

High-Performance Graphics 2023

Delft, The Netherlands

June 26 — 28, 2023

General Chairs

Quirin Meyer, Hochschule Coburg
David McAllister, AMD

Papers Chairs

Jacco Bikker, Traverse Research
Christiaan Gribble, AMD

Preface

We are pleased to present the papers program for High-Performance Graphics (HPG) 2023, the leading international forum for performance-oriented graphics and imaging systems research. The HPG conference series brings together researchers, engineers, and architects to discuss the complex interactions between hardware and software in the design of systems for current and future graphics and visual computing applications.

Formed by the merger of the Eurographics Workshop on Graphics Hardware and the IEEE Symposium on Interactive Ray Tracing in 2009, this year marks the fifteenth in which industrial and academic researchers, practitioners, and students have gathered to advance the field of high-performance computer graphics through HPG's unique platform.

The HPG conference series publishes top-tier research in high-performance graphics, and this year is no exception. We received a total of 28 paper submissions, each of which was reviewed by at least four members of our International Program Committee (IPC). After a short but thorough review process, the IPC recommended 17 innovative papers for this year's conference. The accepted papers went through a second review cycle to ensure the highest quality. Ultimately, seven papers received invitations to appear in a special issue of Eurographics Computer Graphics Forum, and the remaining 10 papers were published in the ACM/Eurographics HPG 2023 Conference Proceedings.

The 2023 paper program covers a wide range of topics including acceleration structures; primitives, surfaces, and appearance modeling; distributed and cloud-based rendering; and GPU computing. Our call for participation particularly encouraged submissions in emerging areas of visual computing—for example, high-performance computer vision, machine learning, and compiler and language technologies—and this year's authors responded in force: Three papers focusing on neural networks for intersection, shading, and anti-aliasing were accepted for publication, and these papers were featured in a dedicated Deep Learning for Graphics session at the conference.

We offer our sincere thanks to the authors of all submitted papers—truly these authors are the lifeblood of HPG's successful papers program. We also offer our deep appreciation for this year's IPC members, each of whom took the time to write detailed and thoughtful reviews and to help guide our efforts in the decision-making process.

Together, the contributions of our authors and committee members here culminate in an exciting selection of papers for High-Performance Graphics 2023.

Sincerely,
Jacco Bikker & Christiaan Gribble
HPG 2023 Papers Chairs

Sponsors





Table of Contents

Acceleration Structures

- Edge-Friend: Fast and Deterministic Catmull-Clark Subdivision Surfaces* e14863
Bastian Kuth, Max Oberberger, Matthäus Chajdas, and Quirin Meyer

Primitives, Surfaces, and Appearance Modeling

- Real-Time Rendering of Glinty Appearances using Distributed Binomial Laws on Anisotropic Grids* e14866
Thomas Deliot and Laurent Belcour
- Sampling Visible GGX Normals with Spherical Caps* e14867
Jonathan Dupuy and Anis Benyoub
- Real-Time Ray Tracing of Micro-Poly Geometry with Hierarchical Level of Detail* e14868
Carsten Benthin and Christoph Peters

Deep Learning for Graphics

- Generative Adversarial Shaders for Real-Time Realism Enhancement* e14870
Arturo Salmi, Szabolcs Cséfalvay, and James Imber

Distributed and Cloud-Based Rendering

- Data Parallel Multi-GPU Path Tracing using Ray Queue Cycling* e14873
Ingo Wald, Milan Jaroš, and Stefan Zellmann

GPU Computing

- GPU-Accelerated LOD Generation for Point Clouds* e14877
Markus Schütz, Bernhard Kerbl, Philip Klaus, and Michael Wimmer

