An Open and Shut Case

Andrew Glassner

Microsoft Research

Click. Whirrrr.

You press the button on your camera, the shutter opens, the film is exposed, the shutter closes, and you have a picture. If anything moved, including you, the image will show *motion blur*.

Motion blur is important. Have you looked at the classic 1933 film "King Kong" recently? The ape doesn't look quite alive in many of the scenes. It's not because of the animation—Kong moves in a wonderful, lifelike way that befits his great size and weight. The problem is that for each frame of the film, Kong was posed by hand, and photographed in that static, rigid pose. Then he was moved to the next position, and photographed again, over and over, once for each and every frame. In each image, Kong was frozen in position.

In modern films, we're used to seeing motion blur: the visible result of combining a finite exposure time with apparent motion in the scene. The most obvious source of motion blur is objects moving in the field of view—for example, a horse running past the camera. But many other sources of motion blur exist. An object out of view can cast a moving shadow, or a mirror out of view can create a moving caustic. A still object undergoing an internal transformation can change its surface appearance, causing changes in color, texture, and motion of highlights. A moving light can cast moving shadows. Even if every object in the scene remains completely static, if the camera moves over the course of a frame, we'll see motion blur as a result.

When we create synthetic images with computer graphics, we have to account for motion blur for three reasons. First, we're used to it from physical motion pictures—audiences have seen so many films shot with real cameras that almost everyone's an expert on how things ought to look when they're moving. When the motion blur looks wrong (or worse, is missing entirely), audiences know it, even if they can't articulate why.

Second, without motion blur a film can suffer from *strobing*. This psychophysical phenomenon has as much to do with the human visual system as film exposures. When you pay attention to a moving object, you unconsciously track its motion. In fact, your visual system predicts where it's going to go based on its past and current speed and velocity. If the object doesn't travel as expected, the visual system has to abandon its prediction and replace it with the real position. If this failure occurs repeatedly, the object appears to *strobe* or jump erratically. Objects that move quickly in a frame but aren't correctly motion-blurred tend to have this problem, particularly if there's strong perspective and the speed of the object changes as it approaches or recedes. You can see this effect in the title sequence of many movies made in the 70s, when it was popular to have the title and actors' names fly out at you. Often these weren't motion-blurred properly, and the titles seem to leap forward in a series of lurching, ragged steps instead of the smooth flight probably intended.

Third, without motion blur an object can *alias* in time. Suppose we're making a scene of a county fair, featuring a carousel. The ride makes one complete revolution every thirty seconds. We set up our motion-picture camera for a time-lapse shot, and just happen to shoot one frame every thirty seconds. The carousel will always get caught at the same position, and it will appear as though it weren't moving. If we take an image every thirty-one seconds, the carousel will appear to move forward very slowly. If we take an image every twenty-nine seconds, the carousel will appear as though it were moving backwards. Like the jaggies on the boundary of an aliased polygon, these effects result from undersampling in time.

Motion blur has been a research topic in computer graphics for many years, and a variety of algorithms can now handle it. My favorite is distribution ray tracing, which elegantly picks up all the motion-blur effects I've encountered. Many commercial rendering systems handle motion blur automatically, so sometimes animators consider it completely taken care of. I'm not so sure.

In a real camera, a mechanical device called the *shutter* controls the exposure of the film. The shutter blocks light from reaching the film before and after the exposure. At the start of an exposure, the shutter has to somehow get out of the way. When the exposure finishes, the shutter has to get back in there again. This physical phenomenon can have quite an influence on how the film gets exposed (and by implication, what audiences know and expect).

The computer graphics literature includes many camera models, lenses, and even discussions of film and its simulation. Surprisingly, I've found little discussion of simulating how a moving shutter actually exposes the film. I find it surprising because modeling the shutter

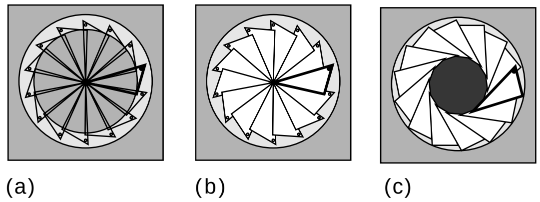

1 A 14-blade leaf shutter. (a) The shutter geometry. The leaves are opaque and overlap. (b) The leaves actually overlap in a circular pile. (c) Each leaf has rotated 30 degrees around its hinge, creating a 14-gon opening in the center.

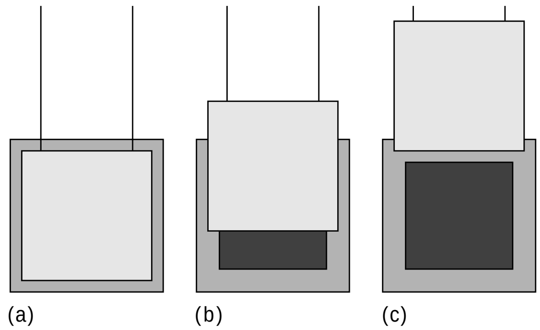

2 A focal-plane shutter. (a) The guillotine shutter covers the film. (b) The shutter rises out of the way. (c) The shutter is completely out of the way and the entire film is exposed.

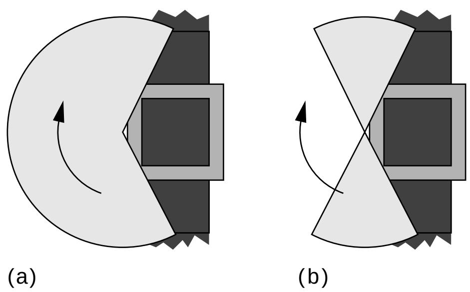

3 A rotary focal-plane shutter for a film camera. (a) A frame of film is being exposed while the shutter spins. (b) A double-sided shutter.

mechanism is easy, and as I'll show in this column, very important to the rendered image.

Physical shutters

Let's begin with the idealized, pinhole camera—a light-tight box with a piece of film taped inside one wall and a pinhole punched in the opposite wall. Normally a piece of metal or cardboard covers up the pinhole. To expose the film, you slide the cover to one side of the hole so that light can pass through and onto the film. You wait a moment, then slide the cover back into place, and voilá, you have a picture.

Most modern cameras are more sophisticated, of course. Today's 35-mm single lens reflex (SLR) cameras typically replace the pinhole with a collection of devices. Almost always these include at least a lens (or series of lenses), an iris (or aperture), and a shutter. The lens collects and focuses the incoming light. The iris blocks some of it, to protect film emulsions that would otherwise be overwhelmed by too much light and to control depth of field. Conceptually, the shutter—an opaque sheet—lies between the iris and the film. When the shutter moves out of the way, it exposes the film. Controlling the amount of time the shutter stays open provides you another handle on how much light reaches the film. It also controls the amount of motion blur. If an object rushes by you, a longer exposure time will result in a larger blur. The shutter can also influence the depth of field in the final image, but I won't go into that here.

