

The background of the cover is a surreal landscape. The ground is a grey, textured surface with a complex, winding maze pattern. In the lower-left foreground, there is a small, white, conical volcano. From the top of the volcano, a dense, chaotic forest of bright red, thin, line-like structures grows upwards, resembling a tangled mass of roots or a digital network. The title text is overlaid on the upper right portion of the image.


THE MAGIC OF COMPUTER GRAPHICS

NORIKO KURACHI | Edited by Michael Stark

The Magic of Computer Graphics

Landmarks in Rendering

Noriko Kurauchi



edited by Michael M. Stark



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

AN A K PETERS BOOK

Translation rights arranged with Ohmsha, Ltd.

Original Japanese edition published as
"CG Magic: Rendering" by Noriko Kurachi
Copyright © 2007 by Noriko Kurachi
Published by Ohmsha, Ltd.,
3-1 Kanda Nishikicho, Chiyodaku, Tokyo, Japan

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

English edition © 2011 by Taylor and Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed in the United States of America on acid-free paper
Version Date: 20110505

International Standard Book Number: 978-1-56881-577-0 (Paperback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Kurachi, Noriko.
[CG magic. English]
The magic of computer graphics : landmarks in rendering / Noriko Kurachi ; edited by
Michael M. Stark.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-56881-577-0 (alk. paper)
1. Rendering (Computer graphics) I. Stark, Michael M. II. Title.

T385.K82813 2010
006.6—dc22

2010029397

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

The Magic of Computer Graphics

Foreword

In broad terms, computer graphics is the use of computers to create images. CG has come to be the dominant form of content creation for video games, movie special effects, and many other forms of entertainment. One of the principal things CG does for creatives is to free them from the bounds of what can be held in front of a camera. From the standpoint of a naive observer, CG really is magic. Vision is entwined with our sense of reality, “seeing is believing” as they say, and when the line between physical experience and imagination is crossed, there is simply no other word for it.

The mathematical basis for computer graphics was established in the late '60s, some might say before that, though there were no computers capable of realizing its potential at the time. The best they could do was draw lines on the screen that were reminiscent of 3-D shapes. Early advances included hidden line removal and the ability to fill polygons with color(!) Later on, basic ray-tracing techniques were introduced to model how light really interacts with surfaces. This was a breakthrough compared to what came before, and as the years and decades passed, graphics researchers continued to push the boundaries of what was possible, and how real faking it could look.

Make no mistake, CG is, has been, and always will be a bag of tricks. CG artists don't reproduce the real world, even in simulation. The physics of light is simply too complex to model exactly. In the time it took you to read the word “physics,” over a trillion photons struck your retina. The computations required to completely reproduce what we see are not only beyond reach, they are largely unnecessary. You didn't notice most of those photons, anyway. Computer graphics is all about finding good approximations to reality, or magic tricks of the light.

Once we have an approximation of reality that fools most of the people most of the time, we can expand our canvas to include things not found in the real world, such as flying logos and even (wait for it) space ships! Besides offering Hollywood the ability to remake every science fiction movie ever made, CG reduces the cost of realistic depictions of the past, of thousands of actors in orc costumes, or knocking down a famous building the producer doesn't have the money to rebuild.

Speaking of buildings, computer graphics is also quite useful to those designing structures and wanting to know how daylight will enter a space before it is built. Many of the same approximations used in movies, with proper care and validation, can be used for accurate lighting predictions as well. Again, we stop short of counting every photon, but a well-designed CG trick not only looks right, but to a reasonable degree, it *is* right. Many of the established methods in CG rendering, such as Monte Carlo ray sampling, radiosity, and image-based lighting, started out as simulation techniques for building design.

Computer graphics is a vast field, and getting larger every day. It is impossible to cover every topic of interest, even within a specialization such as CG rendering. For many years, Noriko Kurachi has reported on the latest developments for Japanese readers in her monthly column for *CG World*. Being something of a pioneer herself, she selected topics that represented original and promising new directions for research, as opposed to the tried and true methods. Many of these novel ideas paid off handsomely, and these are the topics covered in this book.

