(Optimizing) Realistic Rendering with Many-Light Methods

An ACM SIGGRAPH 2012 course presented by

Jaroslav Křivánek
Charles University in Prague

Miloš Hašan
UC Berkeley

Adam Arbree
Autodesk

Carsten Dachsbacher
Karlsruhe Institute of Technology

Alexander Keller
NVIDIA Research

Bruce Walter
Cornell University

This course covers a group of global illumination algorithms known as “many-light methods”, or “VPL-rendering methods”. (VPL = virtual point light)

(Optimizing) Realistic Rendering with Many-Light Methods

Introduction

Jaroslav Křivánek
Charles University in Prague

Our goal is to render realistic images and one of the most important aspects of image realism is *global illumination*, which has been used for rendering all of the example images shown in this slides.

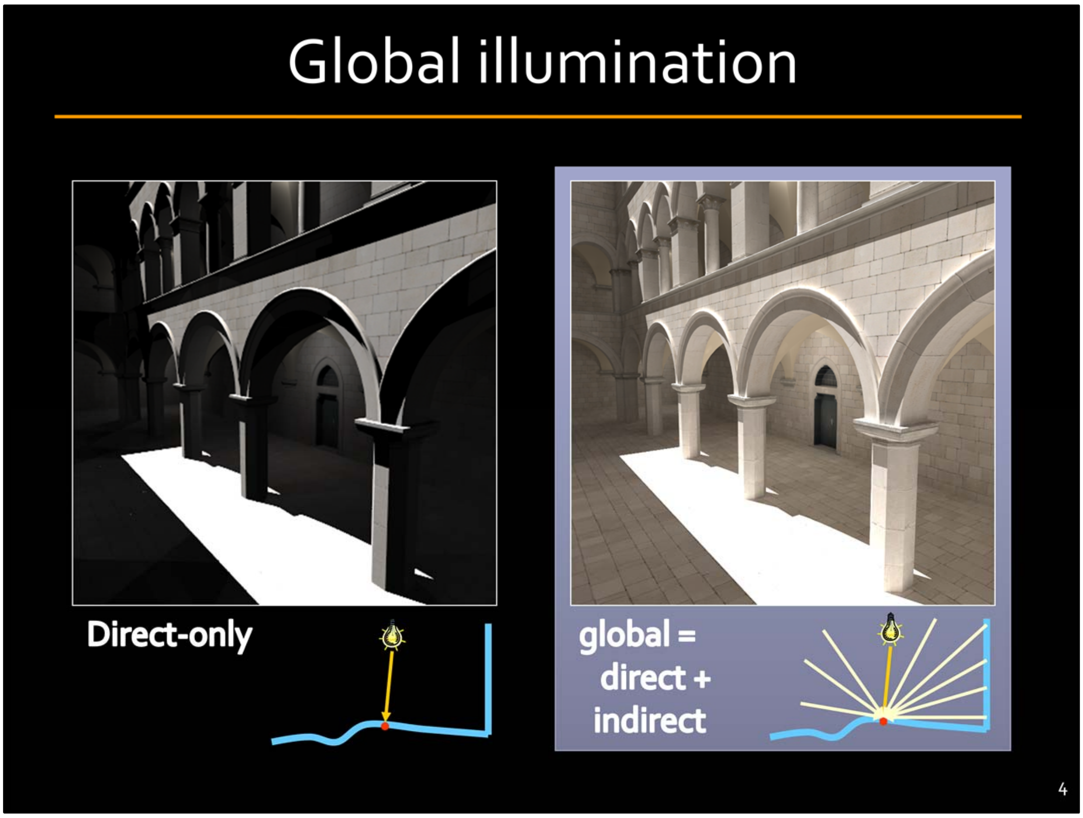

Direct illumination, shown on the left, only considers light arriving at surfaces directly form the light sources. Global illumination includes the *indirect lighting* due to multiple inter-reflections of light on scene surfaces. Doing so significantly improves the realism of the generated images.
Rendering with global illumination involves simulating the light inter-reflections, which is usually a costly process.

