

# PRODUCTION VOLUME RENDERING

## DESIGN AND IMPLEMENTATION



MAGNUS WRENNINGE

 CRC Press  
Taylor & Francis Group  
AN A K PETERS BOOK

# Production Volume Rendering

This page intentionally left blank

# Production Volume Rendering

## Design and Implementation

Magnus Wrenninge



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

Cover design by Vincent Serritella.

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

© 2013 by Magnus Wrenninge  
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works  
Version Date: 20120710

International Standard Book Number-13: 978-1-4398-7363-2 (eBook - PDF)

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) (<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.

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

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

To Chia-Chi.

This page intentionally left blank

# Contents

Preface	xiii
<b>I The PVR System</b>	<b>I</b>
1.1 C++ and Python	1
1.1.1 Use of Namespaces	2
1.1.2 Use of External Libraries	3
1.2 Python Bindings	3
1.3 Rendering with PVR	5
<b>I Fundamentals</b>	<b>II</b>
<b>2 The Basics</b>	<b>13</b>
2.1 Time and Motion Blur	13
2.1.1 Render Globals	15
2.1.2 Function Curves	16
2.2 Cameras	18
2.2.1 Camera Coordinate Spaces	19
2.3 Geometry	20
2.3.1 The <b>Geometry</b> Class	20
2.3.2 The <b>Particles</b> Class	21
2.3.3 The <b>Polygons</b> Class	22
2.3.4 The <b>Meshes</b> Class	23
2.4 Geometry Attributes	24
2.5 Attribute Tables	25
2.5.1 String Attributes	28
2.6 Attribute References	29
2.7 Attribute Iteration	29
<b>3 Voxel Buffers</b>	<b>33</b>
3.1 Introduction to Voxel Buffers	33
3.1.1 Voxel Indexing	34
3.1.2 Extents and Data Window	34
3.1.3 Coordinate Spaces and Mappings	35
3.1.4 What Are the Coordinates of a Voxel?	36



3.2	Implementing a Simple Voxel Buffer . . . . .	37
3.3	Field3D . . . . .	38
3.3.1	The <b>DenseField</b> Class . . . . .	39
3.3.2	The <b>SparseField</b> Class . . . . .	40
3.4	Transformations and Mappings . . . . .	43
3.4.1	Uniform Transforms . . . . .	43
3.4.2	Frustum Transforms . . . . .	44
3.5	Interpolating Voxel Data . . . . .	45
3.5.1	Nearest-Neighbor Interpolation . . . . .	46
3.5.2	Linear Interpolation . . . . .	47
3.5.3	Cubic Interpolation . . . . .	49
3.5.4	Monotonic Cubic Interpolation . . . . .	53
3.6	Filtered Lookups . . . . .	55
3.6.1	Gaussian Filter . . . . .	55
3.6.2	Mitchell-Netravali Filter . . . . .	57
3.6.3	Performance Comparison . . . . .	59
<b>4</b>	<b>Noise</b> . . . . .	<b>61</b>
4.1	Procedural Textures . . . . .	61
4.2	Perlin Noise . . . . .	62
4.3	Noise Functions . . . . .	63
4.4	Fractal Functions . . . . .	64
4.5	The <b>Fractal</b> Base Class . . . . .	66
4.6	Fractional Brownian Motion: fBm . . . . .	66
4.6.1	Octaves . . . . .	67
4.6.2	Scale . . . . .	67
4.6.3	Octave Gain . . . . .	69
4.6.4	Lacunarity . . . . .	70
4.6.5	Value Range of Fractal Functions . . . . .	70
<b>II</b>	<b>Volume Modeling</b> . . . . .	<b>73</b>
<b>5</b>	<b>Fundamentals of Volume Modeling</b> . . . . .	<b>75</b>
5.1	Volume Modeling and Voxel Buffers . . . . .	75
5.2	Defining the Voxel Buffer . . . . .	76
5.2.1	Bounding Primitives . . . . .	76
5.2.2	Boundless Voxel Buffers . . . . .	77
5.3	Volume-Modeling Strategies . . . . .	77
5.3.1	Direct Voxel Access . . . . .	78
5.3.2	Splatting . . . . .	78
5.3.3	Rasterization . . . . .	78
5.3.4	Instantiation . . . . .	79
5.4	Rasterization Primitives . . . . .	79
5.5	Instantiation Primitives . . . . .	80