Starting from the basic behavior of light, Ms. Kurachi introduces the most useful techniques for global and local illumination using geometric descriptions of an environment in the first section. She then goes on in the second section to describe image-based techniques that rely on captured data to do their magic. In the final section, she looks at the synthesis of these two complimentary approaches and what they mean for the future of computer graphics. Being ever careful to check her facts with the original researchers, she offers in these pages a journalist's view of the evolution of computer graphics over the past twenty years in a style that is accessible and thorough, tailored for and by an artist who is also a technician.

—Gregory J. Ward

Preface

The origin of this book was the technical column that I wrote every month in a computer graphics-related Japanese magazine (called *CGWORLD*). My column had two distinguishing characteristics: its topics and style.

I daringly selected topics whose value had not been well-established in the CG community but that nevertheless looked like they would have potential to make the CG scene evolve to the next stage. The writing style of the column was also unique. I communicated personally with the researchers and developers whose methods I described. These interviews enabled me to include added details about what challenges motivated their work and where their inspiration came from, as well as their insights for the future or their suggestions for practical application of their work that they couldn't include in the technical papers.

Sometime after I started this column, people came to expect that its contents would be summed up in the form of a book. The column itself spans a wide range of CG technologies and, in the beginning, it was intended for the book to cover all of them. However, as the writing went on, I discovered that it was unrealistic to cover such a huge volume of material all at once. Therefore I decided to focus on rendering technologies—the first quarter of the original plan. The book, published by the Japanese publisher Ohmsha, was released at the end of 2007 as *CG Magic: Rendering*.

From the responses of researchers and the developers that I had spoken with while writing the book, I knew that even outside of Japan it is uncommon to find a book with contents like this one. I became convinced that it would be meaningful if this book were translated into English and could be read by a wider range of people. Therefore, as soon as the Japanese book was released, I started translating it into English, and when half of translation was completed, I told Greg Ward, who had always been supportive of my writing the book, about my plan of releasing an English edition. He was very positive and gave me the opportunity to meet with Alice Peters at SIGGRAPH 2008, which led to the birth of this book. Even though this started as simply a translated version of the previous Japanese book, I wanted to include descriptions of new developments occurring on the CG scene since the release of the Japanese book. For example, the new sections on

hair rendering were added because hair rendering technologies, from a practical viewpoint, had reached a turning point in physical accuracy. In addition, a new section has been added at the end of each chapter in order to describe how the technologies introduced in the chapter are in progress.

Finally I should note two things. One is that the stories in this book focus mainly on technologies invented after 2000; therefore, I recommend that people refer to more traditional books in order to study the fundamentals of CG rendering. The other is that even though this book describes the content of professional technical papers, it has been adapted to be more accessible to people who are non-professionals in CG rendering. Therefore, I refer those people interested in further study to the original technical papers (listed in the references).

Admitting such imperfections, I hope that the stories about how these new ideas were born and eventually became widespread will be valuable for all those who are interested in the magic that CG enables. I also expect that this book will find a place with people who are interested in visual effects because this book includes many stories about how new CG theories came to be used in this area.

In general, the developing theories and the process of making them practical are very different areas; however, when combined they can make new CG technologies a reality and lead to a true revolution. The stories in this book describe efforts in both areas, and I hope that this feature of the book can inspire those who will lead the progress of CG technologies.

Acknowledgments

This book was well supported by a large number of researchers and developers. Special thanks to Greg Ward, Tony Apodaca, Holly Rushmeier, Henrik Wann Jensen, Toshiya Hachisuka, Timothy Purcell, Pat Hanrahan, Jos Stam, Craig Donner, Eugene D'Eon, Christophe Hery, Peter Sand, Paul Debevec, Mark Sagar, Pierre Buffin, Marc Levoy, Radek Grzeczuk, Ren Ng, Sing Bing Kang, Ravi Ramamoorthi, Sameer Agrawal, Eric Lafortune, Wojciech Matusik, Shree Nayar, Steve Marshner, Arno Zink, Wenzel Jakob, Martin Hill, Gero Muller, Xin Tong, Ko Nishino, Alex Vasilescu, Peter-Pike Sloan, Jan Kautz, Doug James, Changxi Zheng, Mark Sagar, Larry Gritz, Sven Woop, Tomoyuki Nishita, Yoshinori Dobashi, Yosuke Bando, Takeshi Naemura, and Masao Takakuwa for revising the book from a more current perspective, providing new stories and images in some cases, and contributing to the technical reviewing. Those were very time-consuming tasks which required a lot of work, so I honestly feel grateful to have received such contributions.