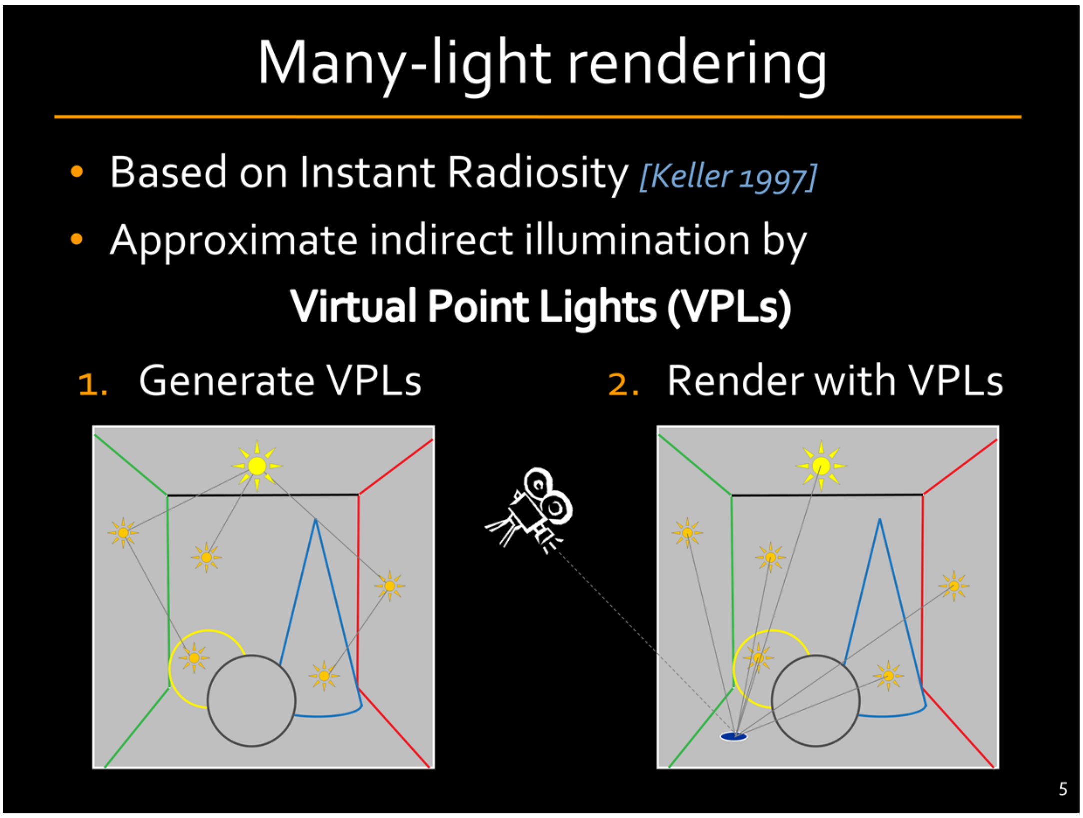

The many-light methods, covered in our course, provide a particularly efficient approach to simulating light inter-reflections.
These methods originate from the Instant Radiosity algorithm proposed by Alexander Keller, one of our presenters today. The main idea is to approximate indirect illumination by a number so called Virtual Point Lights, or VPLs.
A basic VPL rendering algorithm works as follows: In the first step, the VPLs are distributed on scene surfaces by tracing particles from light sources. In the second step, the color of a pixel is computed by summing the contributions from all the VPLs to the surface point(s) visible through that pixel.
In other words, the problem of computing global illumination has been reduced to the computation of direct illumination contributions from many-lights, hence the name.