International Program Committee

Afra, Attila – Intel
Apolinário, Antônio – Federal University of Bahia
Assarsson, Ulf – Chalmers University of Technology
Benthin, Carsten – Intel Corporation
Bikker, Jacco – Utrecht University
Binder, Nikolaus – NVIDIA
Bitterli, Benedikt – NVIDIA
Bittner, Jiří – Czech Technical University in Prague
Boubekeur, Tamy – Telecom Paris
Chajdas, Matthäus – Advanced Micro Devices
Clarberg, Petrik – NVIDIA Corporation
Clua, Esteban – Universidade Federal Fluminense
Dachsbacher, Carsten – Karlsruhe Institute of Technology
Dietrich, Andreas – ESI Group
Doggett, Michael – Lund University
Dupuy, Jonathan – Intel Corporation
Eisemann, Elmar – Delft University of Technology
Hanika, Johannes – Karlsruhe Institute of Technology
Harada, Takahiro – Advanced Micro Devices, Inc.
Havran, Vlastimil – Czech Technical University in Prague
Heitz, Eric – Intel
Herholz, Sebastian – Intel Corporation
Hunt, Warren – Oculus Research
Jarabo, Adrián – Meta Reality Labs Research
Kaplanyan, Anton S. – Facebook Reality Labs
Knoll, Aaron – Advanced Micro Devices
Krüger, Jens – HPC Group, University Duisburg-Essen
Lanman, Douglas – Oculus Research
Lee, Sungkil – Sungkyunkwan University
Lee, Won-Jong – Intel
Leimkuehler, Thomas – MPI Informatik
Luebke, David – NVIDIA
Mantiuk, Rafal – University of Cambridge
Manzke, Michael – Trinity College Dublin
McGuire, Morgan – Roblox
Meister, Daniel – Advanced Micro Devices, Inc.
Membarth, Richard – THI, DFKI
Meng, Xiaoxu – Tencent
Meyer, Quirin – Coburg University

International Program Committee

Mueller, Joerg – Graz University of Technology
Mueller, Thomas – NVIDIA
Mullia, Krishna – Adobe Research
Neyret, Fabrice –
Nimier-David, Merlin – NVIDIA
Oberberger, Max – Advanced Micro Devices
Olano, Marc – UMBC
Ortin, Marta – Universidad de Zaragoza
Owens, John – UC Davis
Patney, Anjul – NVIDIA
Paulin, Mathias – IRIT, Université de Toulouse, CNRS, UPS, France
Pharr, Matt – NVIDIA
Ramamoorthi, Ravi – University of California, San Diego
Reshetov, Alexander – NVIDIA
Ritschel, Tobias – University College London
Salvi, Marco – NVIDIA
Schütz, Markus – TU Wien
Selgrad, Kai – OTH Regensburg
Silvennoinen, Ari – Activision Publishing, Inc.
Sintorn, Erik – Chalmers University
Slusallek, Philipp – DFKI, Saarland University
Spjut, Josef – NVIDIA
Stamminger, Marc – Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
Steinberg, Shlomi – University of California Santa Barbara
Steinberger, Markus – Graz University of Technology, Huawei Technologies
Vardis, Konstantinos – Athens University of Economics and Business
Viitanen, Timo – NVIDIA
Wald, Ingo – NVIDIA
Wang, Beibei – Nanjing University of Science and Technology
Wang, Lili – Beihang University
Wang, Rui – Zhejiang University
Wei, Li-Yi – Adobe Research
Wyman, Chris – NVIDIA
Yan, Ling-Qi – UC Santa Barbara
Yoon, Sungeui – KAIST
Zeltner, Tizian – NVIDIA
Zhao, Shuang – University of California, Irvine
Zhou, Kun – Zhejiang University

Author Index

Belcour, Laurent	e14866	Klaus, Philip	e14877
Benthin, Carsten	e14868	Kuth, Bastian	e14863
Benyoub, Anis	e14867	Meyer, Quirin	e14863
Chajdas, Matthäus	e14863	Oberberger, Max	e14863
Cséfalvay, Szabolcs	e14870	Peters, Christoph	e14868
Deliot, Thomas	e14866	Salmi, Arturo	e14870
Dupuy, Jonathan	e14867	Schütz, Markus	e14877
Imber, James	e14870	Wald, Ingo	e14873
Jaroš, Milan	e14873	Wimmer, Michael	e14877
Kerbl, Bernhard	e14877	Zellmann, Stefan	e14873