Modern still cameras feature one or the other of two popular types of shutters: the leaf shutter and the focal-plane shutter.

A *leaf shutter* acts a lot like the mechanism that adjusts the camera's aperture. Figure 1 shows a typical leaf shutter, which consists of a ring of hinged blades that rotate independently. Normally all overlap, and the shutter admits no light. When activated, each blade of the shutter rotates to open a nearly circular hole in the center of the shutter. When the exposure time completes, the blades rotate back in towards the center, shrinking the hole until it disappears.

A *focal-plane shutter* consists of one or more opaque curtains (or sheets), often arranged in a square or rectangular formation. Figure 2 shows the idea. In a *one-curtain design*, a thin sheet of black metal sits between the aperture and the film, blocking the light. When activated, the curtain moves quickly to one side, exposing the film. After a moment, the curtain slides back into place and the exposure finishes. In a two-curtain design, the first sheet starts in the same place as in the one-curtain form, while the other curtain starts off to one side. When activated, the shutter's first curtain slides away, as before. After a time delay, the second curtain slides into the opening, once again blocking the light. The reason for having two curtains is that the time interval between the motion of the first and second sheets can be very small—the second one can be released while the first is still moving. This effectively creates a moving exposure band across the film. By controlling the timing, this band can be made very small, allowing the film to be exposed to intense amounts of light without saturating. More complex designs involve multiple curtains that travel shorter distances at faster speeds, like rotating vanes on a venetian blind.

Motion-picture cameras employ a rotary form of the focal-plane shutter, shown in Figure 3. The basic idea is to cut a rather large sector out of a spinning disk. While the opaque part of the shutter lies behind the lens, the next frame of film is moved into place and held steady. Then the missing part of the disk comes around, and light streams from the iris to the film. When the disk spins enough that the wedge blocks the light again, the film moves down and the next frame slides into place. Often the disk actually has two sectors taken out of it, so each frame of film is double exposed, capturing the scene at two slightly successive moments.

Some motion-picture cameras actually use two of these disks in series. By rotating one with respect to the

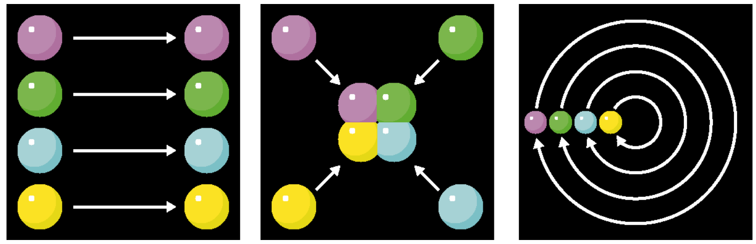

(a) (b) (c)

4 **(a) The linear test case: four balls move from left to right. (b) The radial test case: four balls move from the corners towards the center. (c) The circular test case: four balls move around the center of the frame.**

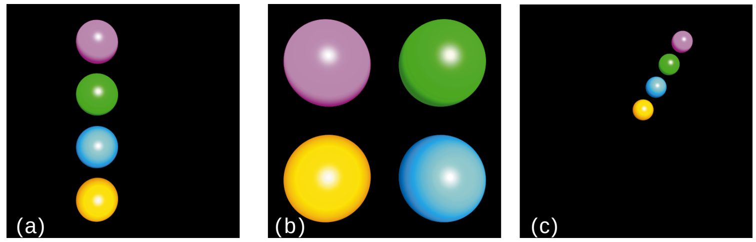

5 **The three test cases as imaged by a perfectly instantaneous shutter at *t* = 0.33: (a) linear, (b) radial, and (c) circular.**

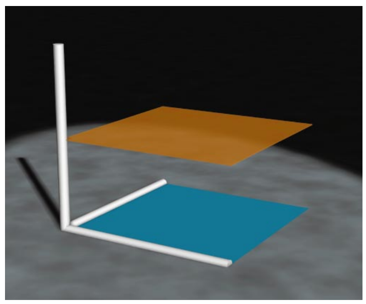

6 **The space-time diagram for the instantaneous shutter. The blue square is the frame, with the vertical time axis at the lower left corner of the frame. The orange volume, in this case a sheet parallel to the film plane, represents the shutter opening.**

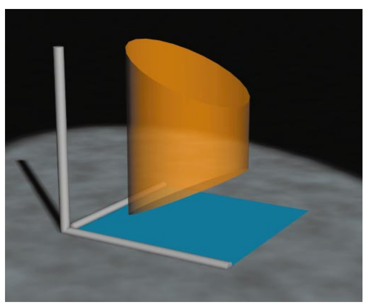

7 **A circular shutter. The shutter cover moves in a left-to-right motion, stays open for a while, then slides back into place.**

other, you can increase or decrease the size of the wedges on a continual basis, even from frame to frame.

Conceptually, these mechanisms all appear to do the same thing: expose the entire emulsion simultaneously, then block the exposure just as uniformly.

The instantaneous shutter

The standard computer graphics shutter—called the *instantaneous* shutter—serves as the default shutter for every rendering system I've ever seen (or written) that doesn't have built-in motion blur. It assumes that the entire film plane is imaged simultaneously, for a single instant.

To show the effects of different shutters, I constructed three very simple test scenes, shown in Figure 4. In Figure 4a we see four spheres moving at a uniform speed from left-to-right over the course of the frame. In Figure 4b, these four balls start just inside the four corners of the frame and move inwards until they're overlapping in the center. In Figure 4c, the four balls (now reduced in size) each make a single full rotation around the center of the frame. I call these the *linear*, *radial*, and *circular* cases. These balls live in a perfectly black environment, lit by a single light. In each test scene, the balls start moving at time $t = 0$ and move with uniform speed until $t = 1$. I'll assume that every exposure in this column takes place in this interval from 0 to 1. Finite shutters will generally expose different parts of the film to different pieces of this time interval.

Figure 5 shows the results of these three scenes with an instantaneous shutter at time $t = 0.33$. Notice they're frozen in place, just as we'd expect. I rendered the images in Figure 5, as well as all the other test scene images in this column, with a distribution ray tracer. That means that sometimes images have speckled regions where objects are moving quickly; more rays (and more computing time) would reduce those effects.

To illustrate the shutters in this column, I'll provide a little 3D space-time diagram for each one, as in Figure 6. The blue square on the ground plane represents the square film emulsion; the vertical axis marks time, and also identifies the lower-left corner of the film. At the ground plane, the time is zero; at the tip of the axis the time is 1. To find the period of time for which any point on the film is exposed, simply locate that point on the blue square and follow it upwards. The orange region above that point represents the exposure interval in time for that spot on the film. For this instantaneous shutter, there's a plane sitting in the cube partway up the time axis, parallel to the film plane. For clarity's sake, the orange solids in these figures don't cast shadows.

