


Preface

This book is an extensive and practical collection of articles about real-time computer graphics, accumulating the knowledge and experience of experts in both industry and academia. Building, in the same style, upon the wealth of the great “Gems” books already available, *GPU Gems* is a collection of short chapters. However, a number of key characteristics make this book unique and valuable to today’s developers as they attempt to harness the ever-increasing power of the graphics processing unit (GPU).

First and foremost, this book focuses squarely on real-time programmable graphics—specifically, on techniques relevant to GPUs. Each chapter was carefully selected to present ideas and techniques that are directly useful in interactive applications, such as computer games. The chapters provide insight and understanding, rather than focusing on low-level API calls or specific mathematical tricks. Furthermore, each chapter is packed with numerous full-color diagrams and images to illustrate and drive home key concepts. Finally, the experience and diversity of the contributors help readers gain a broad understanding, as well as a certain confidence that the advice they are getting comes from experts in the field.

NVIDIA’s strongest asset is its people: the depth and quality of their collective expertise inspired the initial idea for *GPU Gems*. With so much knowledge and expertise at hand, we felt that the thoughts and insights of the teams that brought us many recent advances in real-time graphics would make for a wonderfully instructive book. So, we started the project with an internal call for participation.

Having the good fortune to work with people from leading game development houses, tool developers, film studios, and academic institutions who are shaping the future of real-time computer graphics, we also wanted to highlight their real-world contributions in *GPU Gems*. Hence, a wider, public call for participation allowed us to coalesce a great amount of talent and refreshing perspective into this volume.



Whether you're creating new effects, architecting a graphics engine, or squeezing out the last bits of performance, we hope that this book provides valuable guidance and saves you from some of the challenges the authors faced on their own projects. All of us who worked on *GPU Gems* hope that it will help you to adopt new ideas and take your projects to the next level of graphical realism.

Our Intended Audience


This book provides intermediate and advanced readers with useful information that will help them in their projects. Focusing beyond the fundamentals of high-level shading, *GPU Gems* looks at how to take existing projects further by removing the mystery behind complex effects and advanced GPU programming. With the rapid evolution of real-time shading languages, the collection of algorithms available to real-time graphics developers is larger than ever. By compiling and distributing the information in this book, our goal is to make high-quality, high-performance graphics more accessible to a wider audience that includes game developers, technical directors, professors, and students.

Trying the Examples

Many of the chapters in this book include code samples to make their subject matter more concrete. The authors used whichever shading language they wanted, so the code samples ended up in DirectX 9's High-Level Shader Language (HLSL) or Cg, which were the only two high-level shading languages widely in use during this project. Almost everything that is presented can be applied to either language, as well as to languages that came later, such as the OpenGL Shading Language. The code samples are available on the CD that accompanies this book, along with standalone examples wherever possible. This makes it easy for you to integrate or experiment with the various examples. Updated sample code, as well as additional supplementary materials, is available at the book's Web site: <http://developer.nvidia.com/GPUGems/>.

Acknowledgments

A project the size of *GPU Gems* cannot succeed without the efforts of many people. First, I would like to thank the contributors, without whom this book would not exist. It has been my privilege to work with such an experienced and capable group of people. But part of the challenge when working with a group of this caliber is that everyone is



also exceptionally busy. I am grateful to them for taking the time to work on the project, and for putting in the effort to produce superlative results. Thanks also go to their respective organizations for allowing them to participate in *GPU Gems*. Our appreciation also goes to those who responded to our call for participation and presented worthy proposals for consideration.

The section editors—Kevin Bjorke, Cem Cebenoyan, Sim Dietrich, Simon Green, Juan Guardado, and John Spitzer—contributed enormously by iterating with the authors, helping to shape the manuscript, and tracking down articles to keep them on schedule. I'm grateful to all of them for volunteering to assist with this project. Chris Seitz was also instrumental in the success of this project, taking care of legal issues, reviewing chapters, and always being ready to take a few moments to discuss project concerns with me.

Each chapter underwent an extensive review process, involving comments from peer reviewers, editors, and external reviewers. The contributors and section editors did a wonderful job critiquing and helping to improve the book's content. In particular, I'd like to thank Matt Pharr (who spent weekends above and beyond the call of duty to help review several chapters—in addition to writing three!) and Kevin Bjorke (who was a section editor in addition to writing and contributing to several chapters). Larry Gritz, Eric Haines, and Matthias Wloka were kind enough to serve as critical external reviewers and provided a wealth of insightful comments. Thanks also to our anonymous external reviewers, who did well to tackle the massive amount of material they were given.

Caroline Lie, Spender Yuen, Dana Chan, and Melvin Chong provided their expertise to create the book's cover, template, and diagrams. I would particularly like to thank Spender for his patience as we worked through the more than 100 diagrams in the book (in addition to over 200 screenshots). Catherine Kilkenny, Debra Valentine, and Teresa Saffaie helped improve and clarify the writing. David Kirk lent his insight to produce the book's foreword.

Christopher Keane did a fantastic job of pulling the manuscript through copyediting and composition. Many thanks go to Jacquelyn Doucette, John Fuller, Bernard Gaffney, Curt Johnson, Heather Mullane, and the other folks at Addison-Wesley for their help in this project from start to finish. In particular, Peter Gordon's encouragement and deadlines helped us keep the book on schedule.

From architecting GPUs to helping developers create exciting new content, there are innumerable steps that take place at NVIDIA before the knowledge contained in this



book can be useful to the graphics community. Therefore, I'd like to thank everyone at NVIDIA, because without their hard work there wouldn't be an opportunity to do a project like *GPU Gems*. Specific thanks to Mark Daly, Dan Vivoli, and Jen-Hsun Huang for providing the teams and resources to make it all happen.

Finally, I would like to thank my parents and sister for their extraordinary support, and for making it possible for me to have such exceptional opportunities.

Given the unique combination of effort, creativity, and care that so many talented people contributed to *GPU Gems*, I'm sure this book will serve you well.

Randima (Randy) Fernando
NVIDIA Corporation