5.6	Using Geometry to Guide Volumetric Primitives . . . . .	82
5.6.1	Coordinate Systems . . . . .	82
5.6.2	Local-to-World versus World-to-Local . . . . .	83
5.7	Common Coordinate Systems . . . . .	85
5.7.1	Points/Spheres . . . . .	85
5.7.2	Lines . . . . .	86
5.7.3	Surfaces . . . . .	87
5.8	Procedural Noise and Fractal Functions . . . . .	88
5.8.1	Making Noise Stick . . . . .	88
5.8.2	Density Variation . . . . .	88
5.8.3	Impact on Primitive Bounds . . . . .	90
<b>6</b>	<b>PVR's Modeling Pipeline</b>	<b>91</b>
6.1	Overview . . . . .	91
6.2	The <b>Modeler</b> Class . . . . .	92
6.2.1	Adding Inputs . . . . .	94
6.2.2	Building a Uniform Mapping . . . . .	94
6.2.3	Building a Frustum Mapping . . . . .	95
6.2.4	Executing the Modeling Process . . . . .	96
6.2.5	Accessing the Voxel Buffer . . . . .	97
6.3	Inputs to the Modeler . . . . .	97
6.4	Handling User Parameters . . . . .	98
6.4.1	The <b>ParamMap</b> struct . . . . .	98
6.5	The <b>Primitive</b> Base Class . . . . .	99
6.5.1	Volumetric Primitive versus Underlying Primitive . . . . .	100
6.6	Splatting Data to Voxel Buffers . . . . .	101
6.6.1	Splatting a Point . . . . .	101
6.6.2	Splatting an Antialiased Point . . . . .	102
6.6.3	Splatting and Motion Blur . . . . .	104
<b>7</b>	<b>Rasterization Primitives in PVR</b>	<b>107</b>
7.1	The <b>RasterizationPrim</b> Base Class . . . . .	107
7.1.1	The Rasterization Loop . . . . .	108
7.1.2	Sampling Density from the Subclass . . . . .	109
7.1.3	Optimal Bounds . . . . .	110
7.1.4	Rasterization and Motion Blur . . . . .	110
7.2	Implementing Primitives . . . . .	111
7.2.1	Attributes . . . . .	111
7.2.2	A Design Pattern for Handling Attributes . . . . .	112
7.3	Sphere-Based Primitives . . . . .	114
7.3.1	Bounding the Primitive . . . . .	115
7.4	The <b>Point</b> Primitive . . . . .	116
7.4.1	Executing the Primitive . . . . .	117
7.4.2	Density Function . . . . .	121
7.4.3	Bounding a Single Point . . . . .	122

7.5	The <b>PyroclasticPoint</b> Primitive . . . . .	122
7.5.1	Executing the Primitive . . . . .	124
7.5.2	Density Function . . . . .	125
7.5.3	Bounding a Single Pyroclastic Point . . . . .	128
7.6	Line-Based Primitives . . . . .	128
7.6.1	Bounding a Line . . . . .	130
7.6.2	Executing the Primitive . . . . .	131
7.6.3	Updating the Acceleration Data Structure . . . . .	132
7.6.4	Finding the Closest Point on a Line . . . . .	133
7.6.5	Displacement Bounds . . . . .	134
7.6.6	Interpolating Attributes along the Line . . . . .	135
7.7	The <b>Line</b> Primitive . . . . .	135
7.7.1	Density Function . . . . .	136
7.8	The <b>PyroclasticLine</b> Primitive . . . . .	136
7.8.1	Density Function . . . . .	138
7.8.2	Transforming from World to Local Space . . . . .	141
7.8.3	Updating Per-Polygon Attributes . . . . .	141
7.8.4	Updating Per-Point Attributes . . . . .	142
7.8.5	Displacement Bounds . . . . .	142
8	<b>Instantiation Primitives in PVR</b> . . . . .	145
8.1	The <b>InstantiationPrim</b> Base Class . . . . .	145
8.2	Common Strategies . . . . .	146
8.2.1	Number of Points to Instance . . . . .	146
8.2.2	Local Coordinate Space . . . . .	146
8.2.3	Output from Point-Based Instantiation Primitives . . . . .	146
8.3	The <b>Sphere</b> Instantiation Primitive . . . . .	146
8.3.1	Executing the Primitive . . . . .	148
8.4	The <b>Line</b> Instantiation Primitive . . . . .	152
8.4.1	Executing the Primitive . . . . .	154
8.5	The <b>Surface</b> Instantiation Primitive . . . . .	158
8.5.1	Executing the Primitive . . . . .	160
8.5.2	Eroding the Edges . . . . .	163
III	<b>Volume Rendering</b> . . . . .	165
9	<b>Volumetric Lighting</b> . . . . .	167
9.1	Lighting Fundamentals . . . . .	167
9.2	Absorption . . . . .	169
9.3	Emission . . . . .	171
9.4	Scattering . . . . .	171
9.5	Phase Functions . . . . .	172
9.6	Optical Thickness and Transmittance . . . . .	174
9.7	Wavelength Dependency . . . . .	179
9.8	Other Approaches to Volume Rendering . . . . .	179

