

Physically Based Rendering is a terrific book. It covers all the marvelous math, fascinating physics, practical software engineering, and clever tricks that are necessary to write a state-of-the-art photorealistic renderer. All of these topics are dealt with in a clear and pedagogical manner without omitting the all-important practical details.

pbrt is not just a “toy” implementation of a ray tracer, but a general and robust full-scale global illumination renderer. It contains many important optimizations to reduce execution time and memory consumption for complex scenes. Furthermore, pbrt is easy to extend to experiment with other rendering algorithm variations.

This book is not only a textbook for students, but also a useful reference book for practitioners in the field. The second edition has been extended with sections on Metropolis light transport, subsurface scattering, precomputed light transport, and more.

Per Christensen

Senior Software Developer, RenderMan Products, Pixar Animation Studios

Looking for a job in research or high end rendering? Get your kick-start education and start your own project with this book that comes along with both theory and real examples, meaning real code.

With their second edition, Matt Pharr and Greg Humphreys provide easy access to even the most advanced rendering techniques like Metropolis light transport and quasi-Monte Carlo methods. In addition the framework lets you skip the bootstrap pain of getting data into and out of your renderer.

The holistic approach of literate programming results in a clear logic of an easy-to-read text. If you are serious about graphics, there is no way around this unique and extremely valuable book that is closest to the state of the art.

Alexander Keller

Chief Scientist, Mental Images

Physically Based Rendering

FROM THEORY TO IMPLEMENTATION

SECOND EDITION

MATT PHARR

Intel

GREG HUMPHREYS

NVIDIA



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Morgan Kaufmann is an imprint of Elsevier



MORGAN KAUFMANN PUBLISHERS

REPLACE THIS PAGE WITH THE SEPARATE COPYRIGHT PAGE.

To Deirdre, who even let me bring the manuscript on our honeymoon.

M. P.

To Isabel and Leila, the two most extraordinary people I've ever met. May your pixels never be little squares.

G. H.

ABOUT THE AUTHORS

Matt Pharr is a Principal Engineer at Intel, working as the lead architect in the Advanced Rendering Technology group. He previously co-founded Neoptica, which worked on programming models for graphics on heterogeneous CPU+GPU systems; Neoptica was acquired by Intel. Before Neoptica, Matt was in the Software Architecture group at NVIDIA, co-founded Exluna, and worked in Pixar's Rendering R&D group. He received his Ph.D. from the Stanford Graphics Lab, working under the supervision of Pat Hanrahan. He was also the editor of *GPU Gems 2*.

Greg Humphreys is a member of the OptiX raytracing team at NVIDIA. Previously, he was a professor of Computer Science at the University of Virginia, where he conducted research in both high performance and physically based computer graphics, as well as computer architecture and visualization. Greg has a B.S.E. degree from Princeton, and a Ph.D. in Computer Science from Stanford under the supervision of Pat Hanrahan. When he's not tracing rays, Greg enjoys tournament bridge, knitting, and riding his motorcycle.

Contents

PREFACE

xix

CHAPTER 01. INTRODUCTION 1

1.1	Literate Programming	1
1.1.1	Indexing and Cross-Referencing	3
1.2	Photorealistic Rendering and the Ray-Tracing Algorithm	4
1.2.1	Cameras	5
1.2.2	Ray-Object Intersections	7
1.2.3	Light Distribution	8
1.2.4	Visibility	9
1.2.5	Surface Scattering	10
1.2.6	Recursive Ray Tracing	11
1.2.7	Ray Propagation	13
1.3	pbrt: System Overview	15
1.3.1	Phases of Execution	16
1.3.2	Scene Representation	18
1.3.3	Renderer Interface and SamplerRenderer	24
1.3.4	The Main Rendering Loop	26
1.3.5	Parallelization of pbrt	35
1.3.6	An Integrator for Whitted Ray Tracing	41
1.4	How to Proceed through This Book	47
1.4.1	The Exercises	48
1.5	Using and Understanding the Code	48
1.5.1	Pointer or Reference?	48
1.5.2	Code Optimization	49
1.5.3	The Book Web site	49
1.5.4	Extending the System	49
1.5.5	Bugs	50
	Further Reading	50
	Exercise	52

CHAPTER 02. GEOMETRY AND TRANSFORMATIONS 55

2.1	Coordinate Systems	55
2.1.1	Coordinate System Handedness	56
2.2	Vectors	57
2.2.1	Arithmetic	58
2.2.2	Scaling	59
2.2.3	Dot and Cross Product	60

2.2.4	Normalization	62
2.2.5	Coordinate System from a Vector	63
2.3	Points	63
2.4	Normals	65
2.5	Rays	66
2.5.1	Ray Differentials	68
2.6	Three-Dimensional Bounding Boxes	70
2.7	Transformations	74
2.7.1	Homogeneous Coordinates	75
2.7.2	Basic Operations	76
2.7.3	Translations	77
2.7.4	Scaling	79
2.7.5	x , y , and z Axis Rotations	80
2.7.6	Rotation around an Arbitrary Axis	82
2.7.7	The Look-At Transformation	84
2.8	Applying Transformations	85
2.8.1	Points	85
2.8.2	Vectors	86
2.8.3	Normals	86
2.8.4	Rays	88
2.8.5	Bounding Boxes	88
2.8.6	Composition of Transformations	88
2.8.7	Transformations and Coordinate System Handedness	89
* 2.9	Animating Transformations	90
2.9.1	Quaternions	92
2.9.2	Quaternion Interpolation	94
2.9.3	Animated Transform Implementation	96
2.10	Differential Geometry	101
	Further Reading	103
	Exercises	104
CHAPTER 03. SHAPES		107
3.1	Basic Shape Interface	107
3.1.1	Bounding	109
3.1.2	Refinement	110
3.1.3	Intersection	110
3.1.4	Avoiding Incorrect Self-Intersections	111
3.1.5	Shading Geometry	112
3.1.6	Surface Area	113
3.1.7	Sidedness	113
3.2	Spheres	113
3.2.1	Construction	115
3.2.2	Bounding	116