In addition, there were people who supported me even before the project started, because releasing a book for an English-speaking audience started as my personal plan. Thanks to Nick Foster, Mike Milne, Agent Xray, Andre Corwin Mazzone, Yoshiharu Gotanda, and Katsutsugu Kurachi, who helped my efforts to make the project actually happen.

One significant aspect of the project was the translation into English from Japanese because, in addition to the gap between Japanese and English colloquialisms, the technical nature of the book required the translators to have a high-level knowledge of CG rendering. Thanks to Chuk Tang, Yoshihiro Kanamori, Paul Caristino, Ryusuke Villemin, Witawat Rungjiratananon, Paulo Silva, Kiminori Tsukazaki, and Ayako Ogawa, who accomplished this difficult job.

Many thanks to A K Peters. I greatly appreciate Alice Peters, who made the project happen. Special thanks to Michael Stark, who did an excellent job revising the chapters, uniting all of the voices in the book, and smoothing the English. Thanks to the staff at A K Peters, particularly Sarah Cutler, for leading the project and often giving me advice. Thanks to Ohmsha for clearing the rights issues smoothly. Thanks to Taku Kimura for approving the use of the beautiful cover image in the English edition.

Finally, thanks to my family, friends, and colleagues who kept on encouraging me; I think I couldn't have reached the goal without their warm voices.

Contents

Foreword	vii
Preface	ix
I Geometry-Based Approaches	1
1 Introduction to Photorealistic Rendering	3
1.1 Physical Quantities of Light	3
1.2 Modeling the Behavior of Light	8
1.3 Structure of Rendering	16
2 Global Illumination	23
2.1 The Rendering Equation and Its Solutions	23
2.2 The Radiosity Method	25
2.3 Monte Carlo Ray and Path Tracing	27
2.4 Photon Mapping	38
2.5 Programmable Graphics Hardware and GI	43
3 Volume Rendering and Participating Media	53
3.1 Volume Visualization Techniques	53
3.2 Volume Rendering	55
3.3 Participating Media	62
4 Subsurface Scattering	83
4.1 Subsurface Scattering	83
4.2 Modeling Subsurface Scattering	87
4.3 BSSRDF Models	92
4.4 Recent Progress in Subsurface Scattering	119

II Image-Based Approaches	123
5 Image-Based Rendering	125
5.1 Stereo	125
5.2 Image Warping	137
5.3 Image-Based Modeling and Rendering	138
5.4 The Light Field	147
6 High Dynamic Range Imaging	165
6.1 Response Curves and HDR Imaging	165
6.2 HDR File Formats	173
6.3 HDR Video	179
7 Image-Based Lighting	187
7.1 The Basics of Image-Based Lighting	187
7.2 Environment Maps and IBL	190
III Coalescence of Geometry-Based and Image-Based Approaches	211
8 Reconstruction of Reflectance	213
8.1 BRDF Models	213
8.2 The Reflectance Field	240
8.3 Separation of Direct and Global Illumination	258
8.4 Reflectance of Hair	274
8.5 Multiple Scattering in Hair	282
8.6 Recovering Details of Hair	291
9 Bidirectional Texture Function	305
9.1 Origin of BTFs	306
9.2 BTF Synthesis	309
9.3 BTF Models	317
10 Radiance Transfer	339
10.1 PRT Fundamentals	340
10.2 Advancing Radiance Transfer	351
10.3 PRT and Dynamic Simulation	372
10.4 Precomputed Acoustic Transfer	381
10.5 Precomputed Radiance Transfer in Movies	383
Bibliography	387
Index	405



Geometry-Based Approaches

In computer graphics, *rendering* is the process of creating an image of an object or a general scene from a mathematical representation. Just as styles of painting and drawing in art range from the representational to the abstract, there are different styles of rendering. Early rendering methods, much like early artwork, concentrated on producing recognizable representations of objects—early renderings did not look at all real. *Photorealistic rendering* is the process of reproducing what an object or scene would look like if it were constructed and photographed, and this involves simulating how reflected and scattered light illuminate a real environment.

