I3D 2014 Guest Editor's Introduction

Marc Olano, *Member, IEEE* and Miguel Otaduy, *Member, IEEE*

THIS special section of the *IEEE Transactions on Visualization and Computer Graphics* (*TVCG*) brings you two extended papers based on work first presented at the ACM Symposium on Interactive 3D Graphics and Games (I3D) in 2014. I3D focuses on real-time rendering, animation, and interaction techniques. Games are not the only application for these kinds of interactive methods, but are certainly a significant user of them. Consequently, the quality, time, space, stability, and predictability constraints faced by game development are an important factor in much of the work presented at I3D. This year continues the five year tradition of having the best I3D papers extended with more details and new work for publication in *TVCG*, and with several regular *TVCG* papers being presented at I3D.

The 2014 conference, the 18th for I3D, was held from March 14-16 in San Francisco, California. This continues the I3D custom as a spring conference, with a deadline well timed for summer work, and a conference nestled between the big winter deadlines and summer conference season. The 2014 conference had the added bonus of being co-located with the annual Game Developers conference, continuing to provide a needed bridge between published work on interactive graphics and game development practitioners.

I3D 2014 received 47 paper submissions. Each submission received at least four reviews, by members of the international papers committee or external reviewers. All reviewers participated in online discussions about the papers they reviewed, resulting in a consensus or calls for additional reviews, and ultimately decisions to accept 19 high-quality papers. Of these, a selection of the top reviewed papers received invitations to extend their work for *TVCG*. These papers underwent the full journal review process, with multiple iterations of editing and review. You will find the resulting extended versions of those papers here today.

In addition to the papers program, I3D 2014 also included a posters and demos session, games industry panel, a banquet speech by Chris Wyman of NVIDIA, and keynote talks by Jos Stam of Autodesk and Kun Zhou of Zhejiang University. Together this results in an intimate three-day conference packed with content.

- *M. Olano is with the Computer Science and Electrical Engineering Department, University of Maryland, Baltimore County. E-mail: olano@umbc.edu.*
- *M. Otaduy is with the the Department of Computer Science, Universidad Rey Jaun Carlos-Madrid. E-mail: miguel.otaduy@urjc.es.*

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TVCG.2015.2422231

We are excited to present here extended work building on two papers from that conference, "k+-buffer: An efficient, memory-friendly and dynamic k-buffer framework" by Andreas Vasilakis and Ioannis Fudos, and "More Efficient Virtual Shadow Maps for Many Lights" by Ola Olsson, Erik Sintorn, Viktor Kämpe, Markus Billeter, and Ulf Assarsson.

"k+-buffer: An efficient, memory-friendly and dynamic k-buffer framework" by Vasilakis et al. presents an advanced k-buffer algorithm, which also serves as a case study for per-pixel parallel locking and synchronization primitives in recent graphics hardware. The k-buffer is a fixed size variation on traditional A-buffer method. Since the A-buffer can have unbounded size for deep depth complexity, it is not well suited to hardware implementation. The k-buffer keeps just k fragments, ideally the k closest or k best by some measure. An effective mapping of the k-buffer to rendering hardware could support order-independent rendering of transparent objects, hair, fur, grass, leaves, and a myriad of other layered natural phenomena. Unfortunately, finding the k best fragments requires inserting new fragments into a sorted list of the ones seen so far. The authors present a new version of this algorithm that uses hardware supported synchronization operations to manage the write contention for multiple fragments landing in each pixel. The original I3D paper considered only hardware synchronization primitives available on certain Intel graphics accelerators. The authors have now extended this work to synchronization primitives on NVIDIA hardware as well, and provide a detailed performance evaluation. In addition, they have improved several aspects of their algorithm, improving efficiency and helping to estimate the best value of k. In the resulting paper, the authors provide an exciting new algorithm for multi-layer rendering, and also provide data that we expect to see referenced in many future papers on GPU parallelism.

"More Efficient Virtual Shadow Maps for Many Lights" by Olsson et al. addresses the difficult problem of managing shadows for large numbers of lights in complex scenes. The authors' approach can handle tens to hundreds of lights in scenes with hundreds of thousands to millions of triangles. This is achieved through a combination of clustering of lights and geometry and use of the relatively new hardware support for virtual textures for the cube shadow maps. Virtual textures provide mechanisms to only allocate memory for the portions that are used, thus saving memory and shadow rasterization costs.

Both of these papers explore the performance and algorithmic benefits of emerging graphics hardware features for hard interactive rendering problems. We sincerely hope the readers will find them as exciting and enriching as we have.

Marc Olano received the BS degree in electrical engineering from the University of Illinois in 1990 and the PhD degree in computer science from the University of North Carolina in 1998. He is currently an associate professor in the Computer Science and Electrical Engineering Department, and the director in the Computer Science Game Development Track, University of Maryland, Baltimore County. His PhD dissertation described the first interactive graphics hardware shading language. After receiving the PhD degree, he joined SGI, creating shading languages for commercial graphics hardware. Since 2002, he has been at UMBC, where his research continues to revolve around all aspects of design and use of programmable graphics hardware. He has served as a conference co-chair of the ACM Symposium on Interactive Graphics and Games in 2005 and 2013, and papers co-chair in 2006 and 2014. In addition, he served as the papers co-chair for the ACM/Eurographics Symposium on Graphics Hardware in 2006 and the Student Research Competition chair for ACM SIGGRAPH in 2013 and 2014. He is a member of ACM SIGGRAPH, Eurographics, the IGDA, the IEEE, and the IEEE Computer Society.

Miguel Otaduy received the BS degree in electrical engineering from Mondragón University, in 2000 and the MS and PhD degrees in computer science from the University of North Carolina at Chapel Hill, in 2003 and 2004, respectively. He is an associate professor in the Department of Computer Science's Modeling and Virtual Reality Group, Universidad Rey Juan Carlos (URJC Madrid). From 2005 to 2008, he was a research associate at ETH Zurich. His main research interests include physically based computer animation and computational haptics. He has published more than 80 papers in these fields, and he has served on the editorial board for several journals and conferences, most notably *IEEE Transactions on Haptics*, *IEEE Transactions on Visualization and Computer Graphics*, 2013 and 2015 IEEE World Haptics Conference, 2014 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, and 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. He is a member of the IEEE.