	3.2.3	Intersection	116
	3.2.4	Partial Spheres	119
	* 3.2.5	Partial Derivatives of Normal Vectors	121
	3.2.6	Differential Geometry Initialization	122
	3.2.7	Surface Area	123
3.3	Cylinders		124
	3.3.1	Construction	125
	3.3.2	Bounding	126
	3.3.3	Intersection	126
	3.3.4	Partial Cylinders	127
	3.3.5	Surface Area	129
3.4	Disks		129
	3.4.1	Construction	130
	3.4.2	Bounding	130
	3.4.3	Intersection	131
	3.4.4	Surface Area	132
3.5	Other Quadrics		133
	3.5.1	Cones	133
	3.5.2	Paraboloids	134
	3.5.3	Hyperboloids	134
3.6	Triangles and Meshes		135
	3.6.1	Triangle	139
	3.6.2	Triangle Intersection	140
	3.6.3	Surface Area	145
	3.6.4	Shading Geometry	145
* 3.7	Subdivision Surfaces		148
	3.7.1	Mesh Representation	151
	3.7.2	Bounds	160
	3.7.3	Subdivison	160
		Further Reading	174
		Exercises	175

CHAPTER 04. PRIMITIVES AND INTERSECTION ACCELERATION

183

4.1	Primitive Interface and Geometric Primitives	184	
	4.1.1	Geometric Primitives	187
	4.1.2	Transformed Primitive: Object Instancing and Animated Primitives	189
4.2	Aggregates	191	
	4.2.1	Ray-Box Intersections	193
4.3	Grid Accelerator	195	
	4.3.1	Creation	197
	4.3.2	Traversal	202
4.4	Bounding Volume Hierarchies	208	
	4.4.1	BVH Construction	210
	4.4.2	The Surface Area Heuristic	217

4.4.3	Compact BVH For Traversal	222
4.4.4	Traversal	224
4.5	Kd-Tree Accelerator	227
4.5.1	Tree Representation	229
4.5.2	Tree Construction	231
4.5.3	Traversal	240
4.6	Debugging Aggregates	245
4.6.1	Finding Bugs in Aggregates	246
4.6.2	Fixing Bugs In Aggregates	249
4.6.3	Aggregate Performance Bugs	250
	Further Reading	250
	Exercises	255
 CHAPTER 05. COLOR AND RADIOMETRY		 261
5.1	Spectral Representation	261
5.1.1	The Spectrum Type	263
5.1.2	CoefficientSpectrum Implementation	263
5.2	The SampledSpectrum Class	266
5.2.1	XYZ Color	270
5.2.2	RGB Color	273
5.3	RGBSpectrum Implementation	279
5.4	Basic Radiometry	281
5.4.1	Basic Quantities	282
5.4.2	Incident and Exitant Radiance Functions	286
5.4.3	Luminance and photometry	287
5.5	Working with Radiometric Integrals	288
5.5.1	Integrals over Projected Solid Angle	289
5.5.2	Integrals over Spherical Coordinates	290
5.5.3	Integrals over Area	292
5.6	Surface Reflection	293
5.6.1	The BRDF	294
5.6.2	The BSSRDF	296
	Further Reading	297
	Exercises	298
 CHAPTER 06. CAMERA MODELS		 301
6.1	Camera Model	301
6.1.1	Camera Coordinate Spaces	303
6.2	Projective Camera Models	305
6.2.1	Orthographic Camera	306
6.2.2	Perspective Camera	310
6.2.3	Depth of Field	313
6.3	Environment Camera	318

	Further Reading	320
	Exercises	320
	CHAPTER 07. SAMPLING AND RECONSTRUCTION	323
7.1	Sampling Theory	323
7.1.1	The Frequency Domain and the Fourier Transform	325
7.1.2	Ideal Sampling and Reconstruction	327
7.1.3	Aliasing	331
7.1.4	Antialiasing Techniques	332
7.1.5	Application to Image Synthesis	335
7.1.6	Sources of Aliasing in Rendering	336
7.1.7	Understanding Pixels	337
7.2	Image Sampling Interface	338
7.2.1	Sample Representation and Allocation	342
7.3	Stratified Sampling	346
* 7.4	Low-Discrepancy Sampling	359
7.4.1	Definition of Discrepancy	359
7.4.2	Hammersley and Halton Sequences	361
7.4.3	(0,2)-Sequences	368
7.4.4	The Low-Discrepancy Sampler	372
* 7.5	Best-Candidate Sampling Patterns	378
7.6	Adaptive Sampling	385
7.7	Image Reconstruction	389
7.7.1	Filter Functions	393
7.8	Film and the Imaging Pipeline	402
7.8.1	Film Interface	403
7.8.2	ImageFilm	404
	Further Reading	413
	Exercises	417
	CHAPTER 08. REFLECTION MODELS	423
8.1	Basic Interface	428
8.1.1	Reflectance	430
8.1.2	BRDF \rightarrow BTDF Adapter	430
8.1.3	BxDF Scaling Adapter	431
8.2	Specular Reflection and Transmission	432
8.2.1	Fresnel Reflectance	434
8.2.2	Specular Reflection	439
8.2.3	Specular Transmission	442
8.3	Lambertian Reflection	446
8.4	Microfacet Models	447
8.4.1	Oren-Nayar Diffuse Reflection	449
8.4.2	Torrance-Sparrow Model	452

