

Render smoke and fog without being a computation hog

August 9 2007



The kind of foggy image that the new graphics method can render more quickly and efficiently than other methods. Credit: UC San Diego

Computer scientists from UC San Diego have developed a way to generate images like smoke-filled bars, foggy alleys and smog-choked cityscapes without the computational drag and slow speed of previous computer graphics methods.

"This is a huge computational savings. It lets you render explosions, smoke, and the architectural lighting design in the presence these kinds of visual effects much faster," said Wojciech Jarosz, the UC San Diego computer science Ph.D. candidate who led the study.

Today, rendering realistic computer generated images with smoke, fog, clouds or other "participating media" (some of the light is actually absorbed or reflected by the material, thus the term "participating")



generally requires a lot of computational heavy lifting, time or both.



Each circle represents the areas for which a particular cache point is valid. In order to compute the color arriving at the eye from a particular direction, the method accumulates the lighting from all cache points overlapping with the "line-of-sight" (shown in yellow). Credit: UC San Diego

"Being able to accurately and efficiently simulate these kinds of scenes is very useful," said Wojciech Jarosz (pronounced "Voy-tek Yar-os"), the project leader.

Jarosz and his colleagues from the UCSD Jacobs School of Engineering who are working to improve the process of rendering these kinds of images are presenting their work on Thursday, August 09, 2007 as a part of the "Traveling Light" session at SIGGRAPH, the premier computer graphics and interactive technologies conference.

Imagine a smoke-filled bar. Between your eyes and that mysterious person across the room is a cloud of smoke. In the real world, how well you can see through the smoke depends on a series of factors that you do not have to consciously think about. For graphics algorithms that render these kinds of images, figuring out the lighting in smoky, cloudy or



foggy settings is a computational process that gets complicated fast.

With some existing computer graphics approaches, the system has to compute the lighting at every point along the smoky line of sight between your eyes and the person at the bar – and this leads to computational bottlenecks. The new approach from Wojciech Jarosz and colleagues from the UCSD Computer Science and Engineering Department of the Jacobs School avoids this problem by taking computational short cuts. Jarosz's method, called "radiance caching," makes it easier to actually create realistic images of smoky bars and other images where some material hanging in the air interacts with the light.

With the new approach, when smoke, clouds, fog or other participating media vary smoothly across a scene, you can compute the lighting accurately at a small set of locations and then use that information to interpolate the lighting at nearby points. This approach, which is an extension of "irradiance caching," cuts the number of computations along the line of sight that need to be done to render an image.

In addition, the radiance caching system can identify and use previously computed lighting values.

"If you want to compute all the lighting along a ray, our method saves time and computational energy by considering all the precomputed values that happen to intersect with the ray," said Jarosz.

The computer scientists developed a metric based on the rate-of-change of the measured light. "In areas where the lighting changes rapidly, we can only re-use the cached values very locally, while areas where lighting is very smooth, we can re-use the values over larger distances," said Jarosz.



For example, a uniformly foggy beach scene requires fewer calculations than a foggy scene pierced with many rays of light.

"Our approach can be used for rendering both still images and animations. For still images, our technique gains efficiency by re-using lighting computations across smooth regions of an image. For animations where only the camera moves, we can also re-use the cached values across time, thereby gaining even more speedup," Jarosz said.

In their SIGGRAPH presentation, Jarosz and his co-authors compared their new system to two other approaches, "path tracing" and "photon mapping."

"Our approach handles both heterogeneous media and anisotropic phase functions, and it is several orders of magnitude faster than path tracing. Furthermore, it is view driven and well suited in large scenes where methods such as photon mapping become costly," the authors write in their SIGGRAPH sketch.

In the future, among other projects, the computer scientists would like to develop a way to reuse the same cache points temporally – in different frames of the same animation, for example.

Source: University of California - San Diego

Citation: Render smoke and fog without being a computation hog (2007, August 9) retrieved 1 May 2024 from <u>https://phys.org/news/2007-08-fog-hog.html</u>

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.