For comparison, Figure 7 shows a circular shutter. I'll assume that the iris is a little too close to the film, so that only a circular region of the film actually gets exposed. The figure shows that at the start of the exposure, the shutter is opened by sliding a cover linearly from right to left; first the left-most bit of the circle is exposed, then some to the right, and so on. Finally the cover is all the way to one side, and the entire circular area of the film is exposed. Then the cover slides back into place, completing the exposure.


Image courtesy of Andrew Davidhazy

8 **A water splash caught with the equivalent of a very fast shutter. Photograph made in Andrew Davidhazy's photoinstrumetnation course at the Rochester Institute of Technology.**

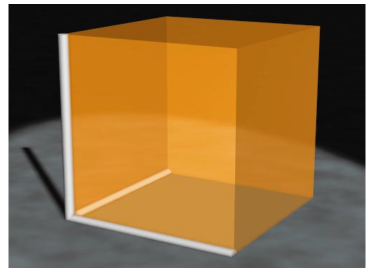

9 **The space-time diagram for the uniform shutter.**

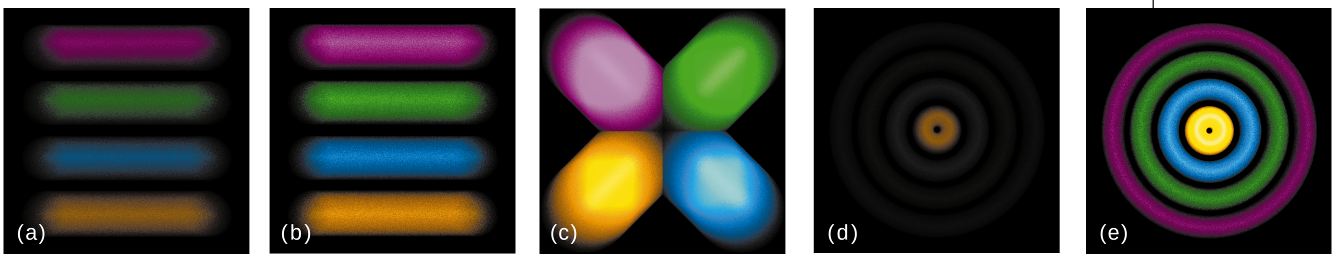

10 **The three test scenes rendered with the uniform shutter. (a) The linear scene. (b) Because the balls cover so much distance, the image in (a) is somewhat dim, so here I increased the brightness of the final image. (c) The radial scene. (d) The circular scene. (e) The circular scene with increased brightness.**

Before leaving the instantaneous shutter, I'd like to mention that professional photographers have a way to simulate it. They set up the shot, turn off all the lights in the room, then open the shutter and hold it open. At the proper moment, they fire a strobe light (or flash) that delivers a short burst of light. When everything is dark again, the shutter closes. In this scenario the entire surface of the film truly is exposed briefly and uniformly at the same time, though it's for a short, finite interval. Figure 8 shows an example of a water splash caught this way.

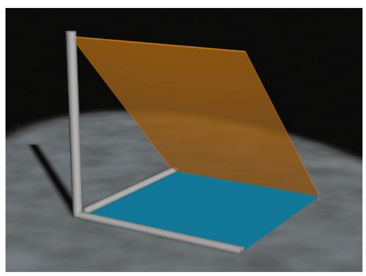

11 **The space-time diagram for the slit focal-plane shutter. The shutter is a thin horizontal line that sweeps the film plane from top to bottom.**

The uniform shutter

The opposite of the instantaneous shutter—the *uniform shutter*—serves as the default shutter for every renderer I've seen (or written) that includes motion blur. Figure 9 shows the space-time diagram. Notice that the entire cube is filled in, indicating that every point on the film is exposed for the entire duration of time. Figure 10 shows the results of our three test scenes imaged with the uniform shutter.

As you might have expected, we can see the entire path of each ball. You might also notice that as a ball moves more quickly, it appears dimmer. This is reasonable when you consider that as an object streams past the camera more quickly than another object, it reflects less light back at the lens and the film, so the film picks up less of that object's image. The circular scene also uses smaller spheres, which reduces their apparent brightness even further. For purposes of illustration, I also provide brightened versions of these images—they weren't rerendered, but simply postprocessed.

You may also notice that the motion blur is symmetrical. That's because the balls are moving at constant speed throughout the course of the exposure. If they started at rest and picked up speed, you'd see a brighter image at the start because they would be in that position longer relative to the end of the exposure. In some commercial packages you can fake motion blur on still images by blurring them, and sometimes you have the option to control how symmetrical the blur appears. If the object moves at uniform speed, the blur is symmetrical. Only if it starts or ends at rest will there be a sharp edge at the start or end of an exposure made with a uniform shutter.

The uniform shutter is a nice little shutter. It's easy to program, and you can easily interpret the results. Of course it's not a real shutter most of the time. The exception is a simulation like the one described above for instantaneous shutters, only instead of a brief pulse of light you leave the light on for the entire duration of the exposure. Since the shutter is fully open during the entire time the light is on, the film is uniformly exposed over the duration of the frame.

The focal-plane shutter

Let's look at a simulation of the focal-plane shutter. For starters, I'll assume a two-curtain shutter where the second shutter lies only a tiny distance behind the first. A simpler mechanical version of this would be a single double-wide curtain with a vertical slit. As the curtain moves from top to bottom, the slit moves across the film plane. Figure 11 shows the space-time diagram.

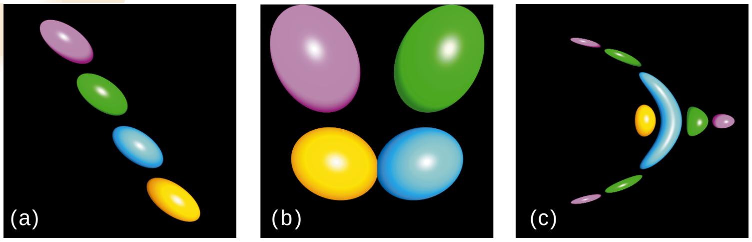

12 The three test scenes rendered with the horizontal slit focal-plane shutter: (a) linear, (b) radial, and (c) circular.