8.4.3	Blinn Microfacet Distribution	455
8.4.4	Anisotropic Microfacet Model	457
8.5	Fresnel Incidence Effects	460
8.6	Measured BRDFs	462
8.6.1	Irregular Isotropic Measured BRDF	463
8.6.2	Regular Halfangle Format	467
	Further Reading	470
	Exercises	472
CHAPTER 09. MATERIALS		477
9.1	BSDFs	477
9.1.1	BSDF Memory Management	482
9.2	Material Interface and Implementations	483
9.2.1	MatteMaterial	484
9.2.2	PlasticMaterial	486
9.2.3	Mix Material	487
9.2.4	Measured Material	489
9.2.5	Additional Materials	490
9.3	Bump Mapping	490
	Further Reading	496
	Exercises	497
CHAPTER 10. TEXTURE		501
10.1	Sampling and Antialiasing	503
10.1.1	Finding the Texture Sampling Rate	503
10.1.2	Filtering Texture Functions	509
* 10.1.3	Ray Differentials for Specular Reflection and Transmission	510
10.2	Texture Coordinate Generation	513
10.2.1	2D (u, v) Mapping	514
10.2.2	Spherical Mapping	515
10.2.3	Cylindrical Mapping	517
10.2.4	Planar Mapping	518
10.2.5	3D Mapping	518
10.3	Texture Interface and Basic Textures	519
10.3.1	Constant Texture	520
10.3.2	Scale Texture	520
10.3.3	Mix Textures	521
10.3.4	Bilinear Interpolation	522
10.4	Image Texture	523
10.4.1	Texture Caching	524
10.4.2	MIP Maps	528
10.4.3	Isotropic Triangle Filter	536
* 10.4.4	Elliptically Weighted Average	539
10.5	Solid and Procedural Texturing	544
10.5.1	UV Texture	545

	10.5.2	Checkerboard	546
	10.5.3	Solid Checkerboard	552
10.6		Noise	553
	10.6.1	Perlin Noise	554
	10.6.2	Random Polka Dots	558
	10.6.3	Noise Idioms and Spectral Synthesis	560
	10.6.4	Bumpy and Wrinkled Textures	565
	10.6.5	Windy Waves	566
	10.6.6	Marble	567
		Further Reading	569
		Exercises	571
* CHAPTER 11. VOLUME SCATTERING			575
11.1		Volume Scattering Processes	575
	11.1.1	Absorption	576
	11.1.2	Emission	578
	11.1.3	Out-Scattering and Extinction	579
	11.1.4	In-scattering	581
11.2		Phase Functions	583
11.3		Volume Interface and Homogeneous Media	587
	11.3.1	Homogeneous Volumes	589
11.4		Varying-Density Volumes	591
	11.4.1	3D Grids	592
	11.4.2	Exponential Density	594
11.5		Volume Aggregates	596
11.6		The BSSRDF	598
	11.6.1	Subsurface Scattering Materials	599
		Further Reading	600
		Exercises	601
CHAPTER 12. LIGHT SOURCES			605
12.1		Light Interface	606
	12.1.1	Visibility Testing	608
12.2		Point Lights	609
	12.2.1	Spotlights	611
	12.2.2	Texture Projection Lights	614
	12.2.3	Goniophotometric Diagram Lights	618
12.3		Distant Lights	621
12.4		Area Lights	623
12.5		Infinite Area Lights	627
		Further Reading	632
		Exercises	633

CHAPTER 13. MONTE CARLO INTEGRATION I:		
BASIC CONCEPTS		637
13.1	Background and Probability Review	638
13.1.1	Continuous Random Variables	639
13.1.2	Expected Values and Variance	640
13.2	The Monte Carlo Estimator	641
13.3	Basic Sampling Of Random Variables	643
13.3.1	The Inversion Method	643
13.3.2	The Rejection Method	650
* 13.4	Metropolis Sampling	652
13.4.1	Basic Algorithm	652
13.4.2	Choosing Mutation Strategies	654
13.4.3	Start-up bias	655
13.4.4	Estimating Integrals With Metropolis Sampling	656
13.4.5	Example: One-Dimensional Setting	656
13.5	Transforming between Distributions	660
13.5.1	Transformation in Multiple Dimensions	661
13.5.2	Example: Polar Coordinates	661
13.5.3	Example: Spherical Coordinates	661
13.6	2D Sampling with Multidimensional Transformations	662
13.6.1	Example: Uniformly Sampling a Hemisphere	663
13.6.2	Example: Sampling a Unit Disk	665
13.6.3	Example: Cosine-Weighted Hemisphere Sampling	668
13.6.4	Example: Sampling a Triangle	670
13.6.5	Example: Piecewise-Constant 2D Distributions	671
	Further Reading	674
	Exercises	675
CHAPTER 14. MONTE CARLO INTEGRATION II:		
IMPROVING EFFICIENCY		679
14.1	Russian Roulette and Splitting	680
14.1.1	Splitting	681
14.2	Careful Sample Placement	682
14.2.1	Stratified Sampling	682
14.2.2	Quasi Monte Carlo	685
14.2.3	Warping Samples and Distortion	685
14.3	Bias	686
14.4	Importance Sampling	688
14.4.1	Multiple Importance Sampling	690
14.5	Sampling Reflection Functions	693
14.5.1	Sampling the Blinn Microfacet Distribution	695
14.5.2	Sampling the Anisotropic Microfacet Model	699
14.5.3	Sampling Fresnel Blend	701
14.5.4	Specular Reflection and Transmission	702