The interaction of light and matter in nature is very complicated, and has been studied in the natural sciences for many years. By the mid-twentieth century it had become commonplace in manufacturing to represent objects using spatial coordinates and mathematical expressions. New approaches for representing curved surfaces emerged by the 1960s. By then, the theoretical foundations of light transport were mostly complete, but the application of physically based light simulation was not feasible until sufficient computer processing power became available in the early 1980s. Since then photorealism has been a major goal of research in computer graphics.

Just like studies in any field of science, researchers in photorealistic rendering took the approach of starting simple. Specifically, basic surface reflection was modeled first, and illumination effects were studied in very simple scenes. More complex environments and volume effects such as smoke and fog were considered later. Research then shifted to scattering and re-emission of light from beneath a surface, an effect observed in hair and skin. Since 2000, much research has focused on generating photorealistic renderings at interactive rates, driven in part by the emergence of high-speed graphics hardware technology. In Part I, methods of photorealistic rendering that work on geometric models are introduced along with methods for implementing them on graphics hardware.

1

Introduction to Photorealistic Rendering

Photorealistic rendering is the process of image generation by simulating the physical interaction of light in an environment. Specifically, the simulation involves how light from *light sources* is reflected and scattered throughout the environment. The environment, or scene, as it is sometimes called, consists of a collection of mathematically described geometric objects. The shape description of the objects is known as the *geometric model* of the environment; however, this is only part of the complete model needed for rendering. Surface properties, such as color and reflection characteristics, as well as volume elements such as smoke and fog, affect the light propagation and are therefore a necessary part of the model. These elements, and a model for the light propagation itself all require a mathematical representation that can be incorporated into the rendering simulation.

In this chapter, the basic physical concepts used in rendering are presented. The model of light propagation most often used in rendering, radiant power carried along light paths in the form of rays, is developed in conjunction with related physical quantities of light measurement. The basic surface reflectance model used in rendering is also introduced, and the chapter concludes with a description of the two primary methods of image generation: scanline rendering and ray tracing.

1.1 Physical Quantities of Light

Modeling light interaction involves the physical quantities associated with light as electromagnetic radiation. Physically, light consists of electromagnetic waves




Figure 1.1 Flux is a measure of radiant power.

formed from simultaneously changing electric and magnetic fields. At a very small scale, light also behaves as if it exists in discrete bundles or packets (photons). However, at the macroscopic scale where rendering is normally done, particle effects really do not come into play, and wave effects have limited impact. In real environments, light is more naturally regarded as a flow of energy, the measurement of which is known as *radiometry*. The physical quantities of light measurement are called *radiometric* quantities.¹

1.1.1 Radiant Energy and Flux

A basic physical (radiometric) quantity of light is *flux* Φ , which is defined as the *radiant energy* Q per time:

$$\Phi = \frac{dQ}{dt}. \quad (1.1)$$

Flux, the temporal flow of radiant energy, is thus a measurement of radiant power and is typically expressed in watts. Flux normally refers to the radiant power reaching or flowing through some real or hypothetical surface. In the case of a light source, the flux through an imaginary sphere centered at the source represents the total flux output of the source (Figure 1.1). This flux divided by 4π , the area of a unit sphere, is the *intensity* of the source.

1.1.2 Irradiance

Flux measures the total radiant power reaching a surface, but without regard to surface position: the power might be concentrated on only part of the surface, or spread evenly across it. A position-dependent measure of radiant power is flux

¹Terminology of radiometry is not entirely standardized. The notation here follows Chapter 2 of *Radiosity and Realistic Image Synthesis* by Cohen and Wallace [Cohen and Wallace 93]; the chapter was written by Pat Hanrahan.