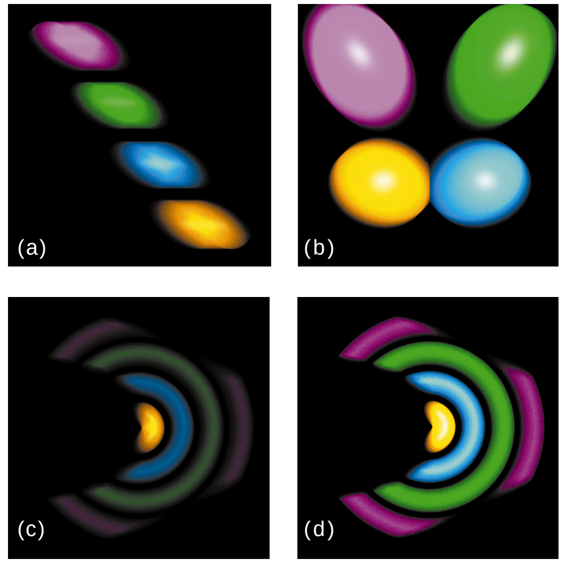

15 The three scenes rendered with the shutter of Figure 14: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).

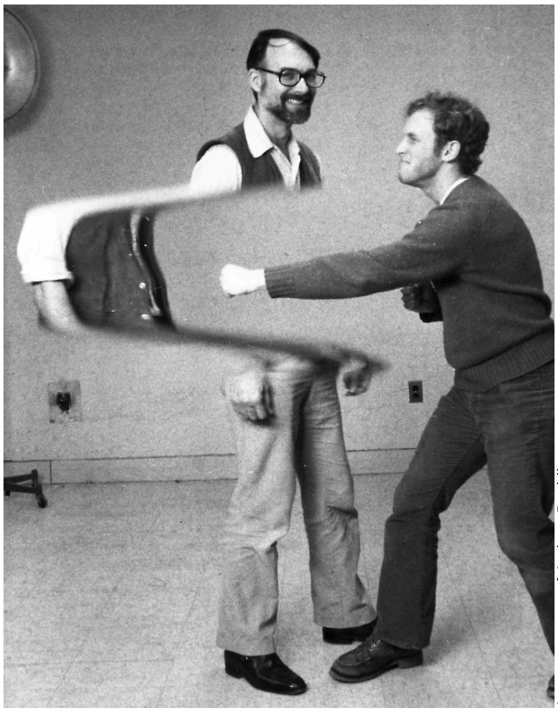

Image courtesy of Andrew Davidhazy

13 An image made with a slow-moving horizontal-slit focal-plane shutter. Photograph made in Andrew Davidhazy's photoinstrumetnation course at the Rochester Institute of Technology.

16 The space-time diagram for the horizontal-wipe focal-plane shutter (the wipe is actually from top to bottom, but it's a horizontal edge that moves down the film plane).

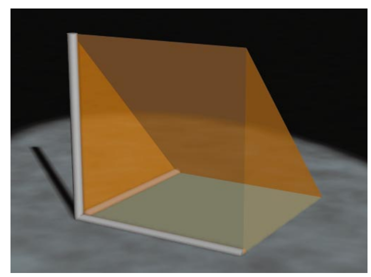

14 The space-time diagram for the horizontal-band focal-plane shutter.

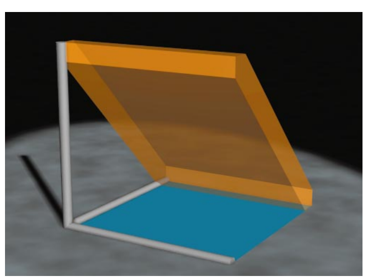

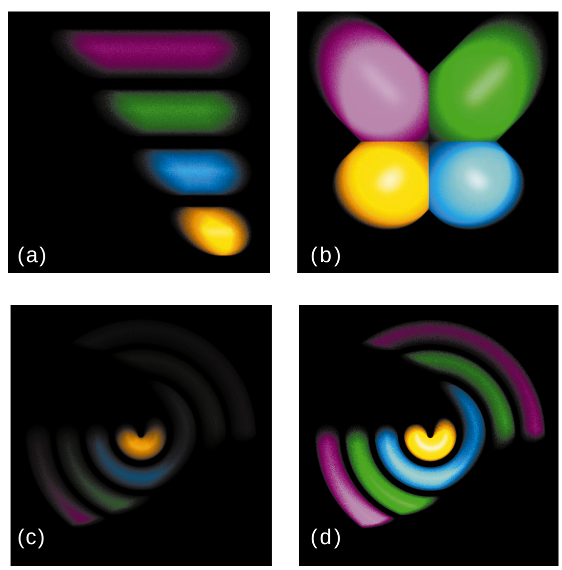

17 The three scenes rendered with the shutter of Figure 16: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).

Figure 12 shows our three scenes using this top-to-bottom, focal-plane slit shutter. The balls are skewed in Figure 12a because as the shutter moves down the frame, each ball moves from left to right, so each scan-line picks up a slightly later moment in time. In Figure 12b, the slit moves along with the red and green balls as they move downward, causing them to stretch out into a distorted, egg-like appearance. By the time the slit reaches the bottom half of the frame, the yellow and

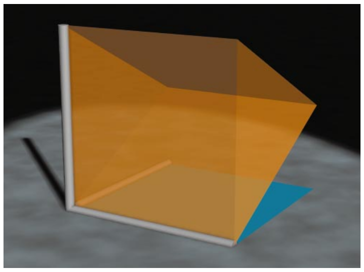

18 **The cyclic horizontal-wipe shutter.**

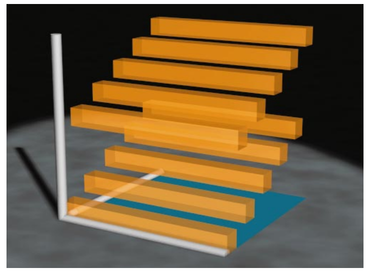

20 **The space-time diagram for a 10-band interlaced shutter.**

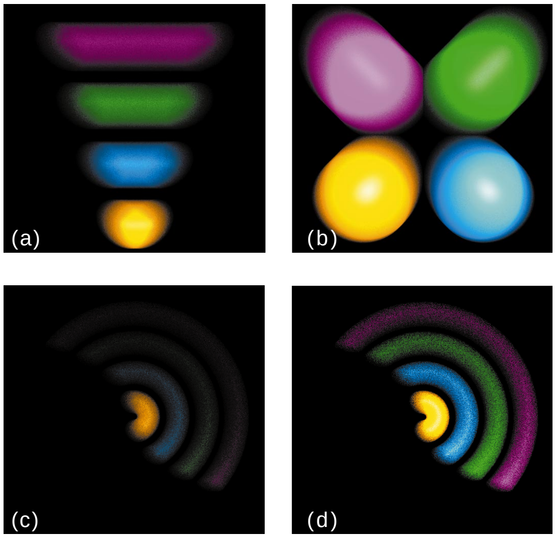

19 **The three scenes rendered with the shutter of Figure 18: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

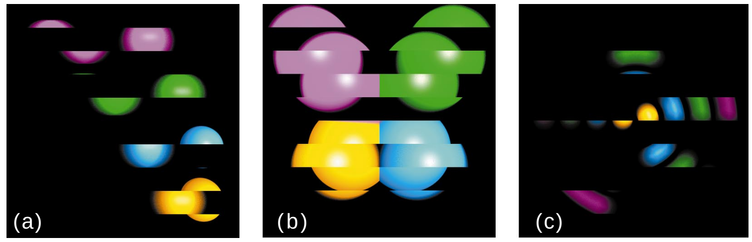

21 **The three scenes rendered with the 10-band interlaced shutter.**

blue balls are about halfway along their paths. Because those balls move upward while the slit moves downward, they're compressed. By the time the slit reaches the lower part of the frame, the balls are already long gone. Figure 12c shows the spinning balls and dramatically points out the interaction between the shutter and the motion of the objects in the frame. As the shutter moves its way down, it picks up the red and green balls at three different places in their orbits; the slower blue and yellow balls are caught only once.