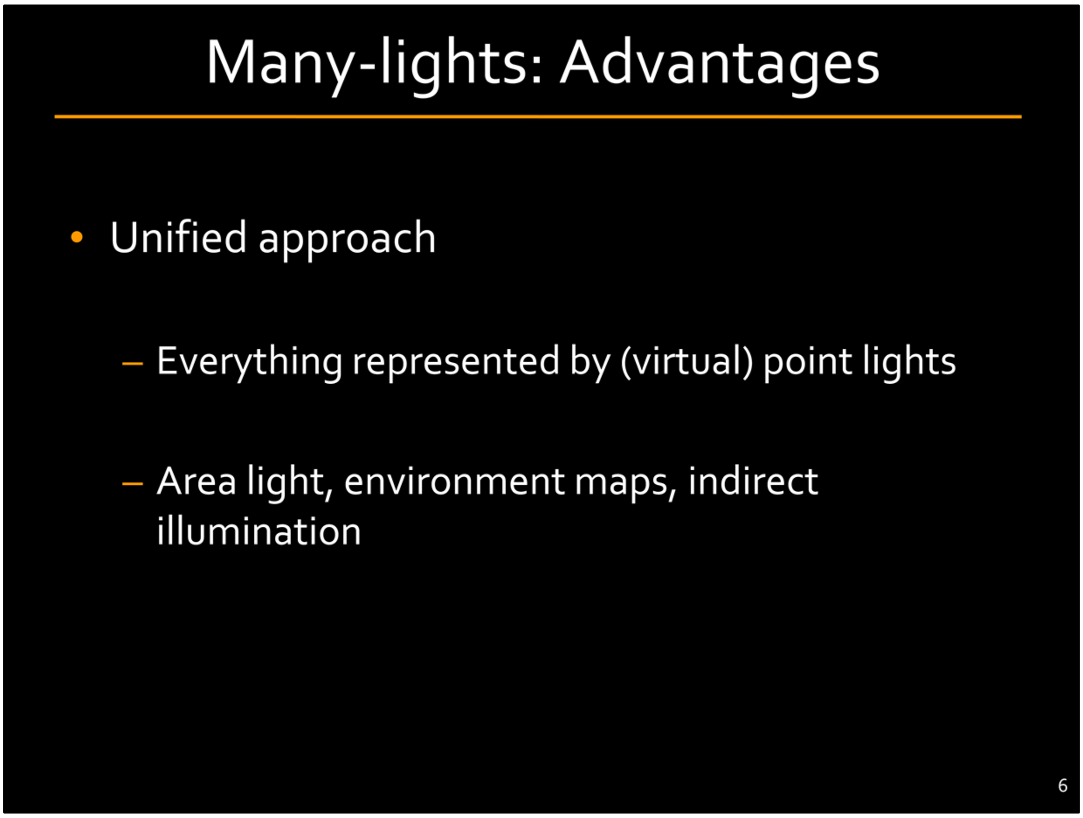

One of the most important advantages of the many-light formulation is that it unifies the rendering problem to computing direct illumination form a (potentially large) number of lights, the VPLs.

Indeed, large area lights, environment maps, and indirect illumination are seamlessly handled in the same way. All that is needed is to convert the illumination to a set of VPLs.