Figure 1.2 Irradiance caused by a light source. (a) The flux at a distance r from the source is spread out over a sphere of radius r . (b) When the flux hits a surface at an angle, the differential surface area dA gets the flux of a smaller differential area on the sphere: the irradiance is reduced by the cosine of the angle of incidence.

per area, the *irradiance* at the area. The irradiance $E(x)$ at a point x on a surface is the differential flux on a differential area dA at x divided by that area:

$$E(x) = \frac{d\Phi}{dA}, \quad (1.2)$$

where dA is perpendicular to the surface normal direction at x , i.e., is parallel to the surface at x . Irradiance is thus a measure of the radiant energy per area per time. The function $E(x)$ is formalized as a function of surface position; it is the limit of the flux in a small region containing x divided by the area of the region, as the region shrinks uniformly to x . Irradiance is also known as *flux density*, as it measures the local density of flux at the surface point.

Radiant emission of a sufficiently small light source can be regarded as a collection of expanding concentric spheres centered at the light source, each having constant flux Φ . When one of these spheres hits a surface point x , its radius r is the distance from x to the source (Figure 1.2(a)). If the light source lies directly above the surface, i.e., in the direction of the surface normal, then the differential flux $d\Phi$ incident on the differential area dA at x is the total flux Φ times the fraction of the sphere area that hits dA , which is $dA/4\pi r^2$. The irradiance at x is therefore

$$E(x) = \frac{d\Phi}{dA} = \frac{\frac{dA}{4\pi r^2} \Phi}{dA} = \frac{\Phi}{4\pi r^2}, \quad (1.3)$$

which is the familiar inverse square law of illumination.

Of course, light sources do not generally lie directly above the illuminated surface; the incident flux sphere usually hits the surface at an angle θ . In this case

the differential area dA corresponds to an area on the sphere that is smaller than dA , according to an effect known as *projective foreshortening*. The ratio of this foreshortening is the cosine of the incident angle θ (Figure 1.2(b)). Consequently, the irradiance at x is

$$E(x) = \frac{\Phi \cos \theta}{4\pi r^2}. \quad (1.4)$$

This cosine factor often appears in radiative transfer equations. It comes about because incident light is spread over an increasingly large area as the angle of incidence increases. This is one reason sunlight is stronger when the sun is high in the sky than when it is low on the horizon. It is also the principal cause of the Earth's seasons: in the winter, the sun is lower in the sky and therefore produces less irradiance.

While irradiance $E(x)$ is the radiant power received at a point x , the radiant power *leaving* a surface at x is known as the *radiant exitance* or *radiosity*, denoted by $M(x)$ or $B(x)$. The definition is the same as that of irradiance, flux per area, but the flux is regarded as leaving the surface rather than arriving. The term “radiant exitance” has become the preferred term in recent years, to avoid confusion with the “radiosity method” for computing global illumination. (described in Chapter 2).

1.1.3 Radiance

Just as flux does not depend on surface position, irradiance and radiant exitance do not depend on direction. But light emission and reflection clearly do depend on direction: the color and strength of light perceived at a photoreceptor in the human eye is dependent on the particular direction. The radiant power exiting from (or incident on) a point x in a direction $\vec{\omega}$ is called the *radiance* $L(x, \vec{\omega})$.²

Radiance is essentially a directional restriction of irradiance or radiant exitance: it is the flux per area in a differential cone of directions (Figure 1.3). As a matter of convenience, the surface through which the flux is measured need not be perpendicular to the direction $\vec{\omega}$. If θ is the angle the surface normal makes with the direction $\vec{\omega}$, then $dA \cos \theta$ is the *projected differential area*—the $\cos \theta$ factor accounts for the projective foreshortening. If $d\omega$ denotes the differential solid angle³ about the direction $\vec{\omega}$, as illustrated in Figure 1.3, radiance can be

²Although directions are often represented by 3D vectors, it is also convenient to represent directions from a point x as points on the unit sphere centered at x . Each point on the unit sphere corresponds to a unique direction. By convention, directions are often denoted by $\vec{\omega}$

³A set of directions corresponds to a subset of the unit sphere; the *solid angle* of the set is the area of the corresponding spherical subset. The set of all directions therefore has solid angle 4π , the area of the unit sphere. The differential solid angle $d\omega$ of a differential cone is the differential area of the unit sphere subtended by the cone.