You can have a lot of fun with this kind of shutter if it moves slowly enough. Figure 13 shows an image made with a very slow-moving horizontal slit shutter. The slit-scan method has been used for years by the special effects industry. It was used to great effect in "2001: A Space Odyssey."

Let's turn to a slightly more physical situation, where the slit is a band. That is, we leave a gap between the trailing edge of the first curtain, exposing the film, and the leading edge of the second curtain, cutting off the light. Figure 14 shows the space-time diagram, and Figure 15 shows our three scenes. As you'd probably expect, they look a lot like Figure 12 and the slit shutter, except that the balls are a little smeared out in space.

A third variation to this approach uses just a single curtain that exposes the film from top to bottom, and then instantly springs shut again—like a *wipe* in video transition terminology. Figure 16 shows the space-time diagram. The top line of the film is exposed at the very start of the frame and remains open to light throughout. Each successive line down the frame starts receiving light a moment after its predecessor and continues to receive light until the end of the frame. The whole thing then shuts down at frame's end—physically implausible, but perfectly easy to program.

Figure 17 shows the three scenes as imaged by this shutter.

Finally, Figure 18 shows a physically based bottom-to-top single-curtain shutter. I call this the *cyclic* version of the shutter, since it returns to its starting point. Figure 19 shows the results of imaging each of the three scenes with this shutter. A variation of this model speeds up the motion of the shutter so that it opens up, and the film is left completely exposed for a while before the shutter closes again (the shutter of Figure 7 follows this approach). In Figure 18, the shutter starts to close immediately after becoming fully open.

The important thing to notice here is that these pictures all differ from the uniform and instantaneous shutters. Of these four variations, the band and cyclic versions probably come closest to approximating real physical shutters in common use. I'll use these four versions—the *line*, *band*, *wipe*, and *cycle* to illustrate most shutter geometries from here on.

The interlaced shutter

Okay, I admit it—this is a weird one. But since today's televisions still scan out an image in an interlaced format, I thought an interlaced shutter would be a good match. And it's even physically plausible. Cut up the film plane into some number of horizontal bands—let's number them starting with 1 at the top. Expose band 1, then bands 3, 5, 7, and so on to the bottom. Then come back up and expose bands 2, 4, 6, on to the bottom of the frame. A television does this on a line-by-line basis, but I thought using a smaller number of bands would illustrate the point better.

Figure 20 shows the space-time diagram for an interlaced shutter using 10 bands. Figure 21 shows the rendered results for our three test scenes.

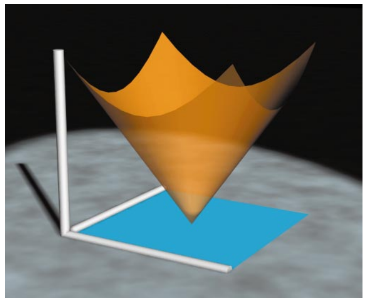

22 **The space-time diagram for the leaf-slit shutter.**

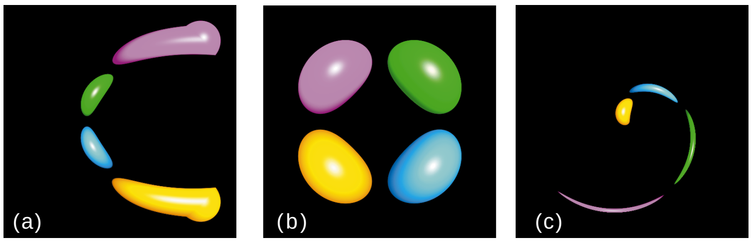

23 **The three test scenes rendered with the leaf-slit shutter: (a) linear, (b) radial, and (c) circular.**

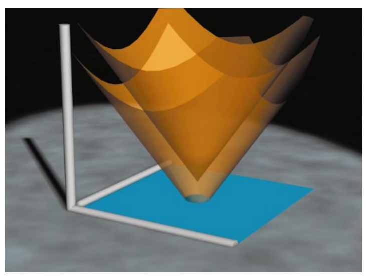

24 **The space-time diagram for the leaf-band shutter. The exposed region lies between the two cones.**

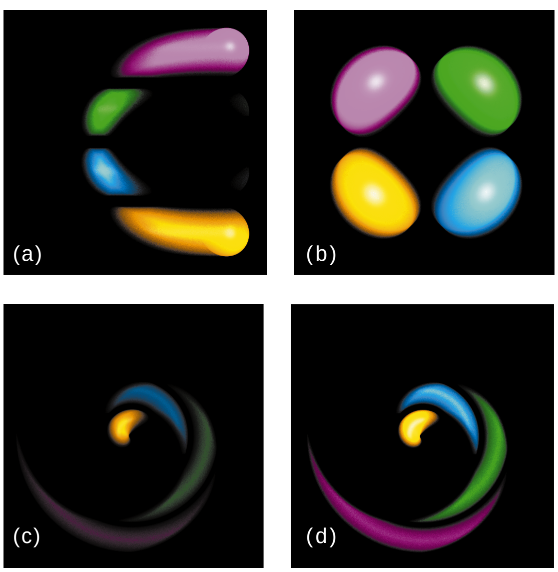

25 **The three test scenes rendered with the leaf band shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

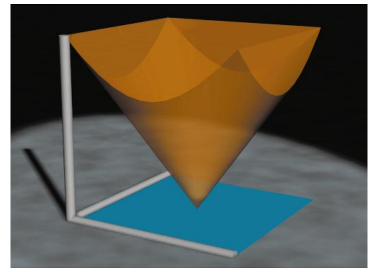

26 **The space-time diagram for the leaf wipe shutter.**

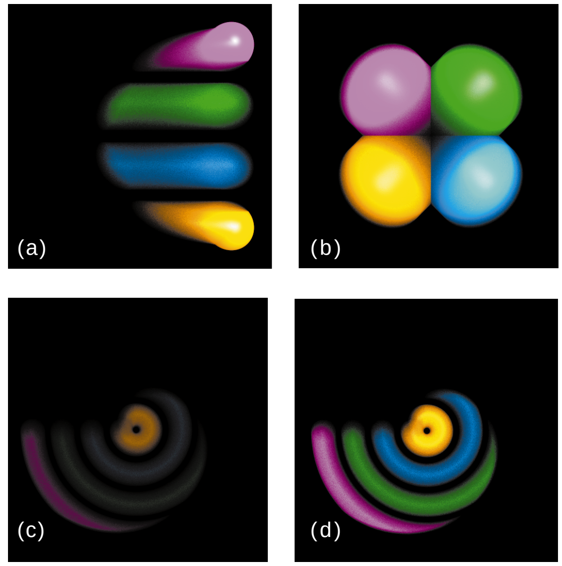

27 **The three test scenes rendered with the leaf-wipe shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

The leaf shutter

Let's now look at the other popular shutter, based on the rotary iris mechanism of Figure 1.