<b>10 Raymarching</b>	<b>181</b>
10.1 An Introduction to Raymarching . . . . .	181
10.2 Lighting and Raymarching . . . . .	185
10.3 Integration Intervals . . . . .	187
10.4 Integration Intervals for Multiple Volumes . . . . .	188
10.5 Integration Intervals for Overlapping Volumes . . . . .	190
10.6 Sampling Strategies . . . . .	191
10.7 Empty-Space Optimization . . . . .	192
10.8 Holdouts . . . . .	194
<b>11 PVR's Rendering Pipeline</b>	<b>203</b>
11.1 The <b>Scene</b> Class . . . . .	203
11.2 The <b>Renderer</b> Class . . . . .	204
11.2.1 Setting Up Rays . . . . .	205
11.2.2 Firing and Integrating Rays . . . . .	207
11.2.3 Executing the Render . . . . .	208
11.2.4 The <b>RayState</b> Struct . . . . .	210
11.2.5 The <b>RenderGlobals</b> Class . . . . .	212
11.3 The <b>Camera</b> Base Class . . . . .	213
11.3.1 Camera Transformations . . . . .	216
11.3.2 Transforming Points . . . . .	218
11.4 The <b>PerspectiveCamera</b> Class . . . . .	220
11.5 The <b>SphericalCamera</b> Class . . . . .	226
11.6 Image Output . . . . .	230
<b>12 PVR Volume Types</b>	<b>233</b>
12.1 Volumes in PVR . . . . .	233
12.1.1 Volume Properties and Attributes . . . . .	235
12.1.2 The <b>VolumeSampleState</b> Struct . . . . .	237
12.1.3 The <b>VolumeSample</b> Struct . . . . .	237
12.2 The <b>ConstantVolume</b> Class . . . . .	238
12.3 The <b>VoxelVolume</b> Class . . . . .	243
12.3.1 Attribute Handling . . . . .	244
12.3.2 Ray Intersection Testing . . . . .	244
12.3.3 Intersecting Uniform Buffers . . . . .	245
12.3.4 Intersecting Frustum Buffers . . . . .	248
12.3.5 Empty-Space Optimization . . . . .	250
12.3.6 Optimizing Sparse Uniform Buffers . . . . .	251
12.3.7 Optimizing Sparse Frustum Buffers . . . . .	257
12.3.8 Sampling the Buffer . . . . .	260
12.4 The <b>CompositeVolume</b> Class . . . . .	262

<b>13 Raymarching in PVR</b>	<b>267</b>
13.1 Introduction	267
13.2 The <b>Raymarcher</b> Base Class	268
13.2.1 The <b>IntegrationResult</b> Struct	269
13.3 The <b>UniformRaymarcher</b> Class	270
13.3.1 Ray Integration	270
13.3.2 Updating Transmittance and Luminance	276
13.4 Integration Intervals	281
13.4.1 Splitting Intervals	281
<b>14 Lighting in PVR</b>	<b>285</b>
14.1 Raymarch Samplers	285
14.1.1 The <b>RaymarchSample</b> Struct	286
14.2 The <b>DensitySampler</b> Class	286
14.3 The <b>PhysicalSampler</b> Class	288
14.4 The <b>Light</b> Base Class	292
14.4.1 Light Intensity	294
14.4.2 Falloff	295
14.5 Point Lights	296
14.6 Spot Lights	298
14.7 Phase Functions in PVR	300
14.7.1 The <b>PhaseFunction</b> Base Class	300
14.7.2 The Isotropic Phase Function	301
14.7.3 The Henyey-Greenstein Phase Function	301
14.7.4 The Double Henyey-Greenstein Phase Function	305
14.7.5 Composite Phase Functions	306
14.8 Occlusion in PVR	311
14.8.1 The Occluder Base Class	311
14.8.2 The <b>OcclusionSampleState</b> Struct	311
14.8.3 The <b>RaymarchOccluder</b> Class	313
<b>15 Precomputed Occlusion</b>	<b>315</b>
15.1 Voxelized Occlusion	315
15.2 Deep Shadows	316
15.3 Strategies for Precomputation	317
15.4 The <b>VoxelOccluder</b> Class	318
15.5 The <b>OtfVoxelOccluder</b> Class	321
15.6 The <b>DeepImage</b> Class	323
15.7 The <b>TransmittanceMapOccluder</b> Class	324
15.8 The <b>OtfTransmittanceMapOccluder</b> Class	327
<b>Bibliography</b>	<b>333</b>
<b>Index</b>	<b>335</b>
<b>Class Index</b>	<b>339</b>

# Preface

Production volume rendering, described in less than 15 words, refers to rendering of nonsurfaces for the purpose of creating images for film or animation production. A slightly longer explanation might include some examples, such as the creation and rendering of virtual smoke, fire, dust, and clouds. The “and” in “creation and rendering” is also an important point: production volume rendering refers just as much to the *modeling* of the effects as it refers to the actual *rendering*, and in fact, the modeling side is often the lesser known part of the two.

No matter which way you look at it, production volume rendering is an esoteric subject. Volume rendering in the general sense encompasses everything from medical visualization to real-time game graphics, but production volume rendering has very little overlap with these subjects. There are only a few publicly available pieces of software, and to make matters more complicated, there is also very little research or other documentation available, due to the fact that most development is done in-house at visual effects and animation production houses around the world.

The consequence of this lack of information is that each person who is tasked with working on a production volume renderer, or who just happens to be interested in the topic, has to start out from scratch, inventing most of the puzzle pieces that he or she needs to make their systems work.

A certain frustration with this lack of documentation and systemization led to a 2010 SIGGRAPH course called *Volumetric Methods in Visual Effects*, and the course, in turn, led to the book you now hold in your hand.

The course tried to give an overview of all the different techniques and ideas that are used in the volume-rendering solutions at some of the largest production facilities, and it also covered a wide range of concrete examples of those techniques, as they are implemented at those facilities.

This book takes over where the first part of the course left off and goes into all the same fundamental topics. But where the course stayed in a

generic context, with simple examples, this book provides and describes a systematic implementation of the techniques.

The book's approach is highly pragmatic. Its scope is limited to the techniques and algorithms that are actively used in production work. It leaves out much of the available research into photorealistic volume rendering, but the reason for doing so is simply that those techniques rarely get used in production work.

