

Faster computer graphics

June 13 2011, by Larry Hardesty



At right, a standard digital animation algorithm has simulated blur by sampling 256 different points on the wings of a moving butterfly for every pixel in the frame. At left is the image produced by sampling one point per pixel. In the center is the result of a new algorithm that samples only one point per pixel but infers the color values of the surrounding points. The result is very close to the 256-sample image but much easier to compute. Images courtesy of Jaakko Lehtinen

Photographs of moving objects are almost always a little blurry — or a lot blurry, if the objects are moving rapidly enough. To make their work look as much like conventional film as possible, game and movie animators try to reproduce this blur. But counterintuitively, producing blurry images is actually more computationally complex than producing perfectly sharp ones.

In August, at this year's Siggraph conference — the premier [computer-graphics](#) conference — researchers from the Computer Graphics Group

at MIT's Computer Science and Artificial Intelligence Laboratory will present a pair of papers that describe new techniques for computing blur much more efficiently. The result could be more convincing video games and frames of digital video that take minutes rather than hours to render.

The image sensor in a digital camera, and even the film in a conventional camera, can be thought of as a grid of color detectors, each detector corresponding to one pixel in the final image. If the objects being photographed are stationary, then during a single exposure, each detector registers the color of just one point on an [object](#)'s surface. But if the objects are moving, light from different points on an object, and even from different objects, will strike a single detector. The detector effectively averages the colors of all the points, and the result is blur.

Digitally rendering a frame of video is a computationally intensive process with several discrete stages. First, the computer has to determine how the objects in the scene are moving. Second, it has to calculate how rays of light from an imagined light source would reflect off the objects. Finally, it determines which rays of light would actually reach an imagined lens. If the objects in the video are moving slowly enough, the computer has to go through that process only once per frame. If the objects are moving rapidly, however, it may have to go through it dozens or even hundreds of times.

Colorfast

Given how difficult blurring is to calculate, you might think that animators would simply ignore it. But that leads to surprisingly unconvincing video. "The motion doesn't look fluid at all," says Jaakko Lehtinen, who worked on both projects as a postdoc in the Computer Graphics Group and is now a senior research scientist with graphics-chip manufacturer Nvidia.

To get a sense of what motion without blur looks like, Lehtinen says, consider the type of clay animation familiar from old movies or Christmas specials such as “Rudolph the Red-Nosed Reindeer.” “This doesn’t have motion blur, because the scene is actually stationary when you take the picture,” Lehtinen says. “It just looks choppy. The motion doesn’t look natural.”

The MIT researchers took two different approaches to simplifying the computation of blur, corresponding to two different stages in the graphics-rendering pipeline. Graduate student Jonathan Ragan-Kelley is the lead author on one of the Siggraph papers, joined by associate professor Frédo Durand, who leads the Computer Graphics Group; Lehtinen; graduate student Jiawen Chen; and Michael Doggett of Lund University in Sweden. In that paper, the researchers make the simplifying assumption that the way in which light reflects off a moving object doesn’t change over the course of a single frame. For each pixel in the final image, their algorithm still averages the colors of multiple points on objects’ surfaces, but it calculates those colors only once. The researchers found a way to represent the relationship between the color calculations and the shapes of the associated objects as entries in a table. For each pixel in the final image, the algorithm simply looks up the corresponding values in the table. That drastically simplifies the calculation but has little effect on the final image.

Adopting the researchers’ proposal would require modifying the architecture of graphics chips. “You can imagine really just going ahead and building what they suggest,” says Henry Moreton, a distinguished engineer at Nvidia. “But I think that the greater value of the paper is that it points at strategies for solving these problems more elegantly, more efficiently, and more practically. Whether they manifest themselves in exactly the fashion that the paper presents is probably not that likely. But what they did is they pointed to a new way of attacking the problem.”

Turning the tables

The second of the Computer Graphics Group's Siggraph papers, led by Lehtinen and also featuring Durand, Chen and two of Lehtinen's Nvidia colleagues, reduces the computational burden of determining which rays of light would reach an imagined lens. To produce convincing motion blur, digital animators might ordinarily consider the contributions that more than 100 discrete points on the surfaces of [moving objects](#) make to the color value of a single pixel. Lehtinen and his colleagues' algorithm instead looks at a smaller number of points — maybe 16 or so — and makes an educated guess about the color values of the points in between. The result: A frame of digital video that would ordinarily take about an hour to render might instead take about 10 minutes.

In fact, both techniques apply not only to motion blur but also to the type of blur that occurs in, say, the background of an image when the camera is focused on an object in the foreground. That, too, is something that animators seek to reproduce. “Where the director and the cinematographer choose to focus the lens, it directs your attention when you're looking at the picture in subtle ways,” Lehtinen says. If an animated film has no such lapses in focus, “there's just something wrong with it,” Lehtinen says. “It doesn't look like a movie.” Indeed, Lehtinen says, even though the paper has yet to be presented, several major special-effects companies have already contacted the researchers about the work.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Faster computer graphics (2011, June 13) retrieved 23 April 2024 from <https://phys.org/news/2011-06-faster-graphics.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.