The first example will be a slit, as before. Here we start with a completely opaque shutter, then a small hole opens in the middle, which enlarges into a circular opening. Note that in the slit version, this isn't a full disk, but just an infinitely thin circle expanding through the sheet, like a water ripple. Though you couldn't build it this way, you can imagine this shutter as a piece of metal with a hole in the center, and the hole steadily grows. Inside the hole lies another black disk, and it grows along with the circular hole so there's just the thinnest circular ring left to expose the film. The missing ring grows larger and larger until the end of the exposure, when it reaches the corners.

Figure 22 shows the space-time diagram. Figure 23 shows our three test scenes. These certainly weren't what I expected from this seemingly innocent circular shutter.

In the linear case, we'd expect to reach the green and blue balls near the center faster than the red and yellow ones at the outside. Indeed we do, since we see images of the blue and green balls to the left of the others. We track the red and yellow balls with the top and bottom of the shutter as it opens, giving us the spread-out trails shown in Figure 23a. In Figure 23b we see the imploding balls are caught uniformly as they approach the center, and they're squished because their radially inward motion runs against the radially outward circle to reduce the amount of each ball visible to the film. In the circular case of Figure 23c, we simply pick up different arcs of the ball's path as the shutter ring expands.

As before, I implemented a band version of the shut-

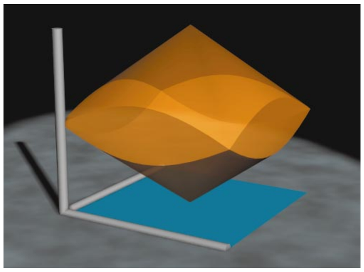

28 **The space-time diagram for the leaf-cycle shutter.**

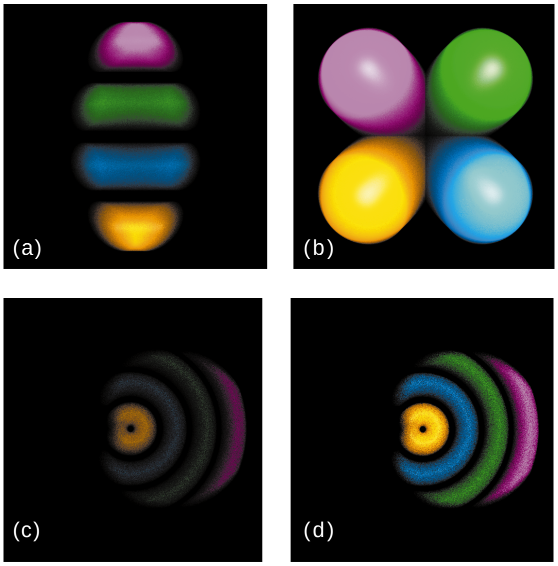

29 **The three test scenes rendered with the leaf cycle shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

ter. Figure 24 shows the space-time diagram. Figure 25 shows the rendered scenes.

Figure 26 shows the integrating, or wipe, version of this shutter. This is much more like the standard leaf shutter—as the shutter opens, the exposed film takes the form of an enlarging disk. Figure 27 shows the rendered scenes.

Finally, Figure 28 shows the symmetrical, or cyclic, version of the shutter—a pretty close approximation of a real leaf shutter. It starts out as a small circle, enlarges, then gets small again. Figure 29 shows our three scenes rendered with this shutter.

The inverse-leaf shutter

Now that we've looked at the two popular physical shutters, let's play around a little with shutters that might be hard to actually build. You can think up all kinds of crazy things, but most are just strange.

One that isn't too weird—the *inverse leaf shutter*—reverses the process of the leaf shutter by exposing the outermost corners, then working inward.

Figure 30 shows the space-time diagram for the line version of the inverse leaf. Figure 31 shows the results for our test scenes. The linear and circular scenes are reasonable reversals of the standard leaf shutter. However, the radial scene differs because the exposure circle pulls in just as the balls move toward the center, so we get the entire track of balls in the final image (the red ball starts just slightly farther out than the others).

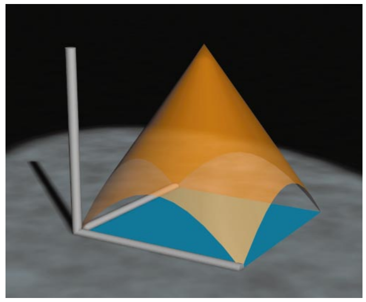

30 **The space-time diagram for the inverse-leaf slit shutter.**

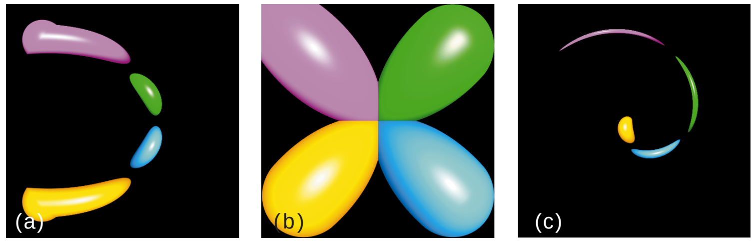

31 **The three test scenes rendered with the inverse-leaf slit shutter: (a) linear, (b) radial, and (c) circular.**

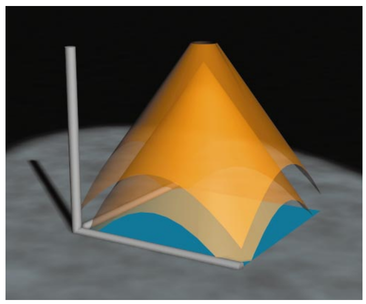

32 **The space-time diagram for the inverse-leaf band shutter. The exposed region lies between the two cones.**

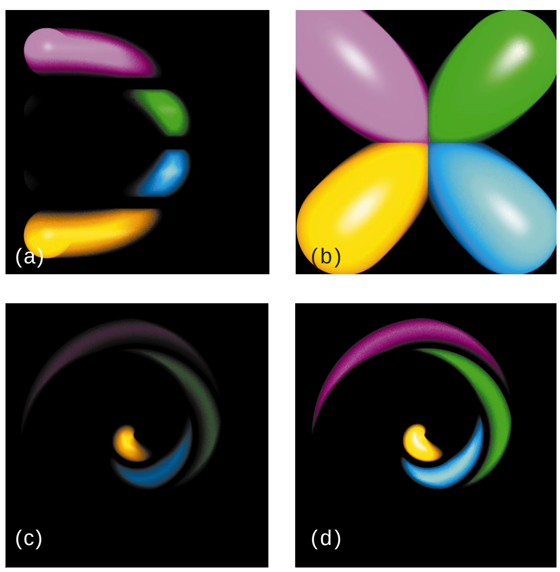

33 **The three test scenes rendered with the inverse-leaf band shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

Figure 32 shows the space-time diagram when we thicken this circle into a band. Figure 33 shows the resulting images. They're mostly smeared versions of the sharper images in Figure 31.