In order to ensure that enough detail is provided on how each technique works, the book is written around an open source renderer called PVR, which can be downloaded, compiled, and modified by the reader. This approach hopefully means that whatever questions are left by the book can be answered by looking at how the code is implemented.

## Goals

The goal of the book is to provide two paths towards understanding production volume rendering. On one side, it describes the techniques used for modern production volume rendering in a generic context. It shows how the techniques fit together, and how the modules that make it up are used to achieve real-world goals. But it wouldn't be a complete book if it did not also describe an implementation of those techniques. Showing how to translate the abstract set of concepts into concrete, working code is an important part of the equation. It shows that the ideas work. And it shows that they work together to create a complete system.

Throughout this book, the illustrations and rendered images are all created using the code that the book describes. In fact, the scripts and the data used for all the examples are freely available (along with the source code) at <http://www.github.com/pvrbook>.

The aim with this approach is to let the readers explore the book in two ways: Someone who is curious about how production volume rendering works in the big picture can start with reading the chapters that describe the fundamental ideas at play. At the same time, if the reader comes across an interesting image, he or she can go straight to the lowest level, starting with the script that created it. The reader can then trace his or her way back through the rendering code and the modeling primitives, which then hopefully illustrates the bigger picture that ties all the steps together.


The primary goal of the book is to be illustrative, and an important step towards that goal is making sure that the translation of the techniques and concepts into the working code is clear and that the code's modular structure properly reflects the ideas presented in the book. Secondly, the concepts outlined in the book and implemented in the code

must be completely integrated, in the sense that each part of the system should work with all the others, using the same fundamental concepts. There should not be any isolated ideas that only work on their own and cannot be integrated into the system.

Efficiency is not the primary concern of the system. In the choice between simpler code or faster performance, PVR strives to be easy to understand. That is not to say that the system's performance is ignored. On the contrary, scalability is a very important aspect of any system that is to be used in production. The goal of the system is instead to show how various design considerations have an impact on both scalability, extensibility, generality, and performance, but to do so in a straightforward and understandable way.



This page intentionally left blank



# The PVR System

The system described in this book is called PVR, standing simply for production volume rendering. Because production volume rendering is a fairly esoteric topic, it was clear early on that a working example of the algorithms and techniques in action had to be included. To not include one would be to rob the reader of the understanding that can sometimes only be had by seeing how a set of techniques covering different areas work in conjunction with one another.

## 1.1 C++ and Python

PVR is implemented entirely in C++. The choice of C++ was simple; although the PVR renderer isn't built primarily for speed, the amounts of data it needs to process makes C or C++ the only real choice. And because the system needed to be modularized so that concepts from the book could map cleanly to code, the object-oriented nature of C++ made it the best candidate.

While C++ has many strengths, runtime configurability isn't necessarily one of them. In order to make it easy for the reader to create scene descriptions, PVR also includes a Python module that wraps the entire code base (or at least the relevant 99% of it) so that scenes can be created and rendered with a few lines of Python code.

To exclude the Python interface (or some other scripting language bindings) would mean that the reader would have to compile a full C++ application in order to create a scene. We assume that the reader takes

no pleasure in decoding GCC's template error messages and, therefore, provide the bindings as a convenient, but powerful, scene-modeling and rendering language.

The intent of this book is to describe a simple but complete volume-modeling and rendering system. The focus is on illustrating the most important concepts as clearly as possible, and on describing a modular system where parts can easily be replaced to show the impact of different techniques and system design decisions. Performance, although very important in a production renderer, is not the primary goal of PVR. The system doesn't use any hard-to-understand or obfuscating optimization techniques, but where algorithmic optimizations are possible, those are discussed and implemented.

The system does use some programming concepts that are advanced in nature, but an effort has been made to keep this in the supporting code. The key parts of the system have all been designed to be as simple and instructive as possible.

### 1.1.1 Use of Namespaces

Namespaces are an important part of writing production code. If we consider some common classes one might write for a render, **Ray**, **Vector**, **Curve**, **Polygon**, etc., it is not unlikely that one of the other libraries that our system uses might have similarly named classes. If this occurs, the compiler will throw up its hands in confusion and tell you that you already created a **Vector** class, and it is not impressed that you are trying to do so again.

C++ provides a concept called *namespaces* that addresses this problem by changing the internal (C++ calls this *mangled*) name of a symbol to include the name of the namespace. This way, the compiler will know the difference between **Imath::Vec3<T>** and your **MyLibrary::Vec3<T>**. For this reason, all of the C++ source code is wrapped in a **pvr** namespace.

PVR also uses namespaces to divide its classes into broad modules, although more for organizational purposes than to prevent symbol name clashes. Representations of geometry, such as polygons, particles, and attributes, are all in the **pvr::Geo** namespace. Modeling-related classes are in **pvr::Model** and so on.

<b>pvr::Accel</b>	Acceleration data structures
<b>pvr::Geo</b>	Geometry-related classes and functions
<b>pvr::Math</b>	Math-related functions
<b>pvr::Model</b>	Volume-modeling classes

<b>pvr::Noise</b>	Noise and fractals
<b>pvr::Render</b>	Rendering-related classes
<b>pvr::Sys</b>	System-related classes
<b>pvr::Util</b>	Various utility functions and classes

### 1.1.2 Use of External Libraries

Where possible, existing libraries have been used to accomplish tasks not central to the functionality of the system and for components not directly related to volume modeling and rendering.