Women in Computer Graphics

Zerrin Yumak – Utrecht University

Dr. Zerrin Yumak is an assistant professor in the Human-Centered Computing Group at the Information and Computing Sciences Department of Utrecht University. She is the scientific director of the Motion Capture and Virtual Reality Lab and founding member of the EDI Committee of her department. She has 15 years of experience in the area of socially interactive 3D digital humans. Her latest work focus on Generative AI methods for motion synthesis in particular non-verbal behaviors such as facial expressions, gestures and gaze.

Noshaba Cheema – Max-Planck-Institut für Informatik

Noshaba Cheema is PhD Candidate at Max-Planck-Institut für Informatik (IMPRS) and researcher at DFKI (ASR). She is dealing with the animation of 3D characters using AI, for computer games, films and so that they can end up interacting with humans.

Seyran Khademi – Delft University

Since April 2021, Seyran Khademi is an Assistant Professor at the Faculty of Architecture and the Built Environment (ABE) and the co-director of the AiDAPT lab (AI for Design, Analysis, and Optimization in Architecture and the Built Environment). She is working as an interdisciplinary researcher between the Computer Vision lab and Architecture Department at ABE. Her research interest lies at the intersection of Data, Computer Vision, and Deep Learning in the context of man-made imagery including illustrations and visual data for Architectural Design. She is particularly interested in the role of AI in the creative process of design. In 2020, she was honored to be the research in residence fellow at the Royal Library of the Netherlands working on visual recognition for children's book collection. In 2017 she was appointed as a postdoctoral researcher at the Computer vision lab working on the ArchiMediaL project, regarding the automatic detection of buildings and architectural elements in visual data focusing on Computer Vision and Deep Learning methods for archival data and street-view imagery. Seyran received her Ph.D. in signal processing and optimization in 2015 from TU Delft, followed by postdoctoral research on Intelligent Audio and Speech algorithms.

Course

A Gentle Introduction to Vulkan for Rendering and Compute Workloads

Lukas Lipp – TU Wien

Bernhard Kerbl – Université Côte d’Azure

Abstract

Join us and take your first steps with the Vulkan API, a modern and open graphics API from the Khronos Group. This Tutorial will cover the very basics but also more complex applications in the fields of real-time-rendering and GPGPU computing. No prior experience is required, though having some basic understanding will make the experience more fun. This tutorial will be presented in a lecture-only format and requires no active coding by the participants. As optional homework, participants are invited to also try out the educational Vulkan framework which was used during the Vulkanised 2023 tutorial to test their newly-acquired knowledge, as well as a separate suite of GPU compute coding tasks.

Short Biographies

Lukas Lipp is a third-year Ph.D. Candidate at TU Wien. His area of research centers around both real-time and offline rendering, with a particular emphasis on differentiable rendering. Besides working on his own research framework which is based on Vulkan, he contributed to the development of a Vulkan programming framework used in an introductory graphics course at TU Wien as well as for a hands-on tutorial at the Vulkanised 2023 conference. He also gave lectures about Vulkan at university and at Vulkanised.

Bernhard Kerbl is a postdoc at Inria, Université Côte d’Azure in the GraphDeco group. He obtained his PhD from the Graz University of Technology, followed by a postdoc position at TU Wien. His past and current research focuses on point-based rendering, high-performance computing, differentiable rendering and novel-view synthesis. Bernhard has taught GPU programming lectures at several Austrian universities.