	14.5.5	Application: Estimating Reflectance	703
	14.5.6	Sampling BSDFs	705
14.6		Sampling Light Sources	708
	14.6.1	Basic Interface	709
	14.6.2	Lights with Singularities	711
	14.6.3	Area Lights	715
	14.6.4	ShapeSet Sampling	722
	14.6.5	Infinite Area Lights	724
* 14.7		Volume Scattering	730
	14.7.1	Sampling Phase Functions	731
	14.7.2	Computing Optical Thickness	731
		Further Reading	733
		Exercises	734
 CHAPTER 15. LIGHT TRANSPORT I: SURFACE REFLECTION			 739
15.1		Direct Lighting	741
	15.1.1	Estimating the Direct Lighting Integral	747
15.2		The Light Transport Equation	751
	15.2.1	Basic Derivation	751
	15.2.2	Analytic Solutions to the LTE	753
	15.2.3	The Surface Form of the LTE	754
	15.2.4	Integral over Paths	755
	15.2.5	Delta distributions in the Integrand	757
	15.2.6	Partitioning the Integrand	758
	15.2.7	The Measurement Equation and Importance	759
15.3		Path Tracing	760
	15.3.1	Overview	762
	15.3.2	Path Sampling	763
	15.3.3	Incremental Path Construction	765
	15.3.4	Implementation	765
	* 15.3.5	Bidirectional Path Tracing	770
15.4		Instant Global Illumination	772
	15.4.1	Creating the virtual light sources	775
	15.4.2	Rendering with virtual light sources	780
15.5		Irradiance Caching	784
	15.5.1	Rendering With The Irradiance Cache	789
	15.5.2	Lookup and Interpolation	791
	15.5.3	Adding new values	794
15.6		Particle Tracing and Photon Mapping	797
	* 15.6.1	Theoretical Basis for Particle Tracing	798
	15.6.2	Photon Integrator	800
	15.6.3	Building the Photon Maps	803
	15.6.4	Using the Photon Map	817
	15.6.5	Photon Interpolation and Density Estimation	827

* 15.7	Metropolis Light Transport	833
	15.7.1 Sample Representation	838
	15.7.2 Mutations	839
	15.7.3 Generating Paths	842
	15.7.4 Path Contributions	845
	15.7.5 MetropolisRenderer Implementation	852
	15.7.6 Rendering	857
	Further Reading	861
	Exercises	865
* CHAPTER 16. LIGHT TRANSPORT II:	VOLUME RENDERING	873
16.1	The Equation of Transfer	873
16.2	Volume Integrator Interface	876
16.3	Emission-Only Integrator	876
16.4	Single Scattering Integrator	882
16.5	Subsurface Scattering	885
	16.5.1 Poisson Distribution of Sample Points	888
	16.5.2 Building the Sample Point Octree	897
	16.5.3 The Dipole Diffusion Approximation	901
	16.5.4 Rendering with Hierarchical Integration	908
	16.5.5 Setting Scattering Properties	912
	Further Reading	913
	Exercises	916
* CHAPTER 17. LIGHT TRANSPORT III:	PRECOMPUTED LIGHT TRANSPORT	925
17.1	Basis Functions: Theory	927
	17.1.1 A Piecewise-Constant Basis	928
	17.1.2 Projection Onto a Basis	930
	17.1.3 Orthonormal Basis Functions	930
17.2	Spherical Harmonics	932
	17.2.1 Efficient Evaluation	935
	17.2.2 Projecting Light Sources	941
	17.2.3 Projecting Incident Radiance Functions	949
	17.2.4 Reducing Ringing	950
	17.2.5 Rotations	951
17.3	Radiance Probes	956
	17.3.1 Creating Radiance Probes	956
	17.3.2 Using Radiance Probes	965
17.4	Precomputed Diffuse Transfer	969
17.5	Precomputed Glossy Transfer	974
	17.5.1 The Transfer Matrix	975

17.5.2	The BSDF Matrix	977
17.5.3	GlossyPRTIntegrator Implementation	979
	Further Reading	982
	Exercises	984

CHAPTER 18. RETROSPECTIVE AND THE FUTURE 989

18.1	Design Retrospective	989
18.1.1	Abstraction versus Efficiency	991
18.1.2	Design Alternative: Triangles Only	991
18.1.3	Increased Scene Complexity	992
18.2	Throughput Processors	993
18.2.1	The Future	995
18.2.2	Further Resources	995
18.3	Conclusion	997

APPENDIXES

A	UTILITIES	999
B	SCENE DESCRIPTION INTERFACE	1045
C	INDEX OF FRAGMENTS	1077
D	INDEX OF CLASSES AND THEIR MEMBERS	1091
E	INDEX OF MISCELLANEOUS IDENTIFIERS	1101
	REFERENCES	1105
	INDEX	1143
	COLOPHON	1169

Preface

[Just as] other information should be available to those who want to learn and understand, program source code is the only means for programmers to learn the art from their predecessors. It would be unthinkable for playwrights not to allow other playwrights to read their plays [or to allow them] at theater performances where they would be barred even from taking notes. Likewise, any good author is well read, as every child who learns to write will read hundreds of times more than it writes. Programmers, however, are expected to invent the alphabet and learn to write long novels all on their own. Programming cannot grow and learn unless the next generation of programmers has access to the knowledge and information gathered by other programmers before them. —Erik Naggum

Rendering is a fundamental component of computer graphics. At the highest level of abstraction, rendering is the process of converting a description of a three-dimensional scene into an image. Algorithms for animation, geometric modeling, texturing, and other areas of computer graphics all must pass their results through some sort of rendering process so that they can be made visible in an image. Rendering has become ubiquitous; from movies to games and beyond, it has opened new frontiers for creative expression, entertainment, and visualization.