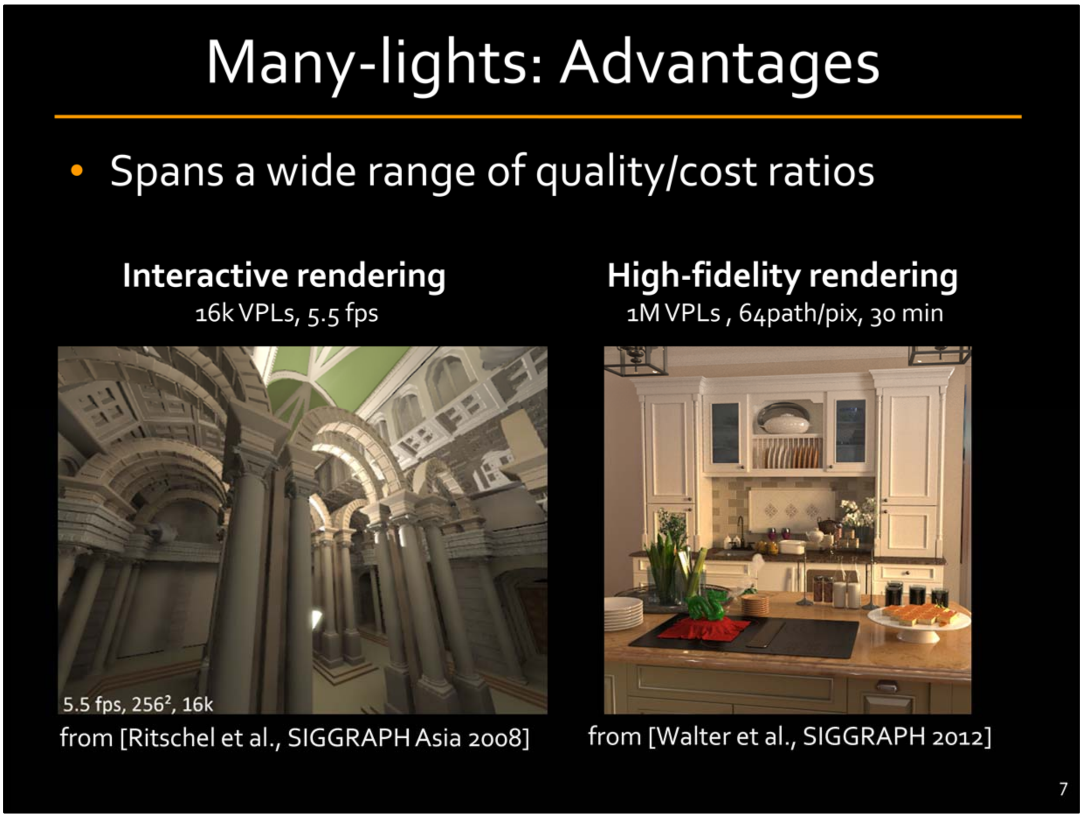

The popularity of many-light methods is also due to their wide applicability, from approximate interactive rendering to high-fidelity offline rendering. This is due to the fact that even with a limited number of VPLs, the generated images, though incorrect, provide a visually plausible approximation of the correct GI solution.

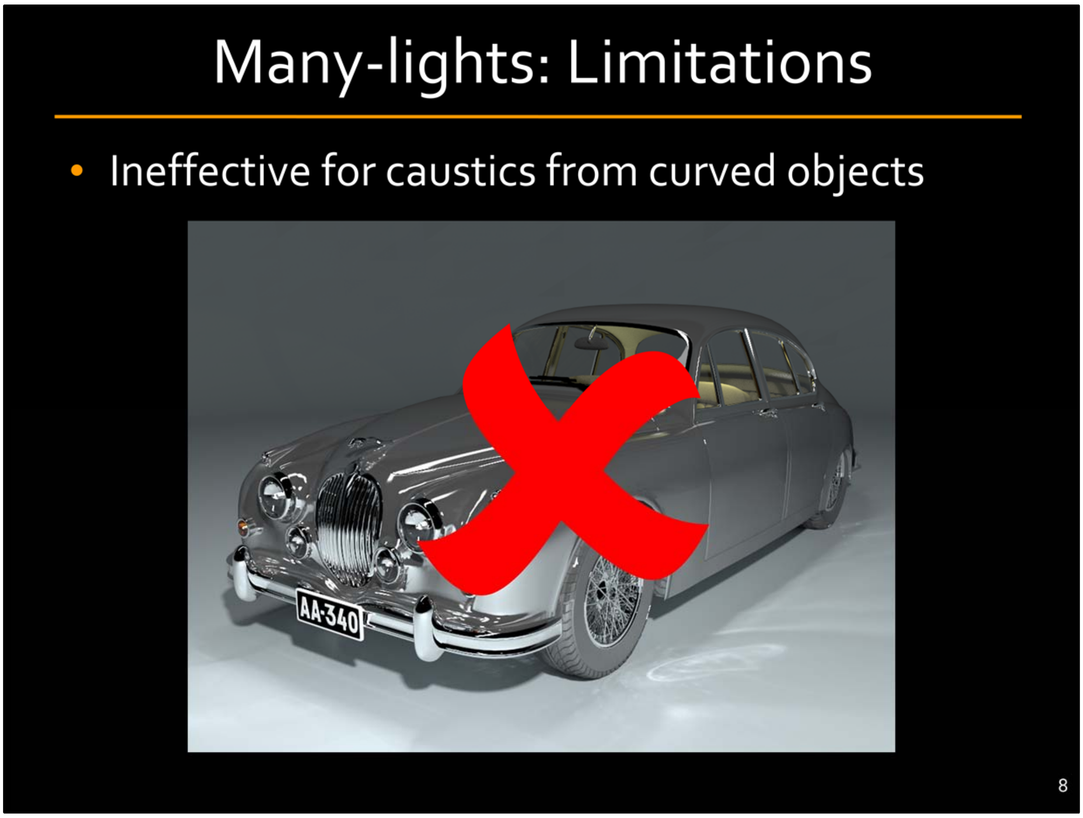

The state-of-the-art many-light methods are, however, not able to efficiently render caustics, and similar caustics-like illumination effects from general curved objects. So they cannot (yet) be considered a complete global illumination solution.

