the speed of the oscillator). The Doppler shift equation is

$$\lambda' = \lambda \frac{1 + \frac{u_r}{c}}{\sqrt{1 - \frac{u^2}{c^2}}} \quad (14)$$

where λ is the initial wavelength, λ' is the final wavelength, u is the speed and u_r

is the radial velocity of the object relative to the observer (positive for receding, negative for approaching).

Note that the above equation correctly implies that when the object is passing an observer in an entirely transverse fashion (i.e. $u_r = 0$), the Doppler shift seen by the observer is entirely caused by the temporal dilation of the emitting object.

Modelling colour

It is usual in computer graphics to model an object's colour as a triplet of red, green and blue values, as these are the only colours displayed by most monitors. Relativistically moving objects will often have their spectra shifted such that the infra-red or ultra-violet portions, or other wavelengths such as radio waves or gamma rays, become frequency shifted to appear as light in the visual region of the spectrum. In order to properly model such objects, their full electromagnetic spectrum should therefore be simulated.

Interesting visual effects demonstrating red and blue shifting can be obtained using simpler models; however, by assuming that the object has only its normal visible colours, and is otherwise completely black, or that it has only a faint reflectance in the non-visible spectrum. Alternatively, the usual $\langle r, g, b \rangle$ triplet of values can be expanded to a (slightly artificial) quintuplet of $\langle \text{infra-red } (i), \text{ red, green, blue, ultra-violet } (u) \rangle$, with the exact wavelength of the i, r , and u, v values being chosen for convenience, and spectrum assumed to be entirely black, or very dark, outside these values (This was the method chosen in the current model).

Regardless of what method is chosen, the object's spectrum will have to be reduced at some stage to an $\langle r, g, b \rangle$ triplet. One straightforward method of doing this is to assign values to the frequencies of the red, green and blue values displayed by the monitor, and simply scale the object's spectrum to these values. Using the modelling schemes above, with either an initial $\langle r, g, b \rangle$ triplet or an $\langle i, r, g, b, u \rangle$ quintuplet, one can easily approximate by using linear interpolation between the final (frequency shifted) values and the actual wavelengths of the display $\langle r, g, b \rangle$ triple.

OUTLINE OF THE COMPLETE RENDERING ALGORITHM

Initialization

1. Set the starting colours of the objects.
2. Set the starting directions and strengths of illuminants.
3. Set the starting positions of the objects.
4. Compute the initial shading of the objects, using equation (13) to determine the direction of the illuminant.
5. Contract the object according to the Lorentz contraction (1) and place it in the scene (the observer's reference frame).

For every frame

1. For each point on each object, with the point's position and velocity, using equation (6), find the visual position of the point.

2. For each point of each object, with the velocity, visual position and radiance of each point, calculate the point's colour (equation (14)).
3. Using the visual position and colour information from above, use a z-buffer or similar to render the frame.
4. Repeat from 'For every frame'.

EXPERIMENTAL RESULTS

The method described was successfully used to simulate a large number of independently moving objects. Using a simple diffuse lighting model, and a (somewhat artificial) modelling of objects as (i, r, g, b, u) quintuplets, produced a number of reasonably fast animations of relativistic scenes. Running on a Silicon Graphics 4D – 310 GTX workstation (with 64Mb of memory running at 33 MHz) the simulation ran at roughly 12,000 polygons/second, whereas if only the geometry was modelled (i.e. the Doppler shift of the object was ignored) the simulation ran at 15,000 polygons/second. On a Silicon Graphics Indigo2 workstation with a 32MB RAM running at 133MHz, full colour animation reached 31,500 shaded polygons/second. Animation speeds vary slightly, since scenes with large numbers of objects stationary in the observer's frame are slightly faster to render.

Plates 1 to 6 show a variety of scenes incorporating relativistic motion, and displaying some common relativistic effects. Note that there may be some lens distortion towards the edge of these images, as Plates 3 to 6 assume a 90° field of view.

Plate 1 shows a non-relativistic view of two spaceships as they would appear at rest.