34 **The space-time diagram for the inverse-leaf wipe shutter.**

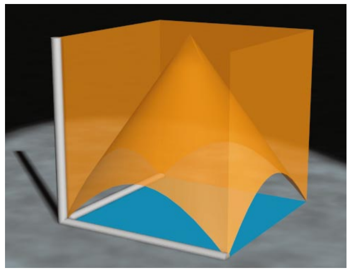

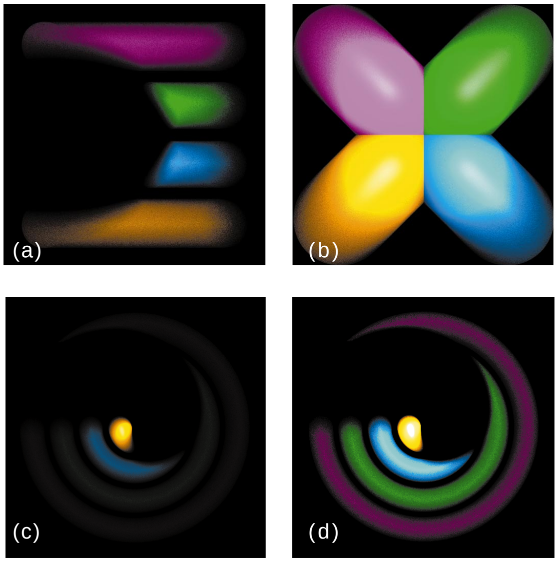

35 **The three test scenes rendered with the inverse-leaf wipe shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

36 **The space-time diagram for the inverse-leaf cycle shutter.**

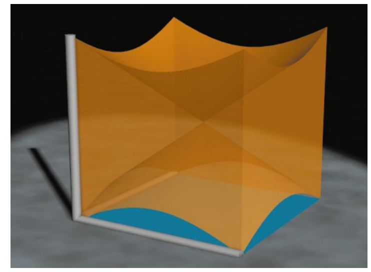

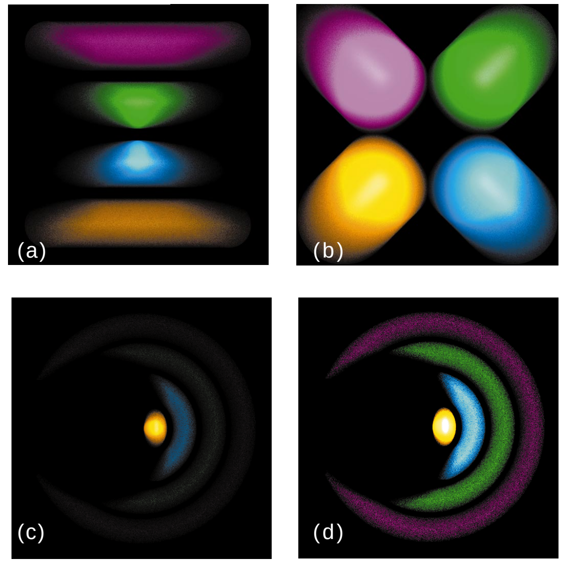

37 **The three test scenes rendered with the inverse-leaf cycle shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

Figure 34 shows the wiping version of this shutter. Figure 35 shows the scenes imaged. The unusual one here is the linear case. I spent some time puzzling over it before seeing that this was really correct. The black disk on the left side of the picture isn't really the shutter itself—it's an artifact of the interaction between the circular shutter shape and the motion of the balls. The red and yellow balls occupy the outer corners at the start of the frame and therefore get slightly imaged. As the circle pulls in, it reveals more and more of the film for the rest of the frame. So the curved edge on the bottom of the red path and the top of the yellow one is the image of the top and bottom of the shutter as it contracts and the balls speed by. The green and blue balls have already left their starting points by the time the shutter reaches that position, so their left side is black. But because the shutter leaves the film open to exposure as it moves, the first few moments of exposure reveal the right side of the frame as well as the left. The left edge of the green and blue trails shows that point in space and time where the contracting right edge of the shutter lies at the same place as the rightward-flying balls.

Finally, Figure 36 shows the cyclic version of this shutter. Figure 37 shows the rendered scenes.

The clock shutter

I call my last imaginary shutter the *clock shutter*. In the slit form, we originally have a slit from the center of the frame to the center of the right edge—the position of a clock's hour hand at 3 pm. As the frame is exposed, the slit runs around counterclockwise through one revolu-

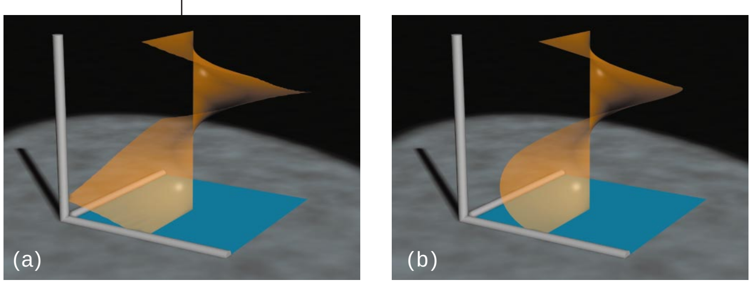

38 **The space-time diagram for the clock-slit shutter. (a) The shutter surface is clipped to the boundaries of the frame. (b) The shutter radius is reduced and not clipped.**

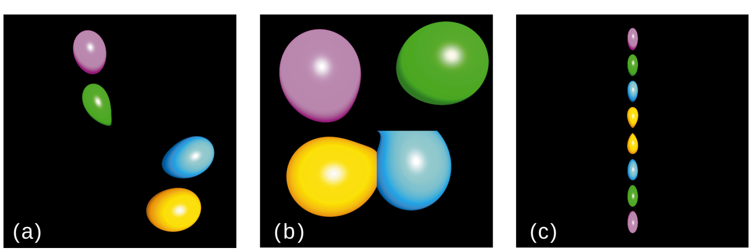

39 **The three test scenes rendered with the clock-slit shutter: (a) linear, (b) radial, and (c) circular.**