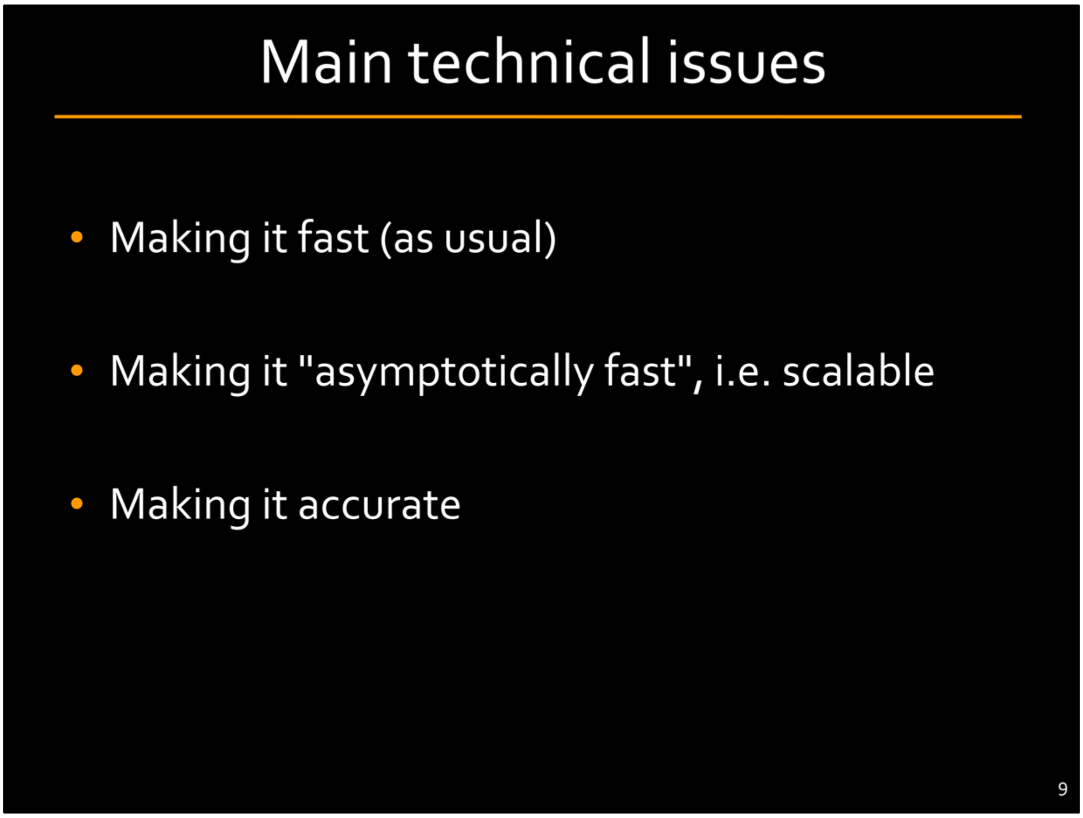

The main technical issues that need to be addressed when using many-light methods in practice are listed on the slide.
First, as usual, we want to make rendering fast. This involves especially accelerating the visibility tests required to compute illumination from VPLs. Such improvements lead to a constant factor speed-up.
However, many-light methods lend themselves to asymptotic speed improvements, i.e. the rendering time can grow much more slowly than the number of VPLs. A prime example of this approach is the Lightcuts algorithm that will be described later in the course by Bruce Walter.
Finally, in the basic form, many-light methods suffer from some approximations. Special care is needed to make these methods applicable in high-fidelity rendering applications.
The goal of our course is to cover all of these issues.

The material will be presented by researchers who were originally involved with the development of the various many-light methods.

Course Overview

- **2:00 (05 min) … Introduction & Welcome** (*Křivánek*)
- **2:05 (30 min) … Instant Radiosity** (*Keller*)
- **2:35 (30 min) … Handling difficult light paths** (*Hašan, Křivánek)*
- **3:05 (25 min) … Scalability with many lights I** (*Walter)*

- **3:30 (15 min) … Break**

- **3:45 (20 min) … Scalability with many lights II** (*Hašan)*
- **4:05 (35 min) … Real-time many-light rendering** (*Dachsbacher)*
- **4:40 (30 min) … ML in Autodesk® 360 Rendering** *(Arbree)*
- **5:10 (05 min) … Conclusion - Q & A** *(All)*

11