In the early years of the field, research in rendering focused on solving fundamental problems such as determining which objects are visible from a given viewpoint. As effective solutions to these problems have been found and as richer and more realistic scene descriptions have become available thanks to continued progress in other areas of graphics, modern rendering has grown to include ideas from a broad range of disciplines, including physics and astrophysics, astronomy, biology, psychology and the study of perception, and pure and applied mathematics. The interdisciplinary nature of rendering is one of the reasons that it is such a fascinating area of study.

This book presents a selection of modern rendering algorithms through the documented source code for a complete rendering system. All of the images in this book, including the one on the front cover, were rendered by this software. All of the algorithms that came together to generate these images are described in these pages. The system, `pbrrt`, is written using a programming methodology called *literate programming* that mixes prose describing the system with the source code that implements it. We believe that the literate programming approach is a valuable way to introduce ideas in computer graphics and computer science in general. Often, some of the subtleties of an algorithm can be unclear or hidden until it is implemented, so seeing an actual implementation is a good way to acquire a solid understanding of that algorithm's details. Indeed, we believe that deep understanding of a small number of algorithms in this manner provides a stronger base for further study of computer graphics than does superficial understanding of many.

In addition to clarifying how an algorithm is implemented in practice, presenting these algorithms in the context of a complete and nontrivial software system also allows us to address issues in the design and implementation of medium-sized rendering systems. The design of a rendering system's basic abstractions and interfaces has substantial implications for both the elegance of the implementation and the ability to extend it later, yet the trade-offs in this design space are rarely discussed.

`pbrt` and the contents of this book focus exclusively on *photorealistic rendering*, which can be defined variously as the task of generating images that are indistinguishable from those that a camera would capture in a photograph, or as the task of generating images that evoke the same response from a human observer as looking at the actual scene. There are many reasons to focus on photorealism. Photorealistic images are crucial for the movie special-effects industry because computer-generated imagery must often be mixed seamlessly with footage of the real world. In entertainment applications where all of the imagery is synthetic, photorealism is an effective tool for making the observer forget that he or she is looking at an environment that does not actually exist. Finally, photorealism gives a reasonably well-defined metric for evaluating the quality of the rendering system's output.

A consequence of our approach is that this book and the system it describes do not exhaustively cover the state-of-the-art in rendering; many interesting topics in photorealistic rendering will not be introduced either because they don't fit well with the architecture of the software system (e.g., finite-element radiosity algorithms) or because we believed that the pedagogical value of explaining the algorithm was outweighed by the complexity of its implementation. We will note these decisions as they come up and provide pointers to further resources so that the reader can follow up on topics of interest. Many other areas of rendering, including interactive rendering, visualization, and illustrative forms of rendering such as pen-and-ink styles, aren't covered in this book at all. Nevertheless, many of the algorithms and ideas in this system (e.g., algorithms for texture map anti-aliasing) are applicable to a wider set of rendering styles.

AUDIENCE

Our primary intended audience for this book is students in graduate or upper-level undergraduate computer graphics classes. This book assumes existing knowledge of computer graphics at the level of an introductory college-level course, although certain key concepts such as basic vector geometry and transformations will be reviewed here. For students who do not have experience with programs that have tens of thousands of lines of source code, the literate programming style gives a gentle introduction to this complexity. We pay special attention to explaining the reasoning behind some of the key interfaces and abstractions in the system in order to give these readers a sense of why the system is structured in the way that it is.

Our secondary, but equally important, audiences are advanced graduate students and researchers, software developers in industry, and individuals interested in the fun of writing their own rendering systems. Although many of the ideas in this book will likely be familiar to these readers, seeing explanations of the algorithms presented in the literate style may provide new perspectives. `pbrt` includes implementations of a number of

advanced and/or difficult-to-implement algorithms and techniques, such as subdivision surfaces, Monte Carlo light transport and Metropolis sampling, subsurface scattering, and precomputed light transport algorithms; these should be of particular interest to experienced practitioners in rendering. We hope that delving into one particular organization of a complete and nontrivial rendering system will also be thought provoking to this audience.

OVERVIEW AND GOALS

pbrt is based on the *ray-tracing* algorithm. Ray tracing is an elegant technique that has its origins in lens making; Carl Freidrich Gauss traced rays through lenses by hand in the 19th century. Ray-tracing algorithms on computers follow the path of infinitesimal rays of light through the scene until they intersect a surface. This approach gives a simple method for finding the first visible object as seen from any particular position and direction, and is the basis for many rendering algorithms.

pbrt was designed and implemented with three main goals in mind: it should be *complete*, it should be *illustrative*, and it should be *physically based*.

Completeness implies that the system should not lack key features found in high-quality commercial rendering systems. In particular, it means that important practical issues, such as antialiasing, robustness, and the ability to efficiently render complex scenes, should all be addressed thoroughly. It is important to consider these issues from the start of the system's design, since these features can have subtle implications for all components of the system and can be quite difficult to retrofit into the system at a later stage of implementation.