<b>OpenEXR/Imath</b>	Used as the main math library. Contains basic classes such as vectors and matrices.
<b>OpenImageIO</b>	A library for image input and output. Gives PVR the ability to read and write a variety of image formats.
<b>Field3D</b>	Provides voxel data structures and routines for storing voxel data on disk.
<b>GPD</b>	A library that enables reading and writing of Houdini's <b>geo</b> and <b>bgeo</b> formats.
<b>boost</b>	After the Standard Template Library, probably the most commonly used library in the world. PVR uses a variety of classes from <b>boost</b> , from smart pointer classes through timing utilities to Boost.Python.

## 1.2 Python Bindings

The use of Python in production pipelines has exploded over the last five to ten years. It is most commonly used for general scripting, but due to the ease of creating bindings for existing C/C++ libraries, it has also replaced the internal scripting languages of several major applications, for example, hscript in Houdini, MEL in Maya, and Tcl in Nuke.

PVR uses Boost.Python for its bindings, which exposes PVR's internal classes and function directly. Boost.Python makes it very simple to build the bindings and supports advanced memory management and object lifetime techniques, such as smart pointers, deep and shallow copying, etc.

As an example, the following code snippet builds a Python class that exposes the **Renderer** class, automatically manages its lifetime using **Renderer::Ptr**, and also replaces Python's default constructor with a custom one.

**Code 1.1.** Creating Python bindings for the **Renderer** class using Boost.Python

```
class_<Renderer, Renderer::Ptr>("Renderer", no_init)
    .def("__init__",      make_constructor(Renderer::create))
    .def("clone",         &Renderer::clone)
    .def("setCamera",     &Renderer::setCamera)
    .def("setRaymarcher", &Renderer::setRaymarcher)
    .def("addVolume",     &Renderer::addVolume)
    .def("addLight",      &Renderer::addLight)
    .def("execute",       &Renderer::execute)
    .def("saveImage",     &Renderer::saveImage)
    ...
;
```

Boost.Python also makes it easy to expand the flexibility of the C++ code when creating bindings. For example, Python has no concept of **const** objects. To make the C++ code play nice with the Python bindings, Boost.Python provides a function that instructs the binding layer that some extra type conversions are legal, for example converting a non-**const** object to a **const** one.

**Code 1.2.** Instructing Boost.Python that a non-**const** pointer is convertible to a **const** pointer

```
implicitly_convertible<Renderer::Ptr, Renderer::CPtr>();
```

Boost.Python also makes it simple to adapt C++'s rigid type structure to work within Python's loosely typed classes. For example, we can tell it to allow conversions from Python's **list** and **dict** types to C++'s **std::vector** and **std::map**.

**Code 1.3.** Converting a Python type to a C++ type

```
template <typename T>
std::vector<T> pyListValues(boost::python::list l)
{
    using namespace boost::python;
    using namespace std;

    vector<T> hits;

    for (boost::python::ssize_t i = 0, end = len(l); i < end; ++i) {
        object o = l.pop();
        extract<T> s(o);
```

```
    if (s.check()) {  
        hits.push_back(s());  
    }  
}  
  
return hits;  
}  
  
// Convert a python list, grabbing only float values  
boost::python::list l = some_list();  
std::vector<float> floatVec = pyListValues<float>(l);
```

---

The bindings for PVR reside in the `libpvr/python` folder.

## 1.3 Rendering with PVR

PVR's Modeling Pipeline, 91  
Voxel Buffers, 33  
PVR's Rendering Pipeline, 203

To give an example of how rendering is accomplished with PVR, we first look at an example of a Python script that uses the **Modeler** class to build a voxel buffer and then renders it using the **Renderer** class. In the example, references have been made to each of the relevant chapters and sections. The final image is shown in Figure 1.1.



**Figure 1.1.** The result of our PVR example.

First, the **pvr** module is imported into the script. The **from pvr import \*** syntax is used instead of **import pvr** so that the symbols in the library become visible without always having to prefix them with the library name.

#### Code I.4. A simple modeling and rendering example

```
#!/usr/bin/env python

from pvr import *
```

The next step is to create instances of the classes that are used in the volume-modeling process. The **Modeler** is responsible for taking inputs in the form of **ModelerInput** instances and turning them into voxel buffers. Each modeler input contains the definition of a *volume primitive* (also known as *volumetric primitive*), in this case, a **PyroclasticPoint**, and a **Geometry** instance, which is the underlying geometric representation of the volume primitive. Pyroclastic points are a type of rasterization primitive, and their geometric representation is a **Particles** instance.

#### Code I.5. A simple modeling and rendering example

```
# Modeling classes
modeler = Modeler()
parts = Particles()
geo = Geometry()
prim = Prim.Rast.PyroclasticPoint()
input = ModelerInput()
```

The particles object is configured with a single point, which by default sits at the origin. The particles instance is then hooked up to the geometry container, and the two are added to the modeler input along with the volume primitive. In a normal render, the volumetric primitive would be configured with a set of *attributes* to drive its appearance. In this simple example the default parameters will be used.