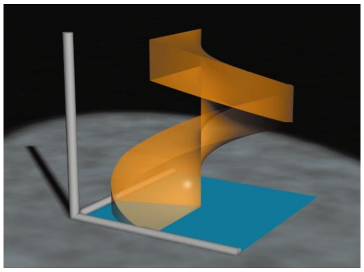

40 **The space-time diagram for the clock-band shutter.**

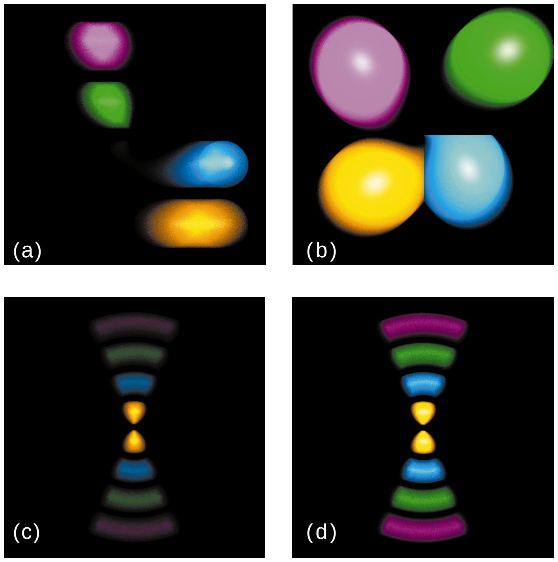

41 **The three test scenes rendered with the clock-band shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

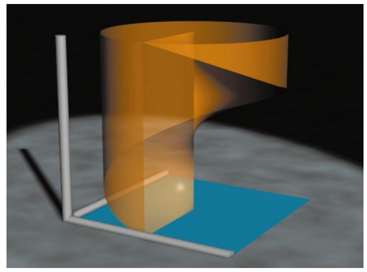

42 **The space-time diagram for the clock wipe shutter.**

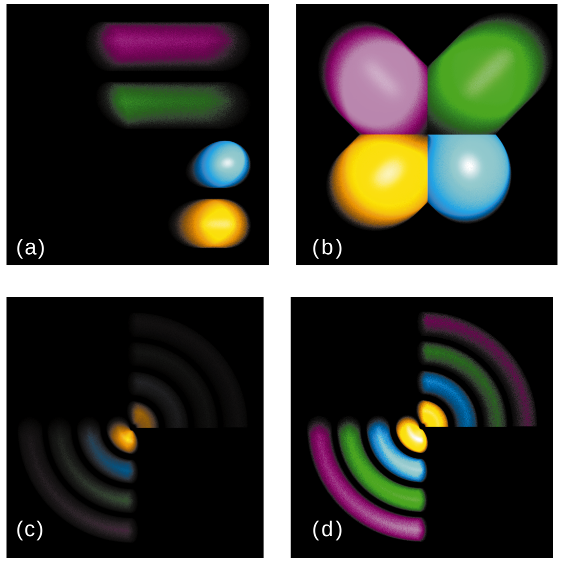

43 **The three test scenes rendered with the clock wipe shutter: (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

tion. This is something like the rotary leaf shutter of Figure 3.

Figure 38a shows the space-time diagram for the slit form. In Figure 38a I clipped the orange exposure solid to the sides of the frame, as I did in the previous figures.

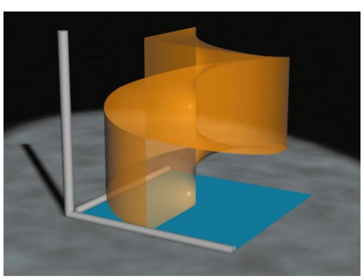

44 **The space-time diagram for the clock cycle shutter.**

In Figure 38b I reduced the radius of the spiral and didn't clip it. I think Figure 38b is clearer and more attractive, so the remaining space-time diagrams for the clock shutter use this reduced radius and no clipping. Note that the rendered images use the full-radius shutter. Figure 39 shows the images for this shutter. It's fun to interpret them. You need to picture the motion of the balls in your mind's eye and superimpose on that the motion of the shutter.

Figure 40 shows the band diagram of the shutter, which created the images in Figure 41. Figure 42 shows the wipe version of the shutter, which was used to make the scene renderings in Figure 43. Finally, Figure 44 shows the cyclic version of this shutter, and Figure 45 its rendered scenes.

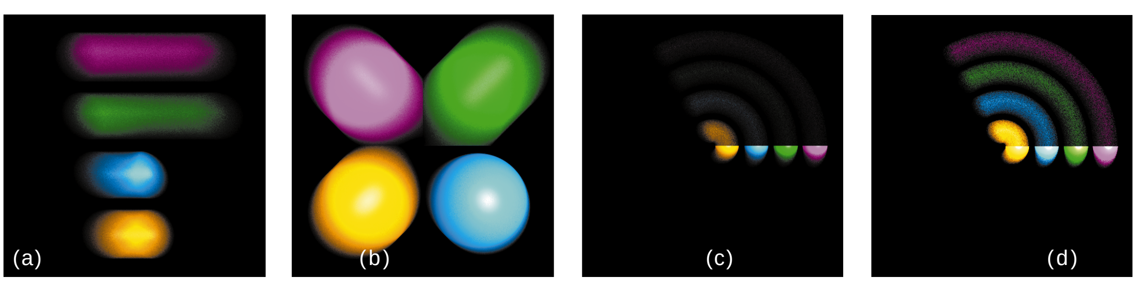

45 **The three test scenes rendered with the clock-cycle : (a) linear, (b) radial, and (c) circular. (d) A brightness-enhanced version of (c).**

Iris out

We can draw three basic principles from these examples regarding the interplay of the shutter geometry and object motion. First, motion perpendicular to the edge of the shutter results in skewed objects. Second, objects that move along with the shutter edge are elongated. Third, objects that move in a way opposite to the shutter edge motion are compressed.

The way that a shutter exposes a piece of film greatly influences how motion is captured and imaged. The instantaneous and uniform shutters that computer graphics uses almost exclusively are convenient to program and think about, but they don't correspond to any commonly available, physical exposure devices.

Just how important is it to model a real shutter? I don't know. I haven't seen any literature on how audiences respond to different shutters, and I haven't had the opportunity to try even some informal tests for myself. So it may be that this isn't a big deal, and we don't need to worry about it. That would surprise me. We know that we're sensitive to the presence or absence of motion blur. Now it may be that any motion blur is as good as correct motion blur. It may also be that I've been unfairly emphasizing edge effects—perhaps in practice the majority of the frame is fully exposed, so the motion and shape of the shutter is relevant for only a very small piece of the total exposure time. But we've all seen so many films that were shot with physical shutters, I'd bet that we're sensitive to their effects. My guess would be that if you create two pieces of animation, one with a uniform shutter and one with a correctly simulated rotary shutter as in Figure 3, viewers would prefer the physically motivated shutter images, even if they couldn't tell you why.

This is important in a different way for video than for film. A video camera doesn't have a mechanical shutter, but the process in which the image is captured, then later played back on an interlaced or progressive raster display, certainly affects how we receive the image.

In any case, it's fun to play around with the exposure model in our synthetic digital cameras and investigate the different ways we can create images without those expensive trips to the camera store.

Click. Whirrr. ■

Acknowledgments

Thanks to Dan Robbins, who helped figure out how to model the clock shutter spiral forms in 3D Studio Max 2.5, Mike Sinclair and John Platt for helpful discussions about shutters, and Andrew Davidhazy for helpful information and for providing Figures 8 and 13.

Readers may contact Glassner at Microsoft Research, e-mail glassner@microsoft.com.