Our second goal means that we tried to choose algorithms, data structures, and rendering techniques with care and with an eye toward readability and clarity. Since their implementations will be examined by more readers than is the case for many other rendering systems, we tried to select the most elegant algorithms that we were aware of and implement them as well as possible. This goal also required that the system be small enough for a single person to understand completely. We have implemented pbrt using an extensible architecture, with the core of the system implemented in terms of a set of carefully-designed abstract base classes, and as much of the specific functionality as possible in implementations of these base classes. The result is that one doesn't need to understand all of the specific implementations in order to understand the basic structure of the system. This makes it easier to delve deeply into parts of interest and skip others, without losing sight of how the overall system fits together.

There is a tension between the two goals of being complete and being illustrative. Implementing and describing every possible useful technique would not only make this book extremely long, but also would make the system prohibitively complex for most readers. In cases where pbrt lacks a particularly useful feature, we have attempted to design the architecture so that the feature could be added without altering the overall system design.

The basic foundations for physically based rendering are the laws of physics and their mathematical expression. pbrt was designed to use the correct physical units and concepts for the quantities it computes and the algorithms it implements. When configured

to do so, pbrt can compute images that are *physically correct*; they accurately reflect the lighting as it would be in a real-world version of the scene. One advantage of the decision to use a physical basis is that it gives a concrete standard of program correctness: for simple scenes, where the expected result can be computed in closed form, if pbrt doesn't compute the same result, we know there must be a bug in the implementation. Similarly, if different physically based lighting algorithms in pbrt give different results for the same scene, or if pbrt doesn't give the same results as another physically based renderer, there is certainly an error in one of them. Finally, we believe that this physically based approach to rendering is valuable because it is rigorous. When it is not clear how a particular computation should be performed, physics gives an answer that guarantees a consistent result.

Efficiency was given lower priority than these three goals. Since rendering systems often run for many minutes or hours in the course of generating an image, efficiency is clearly important. However, we have mostly confined ourselves to *algorithmic* efficiency rather than low-level code optimization. In some cases, obvious micro-optimizations take a backseat to clear, well-organized code, although we did make some effort to optimize the parts of the system where most of the computation occurs.

In the course of presenting pbrt and discussing its implementation, we hope to convey some hard-learned lessons from years of rendering research and development. There is more to writing a good renderer than stringing together a set of fast algorithms; making the system both flexible and robust is a difficult task. The system's performance must degrade gracefully as more geometry or light sources are added to it, or as any other axis of complexity is pushed. Numeric stability must be handled carefully, and algorithms that don't waste floating-point precision are critical.

The rewards for developing a system that addresses all these issues are enormous—it is a great pleasure to write a new renderer or add a new feature to an existing renderer and use it to create an image that couldn't be generated before. Our most fundamental goal in writing this book was to bring this opportunity to a wider audience. Readers are encouraged to use the system to render the example scenes in the pbrt software distribution as they progress through the book. Exercises at the end of each chapter suggest modifications to the system that will help clarify its inner workings, and more complex projects to extend the system by adding new features.

The website for this book is located at www.pbrt.org. The latest version of the pbrt source code is available from this site and we will also post errata and bug fixes, additional scenes to render, and supplemental utilities. Any bugs in pbrt or errors in this text that are not listed at the website can be reported to the email address bugs@pbrt.org. We greatly value your feedback!

CHANGES SINCE THE FIRST EDITION

Six years have passed since the publication of the first edition of this book. In that time, thousands of copies of the book have been sold and the pbrt software has been downloaded thousands of times from the book's website. The pbrt user base has given

us a significant amount of feedback and encouragement, and our experience with the system guided many of the decisions we made in making changes between the version of `pbrt` presented in the first edition and the updated version described here. In addition to a number of bug fixes, we also made several significant design changes and enhancements:

1. Removal of the plugin architecture. The first version of `pbrt` used a run-time plugin architecture to dynamically load code for implementations of objects like shapes, lights, integrators, cameras, and other objects that were used in the scene currently being rendered. This approach allowed users to extend `pbrt` with new object types (e.g., new shape primitives) without recompiling the entire rendering system. This approach initially seemed elegant, but it complicated the task of supporting `pbrt` on multiple platforms and it made debugging more difficult. The only new usage scenario that it truly enabled (binary-only distributions of `pbrt` or binary plugins) was actually contrary to our pedagogical and open-source goals. Therefore, the plugin architecture was dropped in this edition.
2. Removal of the image processing pipeline. The first version of `pbrt` provided a tone-mapping interface that converted high dynamic range (HDR) floating point output images directly into low dynamic range TIFFs for display. This functionality made sense in 2004, as support for HDR images was still sparse. In 2010, however, advances in digital photography have made HDR images commonplace. Although the theory and practice of tone mapping are elegant and worth learning, we decided to focus the new book exclusively on the process of image formation and ignore the topic of image display. Interested readers should read the book written by Reinhard et al. (2005) for a thorough and modern treatment of the HDR image display process.
3. Task parallelism. Multicore architectures are now ubiquitous, and we felt that `pbrt` would not remain relevant without the ability to scale to the number of locally available cores. We also hope that the parallel programming implementation details documented in this book will help graphics programmers understand some of the subtleties and complexities in writing scalable parallel code (e.g., choosing appropriate task granularities or mutex types), which is still a difficult and too-infrequently taught topic.
4. Appropriateness for “production” rendering. The first version of `pbrt` was intended exclusively as a pedagogical tool and a stepping-stone for rendering research. Indeed, we made a number of decisions in preparing the first edition that were contrary to use in a production environment, such as limited support for image-based lighting, no support for motion blur, and a photon mapping implementation that wasn’t robust in the presence of complex lighting. With much improved support for these features as well as support for subsurface scattering and Metropolis light transport, we feel that `pbrt` is now much more suitable for rendering very high-quality images of complex environments as it is presented here. The tradeoff in making these improvements is that as the system becomes more feature-complete, it may be harder for instructors to use the new software to create manageable assignments for students. While this is a real concern, we had similar reservations about the first version of `pbrt` “relieving” students of the burden and benefits of writing their own ray-tracing system from scratch. With

experience from the first edition being used at many universities, we have come to believe that this tradeoff was a good one, and we hope and expect that the new edition will continue to enable high-quality rendering courses.