#### Code I.6. A simple modeling and rendering example

```
# Create a modeling primitive with a single input point
parts.add(1)
geo.setParticles(parts)
input.setGeometry(geo)
input.setVolumePrimitive(prim)
```

The final step in the volume-modeling process is to add the **ModelerInput** to the **Modeler** and then update the *bounds*, which configure the

Fundamentals of Volume Modeling, 75

**Modeler**, 92

**ModelerInput**, 97

**Rast::PyroclasticPoint**, 122

**Geometry**, 20

Rasterization Primitives, 79

Using Geometry to Guide Volumetric Primitives, 82

**Particles**, 21

Geometry Attributes, 24

**ModelerInput**, 97

**Modeler**, 92

Defining the Voxel Buffer, 76

voxel buffer so that it encloses all of the primitives in the modeler's list of inputs. Once the final resolution of the voxel buffer has been set, the modeler is executed, which runs each of its inputs and lets them write their data to the voxel buffer.

**Code 1.7.** A simple modeling and rendering example

```
# Add input to modeler and rasterize
modeler.addInput(input)
modeler.updateBounds()
modeler.setResolution(200)
modeler.execute()
```

**Renderer**, 204

Volumetric Lighting, 167

**PerspectiveCamera**, 220

Cameras, 18

Raymarching in PVR, 267

**Raymarcher**, 268

**RaymarchSampler**, 285

**Light**, 292

**Occluder**, 311

**VoxelVolume**, 243

Next, we create the rendering-related objects. The **Renderer** is the most important one, responsible for firing the rays that are turned into the pixel values of the final image. Any render requires that a camera be present; in this example, a **PerspectiveCamera** is used.

PVR uses raymarching to integrate the volumetric properties in the scene, but the task is broken into two separate parts. The **Raymarcher** takes steps along each ray fired from the final image's pixels, but the radiance and transmittance change at each step is determined by a **RaymarchSampler**.

Various types of light sources may be used in PVR to produce illumination of the scene. The task is divided between two classes. Each **Light** determines the illumination intensity and direction present at various points in the scene, but answering queries about how much light actually arrives at any given point is done by the **Occluder** rather than the light itself.

Although the modeler produces a voxel buffer after the **execute()** method has been called, it is not in a form that the renderer can handle. Instead, we first wrap it in a **VoxelVolume** instance, which makes the buffer renderable.

**Code 1.8.** A simple modeling and rendering example

```
# Rendering classes
renderer      = Renderer()
camera        = PerspectiveCamera()
raymarchSampler = PhysicalSampler()
raymarcher    = UniformRaymarcher()
occluderCamera = SphericalCamera()
occluder      = OtfTransmittanceMapOccluder(renderer, occluderCamera, 8)
light         = PointLight()
volume        = VoxelVolume()

# Configure rendering objects
camera.setPosition(V3f(0.0, 0.25, 6.0))
camera.setResolution(V2i(320, 240))
```



```

raymarcher.setRaymarchSampler(raymarchSampler)
light.setPosition(V3f(10.0, 10.0, 10.0))
light.setIntensity(Color(1.5))
occluderCamera.setPosition(light.position())
light.setOccluder(occluder)
volume.addAttribute("scattering", V3f(4.0, 6.0, 8.0))
volume.setBuffer(modeler.buffer())

```

Before the final image can be rendered, the camera, raymarcher, volume, and light all need to be added to the **Renderer**. Taken together, the **Volume** and **Light** instances make up the **Scene** of the render.

The last step before executing the render is to print the scene information. The `execute()` call will then start firing rays through the scene, which ultimately results in an image that can be saved to disk.

**Renderer**, 204

**Volume**, 233

**Light**, 292

**Scene**, 203

#### Code I.9. A simple modeling and rendering example

```

# Connect the renderer with the raymarcher, camera, volume and light
renderer.setRaymarcher(raymarcher)
renderer.setCamera(camera)
renderer.addVolume(volume)
renderer.addLight(light)

# Print scene structure and start render
renderer.printSceneInfo()
renderer.execute()

# Save image to disk
renderer.saveImage("out/image.png")

```

PVR reports progress and status information at each step in the pipeline. The example code above generates the following log output, which gives feedback on each of the steps in the process.

#### Code I.10. The log file generated by the above rendering example

```


23:01:33 [pvr] Updated bounds to: ((-3 -3 -3), (3 3 3))
23:01:33 [pvr] Creating dense buffer
23:01:33 [pvr] Using uniform/matrix mapping
23:01:33 [pvr] Setting voxel buffer resolution to: (200 200 200)
23:01:33 [pvr] Pyroclastic point primitive processing 1 input points
23:01:35 [pvr]   Rasterization: 22.68%
23:01:38 [pvr]   Rasterization: 45.09%
23:01:40 [pvr]   Rasterization: 67.70%
23:01:43 [pvr]   Rasterization: 90.19%
23:01:44 [pvr]   Time elapsed: 11.2419996
23:01:44 [pvr]   Voxel buffer memory use: 91MB
23:01:44 [pvr]   Scene info:
23:01:44 [pvr]     (VoxelVolume)
23:01:44 [pvr]       a scattering