Plate 2 shows the same spaceships moving rapidly towards and to the right of the observer. The front spaceship moves at $0.8c$, while the rear spaceship moves at $0.99c$. The colour spectrum shift has been reduced by a factor of two to allow us to view the colour gradations across the front spaceship, however the rear spaceship is entirely red-shifted out of the visible spectrum. We are still able to see the rear ship because of the spectrum of the illuminant. Notice also the extreme 'Terrell' rotation of the rear ship, due to its great speed.

Plate 3 shows a stationary view of a city. Notice particularly the green trees on the right hand side, and the general rectangular shape of the buildings and of the square tiled roads.

Plate 4 shows the city as we travel through it at $0.9c$. Since these speeds cause a blueshift beyond the modelled spectrum in this simulation, this image are shown in false colour—the frequency shift has been compressed by a factor of 2 (as above). Notice the distortion of the trees, and how the city seems to 'stretch' before us, with the buildings appearing further away, and curved.*

Plate 5 shows a collection of objects on a tiled board. There are a number of spheres, and a number of rectangular blocks.

Plate 6 shows the same collection in motion. All of them have different velocity vectors. Notice how the spheres intersect the tiled board with a non-circular cross section (due to Lorentz contraction) even though they still appear circular. (An interesting result in special relativity is that a sphere will always appear to have a circular outline, no matter how quickly it is moving in an observer's frame.) This

*It is interesting to compare these images with the analytically-derived drawings in Reference 18.

picture is framed by two of the blocks which are transversely passing the observer. Once again, the image is in false colour, as the frequency shift has been halved to make the objects visible.

DISCUSSION

Accuracy

In order to produce a realtime simulation, we have had to make a number of approximations. Firstly, by using a polygon based system rather than a pixel based system (such as in a ray-tracer) we introduce geometrical inaccuracies. Straight lines should appear as hyperbolae, however polygon edges will always be straight lines. This is the same difficulty all polygon-based systems experience when modelling curved surfaces, and the solution is the same—we can increase the accuracy of the simulation to any arbitrary level by simply increasing the number of polygons; however, this will slow down the simulation.

Aside from this, the simulation is quite general. It handles any collection of independently moving objects under the conditions of special relativity, as long as all parts of each object are moving with the same constant velocity. Currently the simulation does not directly handle deforming objects, such as exploding stars and so on, which can however be modelled as a large number of component objects or individual polygons. The techniques described here are easily expanded to include individual velocities for all component points.

The main area where liberties are taken lies in the illumination model. Purists may prefer to consider the illumination model as simply making the objects self-luminous, in which case the physically unlikely infinitely distant Rayleigh–Jeans illuminants can be ignored. The current methods does, however, show a number of interesting effects, such as shading on relativistic objects, and the apparent change in the direction of illumination of such objects.

Extensions

A number of simple extensions are possible. More sophisticated pictures can be produced by using a more advanced illumination model (e.g. a specular reflection model), rather than simple diffuse reflection, as well as modelling more accurately the spectra of both the object and the illuminant. This would require modelling the frequency shift and intensity change of light passing from the illuminant to the object, as has been done by a number of other authors in their ray-tracing models.

Another possible extension is shadow handling. This could be done by moving, in addition to the objects, shadow regions. The shapes of such regions could be calculated before the start of the simulation, and treated in the same way as objects are for the purpose of finding their apparent visual shape, and hence the objects apparently in shadow at any given moment could be found.

It is possible to model images with rotating parts. Although rotation comes under the purview of general relativity, if the parts being rotated are not physical, but rather some pattern of localized surface features of an object, it can be simulated under the laws of special relativity. An example of this would be a (very large)

clock with an LCD display, each cell of which was independently timed to produce the illusion of hands moving across the clock's face. Unfortunately an analytical solution to the intersection of a rotating point with the observer's light cone does not seem possible; however, a numerical approach is quite feasible. Preliminary results show clocks with highly curved hands, the minute hand, for example, pointing to the three at the centre of the clock, but ending up showing twelve at the rim. In extreme cases, where the apparent velocity of the tip of the minute hand exceeds the speed of light, multiple (bent) minute hands can be observed.