ACKNOWLEDGMENTS

Pat Hanrahan has contributed to this book in more ways than we could hope to acknowledge; we owe a profound debt to him. He tirelessly argued for clean interfaces and finding the right abstractions to use throughout the system, and his understanding of and approach to rendering deeply influenced its design. His willingness to use `pbrt` and this manuscript in his rendering course at Stanford was enormously helpful, particularly in the early years of its life when it was still in very rough form; his feedback throughout this process has been crucial for bringing the text to its current state. Finally, the group of people that Pat helped assemble at the Stanford Graphics Lab, and the open environment that he fostered, made for an exciting, stimulating, and fertile environment. We feel extremely privileged to have been there.

We owe a debt of gratitude to the many students who used early drafts of this book in courses at Stanford and the University of Virginia between 1999 and 2004. These students provided an enormous amount of feedback about the book and `pbrt`. The teaching assistants for these courses deserve special mention: Tim Purcell, Mike Cammarano, Ian Buck, and Ren Ng at Stanford, and Nolan Goodnight at Virginia. A number of students in those classes gave particularly valuable feedback and sent bug reports and bug fixes; we would especially like to thank Evan Parker and Phil Beatty. A draft of the manuscript of this book was used in classes taught by Bill Mark and Don Fussell at the University of Texas, Austin, and Raghu Machiraju at Ohio State University; their feedback was invaluable, and we are grateful for their adventurousness in incorporating this system into their courses, even while it was still being edited and revised.

Matt Pharr would like to acknowledge colleagues and co-workers in rendering-related endeavors who have been a great source of education and who have substantially influenced his approach to writing renderers and his understanding of the field. Particular thanks go to Craig Kolb, who provided a cornerstone of Matt's early computer graphics education through the freely available source code to the `rayshade` ray-tracing system, and Eric Veach, who has also been generous with his time and expertise. Thanks also to Doug Shult and Stan Eisenstat for formative lessons in mathematics and computer science during high school and college, respectively, and most importantly to Matt's parents, for the education they've provided and continued encouragement along the way. Finally, thanks also to Nick Triantos, Jayant Kolhe, and NVIDIA for their understanding and support through the final stages of the preparation of the first edition of the book.

Greg Humphreys is very grateful to all the professors and TAs who tolerated him when he was an undergraduate at Princeton. Many people encouraged his interest in graphics, specifically Michael Cohen, David Dobkin, Adam Finkelstein, Michael Cox, Gordon Stoll, Patrick Min, and Dan Wallach. Doug Clark, Steve Lyon, and Andy Wolfe also supervised various independent research boondoggles without even laughing once. Once, in a group meeting about a year-long robotics project, Steve Lyon became exasperated and yelled, "Stop telling me why it can't be done, and figure out how to do it!"—an im-

promptu lesson that will never be forgotten. Eric Ristad fired Greg as a summer research assistant after his freshman year (before the summer even began), pawning him off on an unsuspecting Pat Hanrahan and beginning an advising relationship that would span 10 years and both coasts. Finally, Dave Hanson taught Greg that literate programming was a great way to work, and that computer programming can be a beautiful and subtle art form.

We are also grateful to Don Mitchell, for his help with understanding some of the details of sampling and reconstruction; Thomas Kollig and Alexander Keller, for explaining the finer points of low-discrepancy sampling; and Dave Eberly, “Just d’FAQs,” Hans-Bernhard Broeker, Steve Westin, and Gernot Hoffmann, for many interesting threads on *comp.graphics.algorithms*. Christer Ericson had a number of suggestions for improving our kd-tree implementation. Christophe Hery helped us with understanding the nuances of subsurface scattering and Peter-Pike Sloan was kind enough to carefully review Chapter 17 on precomputed light transport algorithms.

Many people and organizations have generously supplied us with scenes and models for use in this book and the *pbrt* distribution. Their generosity has been invaluable in helping us create interesting example images throughout the text. The bunny, Buddha, and dragon models are courtesy of the Stanford Computer Graphics Laboratory’s scanning repository at graphics.stanford.edu/data/3Dscanrep/. The ecosystem scene was created by Oliver Deussen and Bernd Lintermann for a paper by them and collaborators (Deussen, Hanrahan, Lintermann, Mech, Pharr, and Prusinkiewicz 1998). The “killeroo” model is included with permission of Phil Dench and Martin Rezard (3D scan and digital representations by headus, design and clay sculpt by Rezard). The physically accurate smoke data sets were created by Duc Nguyen and Ron Fedkiw. Nolan Goodnight created environment maps with a realistic skylight model. The Cornell Program of Computer Graphics Light Measurement Laboratory allowed us to include measured BRDF data, and Paul Debevec provided numerous high dynamic-range environment maps. Marc Ellens provided spectral data for a variety of light sources, and the spectral RGB measurement data for a variety of displays is courtesy of Tom Lianza at X-Rite.