```

```
23:01:44 [pvr]      i scattering : (4 6 8)
23:01:44 [pvr]      i Empty space optimization disabled
23:01:44 [pvr]      p Isotropic
23:01:44 [pvr]      (PointLight)
23:01:44 [pvr]      i (1.5 1.5 1.5)
23:01:44 [pvr]      o OtfTransmittanceMapOccluder
23:01:44 [pvr] Rendering image (320 240) (1 x 1)
23:01:46 [pvr]      44.23%
23:01:49 [pvr]      88.77%
23:01:49 [pvr]      Time elapsed: 5.51499987
23:01:49 [pvr] Writing image: out/image.jpg
23:01:49 [pvr]      Done.
```


---

This page intentionally left blank



# Fundamentals

This page intentionally left blank



# The Basics

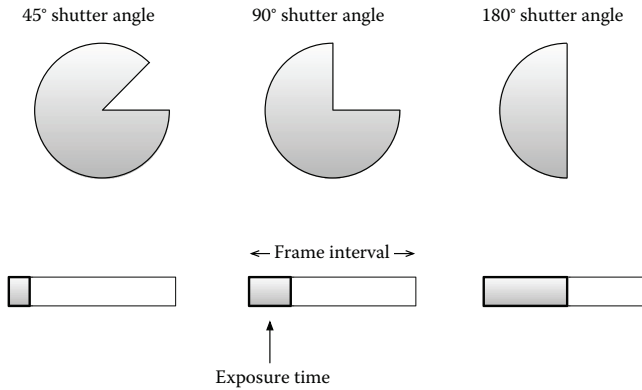
## 2.1 Time and Motion Blur

One of the key requirements of a production renderer is that it must handle the effects of time-varying properties robustly. Because a camera shutter (see Figure 2.1) stays open over a finite amount of time, we must account for changes in the scene during that time period. Attributes that often change with time are object position (translation motion blur), object shape (deformation motion blur), and camera position and orientation (camera motion blur). For a rendering solution to be useful in a production context, each of these must be handled.

Time itself can have multiple different reference frames, each of which can be useful in a renderer. Film and television rely on showing a sequence of images that change quickly enough to give the illusion of motion. Each of these images is referred to as a *frame*, and an integer number can be assigned to each frame, making for a convenient reference frame for time.

Time can, of course, also be measured in seconds. If we consider an animation containing 48 frames, running at 24 frames per second, we have two seconds of material, given that the *frame rate* is the standard 24 frames per second (often abbreviated fps). If we start numbering frames at 1, we can convert between time ( $t$ ) and frame number ( $F$ ) using a simple formula:

$$t = \frac{F - 1}{\text{frames per second}}.$$



**Figure 2.1.** A motion picture camera shutter uses a rotating disk as its shutter. By varying the size of the open section, a certain shutter angle and exposure time is achieved.

A third way of looking at time is to define it in the time interval bounded by the camera shutter opening and closing. Here we can consider a parametric measure of time that is unitless and is defined only within a  $[0, 1]$  interval.

The shutter is usually not open for the entire duration of a frame. Most motion pictures are filmed using a *shutter angle* of  $180^\circ$  [Borum 07], a notion that warrants some explaining. A motion picture camera's shutter is most often a rotating disk with an adjustable slot, which in turn controls how long the exposure of each film frame is. The degree measure refers to the angle the slot occupies on the disk, meaning that at  $180^\circ$  the shutter is open 50% of the time and is closed 50% of the time. In computer graphics, this is more often expressed as a fraction called *shutter length* or *motion blur length*, and the corresponding value to  $180^\circ$  is 0.5.

We often refer to the actual time that the camera shutter is open as  $dt$ , and we can find it using the frame rate and shutter angle or motion blur length:

$$dt = \frac{\text{motion blur length}}{\text{frames per second}} = \frac{\text{shutter angle}/360^\circ}{\text{frames per second}}.$$

A strict and consistent definition and handling of the various time frames is important in a renderer, and in PVR, they are actual class types, which prevents accidental misinterpretation. Time in seconds is defined by the **Time** class, whereas shutter open/close time is defined by **PTime**. PVR does not deal with frame time directly.

**Code 2.1.** The `Time` class

```

class Time
{
public:
    explicit Time(const float t)
        : m_value(t)
    { }
    operator float() const
    { return m_value; }
    float value() const
    { return m_value; }
private:
    float m_value;
};

```

**Code 2.2.** The `PTime` class

```

class PTime
{
public:
    explicit PTime(const float t)
        : m_value(t)
    { }
    operator float() const
    { return m_value; }
    float value() const
    { return m_value; }
private:
    float m_value;
};

```

Geometry Attributes, 24

In PVR, the motion of particles and geometry is represented using the  $v$  (for *velocity*) point attribute (see Section 2.4). Velocity vectors are always defined in m/s, meters per second. *Motion vectors* are different from velocity vectors; they refer to motion in the current shutter open/close interval and thus only have the length unit  $m$ . Any velocity vector can be converted to a motion vector through the formula

$$\vec{m} = \vec{v} \cdot dt.$$

### 2.1.1 Render Globals

When designing a renderer, there is always a certain amount of information that pertains to all parts of the renderer. Rather than passing a set of variables to every single rendering function, it is quite common to use an



object that is globally accessible to store such information. In PVR, this class is called **RenderGlobals**.