SOFTWARE AVAILABILITY

Both the program (running on most silicon graphics machines) and the source code, along with example animation files (gif89a), are freely available, either by emailing the author at cbetts@bruce.cs.monash.edu.au or through the author's web home page <http://www.cs.monash.edu.au/~cbetts/>. Also on the author's home page are some sample images.

CONCLUSION

A simple 'visual lag' transformation allows the use of traditional graphics engines (such as the z-buffer) to be used to swiftly simulate the visual appearance of objects moving under the conditions of special relativity. Firstly the simulation must be initialized, with the initial illumination effects and the Lorentz contraction being calculated before animation begins. Once this first step has been completed the additional overhead for the frame-by-frame relativistic calculations, although significant, is not debilitating, and still allows swift rendering, especially where there is existing built-in graphics hardware. Fast animation of relativistic effects, at real-time speeds for small objects, can thus be readily performed.

acknowledgements

I would like to thank Professor Chris Wallace for supervising the work this article is based on, and for his assistance in the preparation of this article. I would also like to thank Dr. Damian Conway for his encouragement, proofreading, programming assistance and extensive suggestions.

REFERENCES

1. J. Terrell, 'Invisibility of the Lorentz contraction', *Physical Review*, **116**, (4), 1041 (1959).
2. R. Penrose, 'The apparent shape of a relativistically moving sphere', *Proceedings of the Cambridge Philosophical Society*, **55**, 137 (1959)—cited in Reference 9.
3. J. Schwartz and M. McGuinness, *Einstein for Beginners*, Writers and Readers Publishing Cooperative, Ltd., London, 1979.
4. H. Bondi, *Relativity and Common Sense*, Dover Publishing, New York, 1962.
5. W. Rindler, *Introduction to Special Relativity*, (2nd edn) Oxford University Press, New York, 1991.
6. Bertrand Russell, *The ABC of Relativity*, (3rd edn) Allen & Unwin, London, 1969.
7. 'Relativity', *Everyman's Encyclopedia*, sixth edn J. M. Dent and Sons, London, 1978.
8. D. F. Lawden, *An Introduction to Tensor Calculus, Relativity and Cosmology*, (3rd edn, reprint) Chapman and Hall Ltd., London, 1986.
9. G. D. Scott and H. J. Van Dreil, 'Geometrical appearances at relativistic speeds', *American Journal of Physics*, **38**, (8), 971 (1970).

10. F. R. Hickey, 'Two-dimensional appearance of a relativistic cube', *American Journal of Physics*, **47**, (8), 711 (1979).
11. K. G. Suffern, 'The apparent shape of a rapidly moving sphere', *American Journal of Physics*, **56**, (8), 728 (1979).
12. M. L. Boas, R. C. Calhoun and O. Joran, 'Binocular observation of moving objects', *American Journal of Physics*, **39**, (7), 782 (1971).
13. G. D. Scott and H. J. van Dreil, 'Geometric appearances at relativistic speeds', *American Journal of Physics*, **38**, (8), 971 (1970).
14. Huu Phu Nguyen, 'Implementation of two relativistic renderers', Master's Project Report under Direction of Professor Carlo H. Sequin, *Report No. UCB/CSD 88/481*, Computer Science Division, Dept. of Electrical Engineering and Computer Sciences, University of California, Berkeley, 1988.
15. Ping-Kang Hsiung and Robert H. P. Dunn, 'Visualising relativistic effects in spacetime', *Proceedings of the Supercomputing '89 Conference*, 13–17 November, 1989—cited in Reference 16.
16. Ping-Kang Hsiung, Robert H. Thibadeau and Michael Wu, 'T-buffer: fast visualization of relativistic effects in spacetime', *Computer Graphics*, **24**, (2), 83 (1990).
17. H. Blatter and T. Greber, 'Aberration and Doppler shift: an uncommon way to relativity', *American Journal of Physics*, **56**, (4), 333 (1988).