We are particularly grateful to Guillermo M. Leal Llaguno of Evolución Visual, www.evvisual.com, who modeled and rendered the San Miguel scene featured on the cover and in numerous figures in the book. We would also especially like to thank Marko Dabrovic (www.3lhd.com) and Mihovil Odak at RNA Studios (www.rna.hr), who supplied a bounty of excellent models and scenes, including the Sponza atrium, the Sibenik cathedral, and the Audi TT car model. Many thanks are also due to Florent Boyer (www.florentboyer.com), who provided the contemporary house scene used in many of the images in Chapter 15.

We would also like to thank the book’s reviewers, all of whom had insightful and constructive feedback about the manuscript at various stages of its progress. We’d particularly like to thank the reviewers who provided feedback on both editions of the book: Ian Ashdown, Per Christensen, Doug Epps, Dan Goldman, Eric Haines, Erik Reinhard, Pete Shirley, Peter-Pike Sloan, Greg Ward, and a host of anonymous reviewers. For the second edition, Janne Kontkanen, Nelson Max, Bill Mark, and Eric Tabellion also contributed numerous helpful suggestions.

We'd like to thank the faculty members at various universities who have used `pbrt` in their courses, including Emmanuel Agu, Dirk Arnold, Stephen Chenney, Yung-Yu Chuang, Don Fussell, Pat Hanrahan, Bill Mark, Nelson Max, Gary Meyer, Torsten Möller, Rick Parent, Sumanta Pattanaik, and Luiz Velho.

Many people have contributed to not only `pbrt` but to our own better understanding of rendering through bug reports, patches, and suggestions about better implementation approaches. A few have made particularly substantial contributions—we would especially like to thank Kevin Egan, John Danks, Volodymyr Kachurovskyi, Solomon Boulos, and Stephen Chenney. In addition, we would also like to thank Rachit Agrawal, Frederick Akalin, Mark Bolstad, Thomas de Bodt, Brian Budge, Mark Colbert, Shaohua Fan, Nigel Fisher, Jeppe Revall Frisvad, Robert G. Graf, Asbjørn Heid, Keith Jeffery, Greg Johnson, Aaron Karp, Donald Knuth, Martin Kraus, Murat Kurt, Larry Lai, Craig McNaughton, Swaminathan Narayanan, Anders Nilsson, Jens Olsson, Vincent Pegoraro, Nils Thuerey, Xiong Wei, Wei-Wei Xu, Arek Zimny, and Matthias Zwicker for their suggestions and bug reports. Finally, we would like to thank the *LuxRender* developers and the *LuxRender* community, particularly Terrence Vergauwen, Jean-Philippe Grimaldi, and Asbjørn Heid; it has been a delight to see the rendering system they have built from `pbrt`'s foundation, and we have learned from reading their source code and implementations of new rendering algorithms.

For the production of the first edition, we would also like to thank Tim Cox (senior editor), for his willingness to take on this slightly unorthodox project and for both his direction and patience throughout the process. We are very grateful to Elisabeth Beller (project manager), who has gone well beyond the call of duty for this book; her ability to keep this complex project in control and on schedule has been remarkable, and we particularly thank her for the measurable impact she has had on the quality of the final result. Thanks also to Rick Camp (editorial assistant) for his many contributions along the way. Paul Anagnostopoulos and Jacqui Scarlott at Windfall Software did the book's composition; their ability to take the authors' homebrew literate programming file format and turn it into high-quality final output while also juggling the multiple unusual types of indexing we asked for is greatly appreciated. Thanks also to Ken DellaPenta (copyeditor) and Jennifer McClain (proofreader) as well as to Max Spector at Chen Design (text and cover designer), and Steve Rath (indexer).

For the second edition, we'd like to thank Greg Chalson who talked us into expanding and updating the book; Greg also ensured that Paul Anagnostopoulos at Windfall Software would again do the book's composition. We'd like to thank Paul again for his efforts in working with this book's production complexity. Finally, we'd also like to thank Todd Green, Paul Gottehrer, and Heather Scherer at Elsevier.

ABOUT THE COVER

The “San Miguel” scene on the cover of the book was modeled and then rendered by Guillermo M. Leal Llaguno of Evolución Visual, www.evvisual.com, based on a hacienda that he visited in San Miguel de Allende, Mexico. The scene was modeled in *3ds max* and exported to the `pbrt` file format with a custom script written by Guillermo. The scene features just over 2.5 million unique triangles and has a total geometric complexity of

10.7 million triangles due to the use of object instancing; the `pbrt` files that describe the scene geometry require 620 MB of on-disk storage. There are a total of 354 texture maps, representing 293 MB of texture data. Final rendering of the cover image at 1496 by 2235 resolution using `pbrt` took over 40 hours of computation on an eight-core Mac Pro computer. The scene is available in the `scenes/sanmiguel` directory of the `pbrt` distribution.

ADDITIONAL READING

Donald Knuth's article *Literate Programming* (Knuth 1984) describes the main ideas behind literate programming as well as his web programming environment. The seminal \TeX typesetting system was written with web and has been published as a series of books (Knuth 1986, Knuth 1993a). More recently, Knuth has published a collection of graph algorithms in literate format in *The Stanford GraphBase* (Knuth 1993b). These programs are enjoyable to read and are excellent presentations of their respective algorithms. The website www.literateprogramming.com has pointers to many articles about literate programming, literate programs to download, and a variety of literate programming systems; many refinements have been made since Knuth's original development of the idea.

The only other literate programs we know of that have been published as books are the implementation of the `lcc` compiler, which was written by Christopher Fraser and David Hanson and published as *A Retargetable C Compiler: Design and Implementation* (Fraser and Hanson 1995), and Martin Ruckert's book on the `mp3` audio format, *Understanding MP3* (Ruckert 2005).