The **RenderGlobals** class stores pointers to the current scene as well as to the render camera. Information relating to time is also available in the class. When a new render frame is initiated, the rendering class configures **RenderGlobals** by calling the static **setupMotionBlur()** method. After that, any of PVR's components can access *dt* through the **RenderGlobals::dt()** method.


### Code 2.3. RenderGlobals

```
class RenderGlobals
{
public:
    // Typedefs
    typedef boost::shared_ptr<const pvr::Render::Scene> SceneCPtr;
    typedef boost::shared_ptr<const pvr::Render::Camera> CameraCPtr;
    // Exceptions
    DECLARE_PVR_RT_EXC(BadFpsException, "Bad frames per second value:");
    DECLARE_PVR_RT_EXC(BadShutterException, "Bad shutter value:");
    // Main methods
    static void      setupMotionBlur(const float fps, const float shutter);
    static void      setScene(SceneCPtr scene);
    static void      setCamera(CameraCPtr camera);
    // Accessors
    static float      fps();
    static float      shutter();
    static float      dt();
    static SceneCPtr  scene();
    static CameraCPtr camera();
private:
    // Data members
    static float      ms_fps;
    static float      ms_shutter;
    static float      ms_dt;
    static SceneCPtr  ms_scene;
    static CameraCPtr ms_camera;
};
```

## 2.1.2 Function Curves

So far, we have mentioned motion blur, but we have not yet showed how time-varying properties can be represented. PVR treats properties that change with time as one-dimensional functions that are linearly *interpolated* to find sub-sample values. (See, for example, Figure 2.2.)

In most cases, time-varying properties change gradually and smoothly over time and can be described well using only two time samples: one at the shutter open time and one at shutter close. In other cases, such as




**Figure 2.2.** A function curve using four samples and linear interpolation.

a camera with “shake” applied, the changes can be drastic and sudden, requiring more than two samples to accurately describe the change in position, orientation, etc. In order to support arbitrary intra-frame motion, PVR handles these time-varying properties using the `Util::Curve` class. Each `Curve` can contain as many samples as is necessary, and in-between values are then linearly interpolated.

#### Code 2.4. Curve

```
class Curve
{
public:
...
    // Main methods
    void addSample(const float t, const T &value);
    T interpolate(const float t) const;
    size_t numSamples() const;
    const SampleVec& samples() const;
    std::vector<float> samplePoints() const;
    std::vector<T> sampleValues() const;
    void removeDuplicates();
    static CPtr average(const std::vector<CPtr> &curves);
private:
...};
```



**Figure 2.3.** PVR's temporal coordinate frames.

The **Curve** class is templated and can theoretically be used to store any type of data that can be interpolated. The most commonly used types, however, are

```
typedef Curve<float>   FloatCurve;
typedef Curve<Color>   ColorCurve;
typedef Curve<Vector>  VectorCurve;
typedef Curve<Quat>    QuatCurve;
```

When a **Curve** is used in PVR, the time dimension is assumed to line up with **PTime**, such that  $t = 0.0$  refers to the start of the frame and the shutter open time. The end of the shutter interval falls at  $t = 1.0$ , which matches the **PTime** definition. When constructing a curve, it is most common to have information about how a given attribute changes from frame to frame. If we let  $t_p^i = 0.0$  be the **PTime** of the current frame start, then the equivalent **PTime** for the next frame start can be found through the expression

$$t_p^{i+1} = \frac{1}{\text{motion blur length}}.$$

**PTime**, 15

Figure 2.3 illustrates PVR's definition of time reference frames for a motion blur length of 0.5.

## 2.2 Cameras

The camera classes in PVR are very simple, and their implementations should provide no surprises to the reader. Most parameters are assumed

to be temporally varying, meaning that properties such as position, orientation, and field of view may change over the course of the current frame. In support of this, all of the calls to the camera that depend on time use the **PTime** concept, where 0.0 is assumed to be the start of the current frame, i.e., when the shutter opens, and 1.0 is the time that the shutter closes.


This book assumes the reader is familiar with common computer graphics conventions for camera projection calculations. For a thorough introduction to cameras in computer graphics, [Watt 00] and [Pharr and Humphreys 10] are good sources.

### 2.2.1 Camera Coordinate Spaces

PVR's world space is right-handed, but its camera space is left-handed. By default, a camera with no rotation looks down the negative  $z$ -axis in world space, and down positive  $z$  in camera space. The camera's  $x$ - and  $y$ -axes line up with the world space, with  $x$  to the right and  $y$  facing up. (See Figure 2.4.)

Screen space defines the projected view of the camera, with  $x = -1$  indicating the left edge of the camera's view,  $x = 1$  the right edge,  $y = -1$  the bottom edge, and  $y = 1$  the top edge. The depth dimension has no negative range, instead putting  $z = 0$  at the near plane of the projection and  $z = 1$  at the far plane. (See Figure 2.5.)

PVR also uses a second projection space, which is called NDC space. NDC stands for normalized device coordinates and changes the  $x$  and  $y$  ranges of the projection to the  $[0, 1]$  range. (See Figure 2.6.)



**Figure 2.4.** PVR's camera space conventions.