Keynote

High Performance Graphics at Guerrilla, an Art Director's Perspective

Jan-Bart van Beek
Guerrilla

Short Biography

As Studio Art Director, Jan-Bart van Beek oversees the art and animation quality of all Guerrilla projects at both studios. A photography student at the Royal Academy of Fine Arts in The Hague, Jan-Bart turned to computer graphics software to enhance his work. After graduation he applied his knowledge as a CG artist in the field of advertising. His ambition to create something more enduring led him to join Guerrilla, where he quickly became the Lead Artist on Killzone and the Art Director on Killzone 2 and beyond, before taking on the roles of Studio Art Director on the Horizon franchise and Studio Director at Guerrilla.

Keynote

Pursuing High-Performance in Photorealistic Rendering

Lingqi Yan

UC Santa Barbara

Abstract

In today's world, the aspiration to achieve photorealistic computer-generated imagery has become increasingly prominent. However, the need for realism places additional burdens on the requirements of high-performance graphics, involving detailed materials, complex geometry, and costly light transport. This keynote presentation aims to explore the challenges encountered in the quest for realism and share preliminary solutions based on recent research from our group. Both offline rendering and real-time rendering will be examined, highlighting the ways in which traditional mathematical/physical approaches and modern neural networks can contribute to both domains.

Short Biography

Lingqi Yan is an Assistant Professor of Computer Science at UC Santa Barbara, a co-director of the MIRAGE Lab, and affiliated faculty in the Four Eyes Lab. Lingqi strives to achieve photo-realistic rendering. His research focuses on physically-based rendering in both micro and macro ends, encompassing areas such as appearance modeling and representation, physical light transport theory, and real-time ray tracing practices. Lingqi's outstanding contributions have been recognized through various accolades, including the SIGGRAPH 2019 Outstanding Doctoral Dissertation Award and a SIGGRAPH 2022 Best Paper Honorable Mention.

Hot3D

Ray Tracing with Imagination: From Caustic and Wizard to Photon

Simon Fenney

Imagination Technologies

Abstract

Imagination has been developing and producing ray tracing hardware for around 12 years. In this talk we will look back at some of the early work that started at Caustic Graphics, e.g. the Caustic 2500 accelerator card of 2013, moving through “Wizard/Plato” that was demonstrated circa 2015, and on to “Photon”, the latest technology we offer to our silicon partners. We will look at some of the principles that have remained in common but also examine a few of the differences that have evolved over the years.

Short Biography

Simon Fenney is a Research Fellow at Imagination Technologies, having been a member of the PowerVR division since its inception in 1992. In the mid-1980s he studied Computer Science at the University of Queensland, specialising in graphics and parallel computing, before taking a role at Integrated Arts, developing a ray tracer for transputer arrays and, later, working on a 2D vector graphics animation system. His interests include image and texture compression, rasterisation, ray tracing, and the odd spot of gardening.

Keynote

Multum In Parvo: Level of Detail and Approximation Models at the Graphics Nexus

Tamy Boubekeur
Adobe Research

Abstract

MIP maps, level-of-details, approximation models, multiresolution processing and multiscale analysis play a central role in 3D computer graphics. In this talk, I will delve into some aspects of our recent contributions to the 3D digital content creation pipeline. From material capture to real-time ray tracing, through surface reconstruction, I will emphasize how these methodologies not only allow for faster and more scalable 3D graphics, but indeed play a major role in the way we design algorithms in 3D and, ultimately, are a fundamental aspect of what we do in our field and how we can impact other areas of applied sciences.

Short Biography

Tamy Boubekeur is a researcher in computer science, specialized in 3D computer graphics. He is currently a Lab Director and Senior Principal Research Scientist at Adobe Research, as well as a Professor at the Computer Science Department of Ecole Polytechnique, Institut Polytechnique de Paris. He is also a Professor (on leave) at Telecom Paris, Institut Polytechnique de Paris. He was previously the founder and head of the Computer Graphics Group at Telecom Paris and chief scientist at Allegorithmic. He was also an Associate Researcher at TU Berlin and a PhD student at INRIA. His lab at large works on 3D visual computing. His personal research areas focus on 3D Computer Graphics, with a special interest in Modeling, Rendering and Learning efficiently 3D data.