

OpenCL Tutorial presenter biographies

Neil Trevett is Vice President of Mobile Content at NVIDIA, and President of the Khronos Group. Neil has spent over twenty five years in the 3D graphics industry and is responsible at NVIDIA for enabling and encouraging compelling applications on mobile devices and Smartbooks. Previously, as Vice President of 3Dlabs, Neil was at the forefront of the silicon revolution bringing interactive 3D to the PC. Neil is currently the elected President of the Khronos Group where he initiated the OpenGL ES and OpenKODE working groups and chairs the OpenCL and EGL working groups that are defining industry standards for advanced compute, graphics and media processing on a wide range of mobile, embedded and desktop systems. Neil was elected President for eight consecutive years of the Web3D Consortium dedicated to creating open standards for communicating real-time 3D on the Internet.

Mike Houston is a Senior System Architect in the Advanced Technology Development group at AMD in Santa Clara, working in architecture design and programming models for parallel architectures. He received his PhD in Computer Science from Stanford University in 2008, focusing on research in programming models, algorithms, and runtime systems for parallel architectures including GPUs, Cell, multi-core, and clusters. His dissertation includes the Sequoia runtime system, a system for programming hierarchical memory machines. He received his BS in Computer Science from the University of California, San Diego in 2001.

Tim Mattson (Ph.D. theoretical Chemistry, 1985) is a principal engineer at Intel Corporation. He is a parallel programmer. He started with the Caltech/JPL cosmic cube in 1985 and over the years worked on every major class of parallel system including MPP, VLIW, SIMD, clusters, SMP, NUMA and vector supercomputers. Noteworthy projects include the world's first TFLOP computer, OpenMP, the first generally programmable TFLOP chip (Intel's 80 core research chip), OpenCL, and pioneering work on design patterns for parallel programming. Tim is also a professional kayak instructor (advanced open ocean) and a kayak instructor trainer.

Chris Lamb manages GPU computing system software at NVIDIA, responsible for GPU computing initiatives including OpenCL. A graduate of the Computer Engineering program at the University of Illinois at Urbana-Champaign, he is also the co-founder of a small, but successful, microwave wireless communications company. Most recently prior to joining NVIDIA he was an architect at the many-core physics startup Ageia working on software/hardware co-design, hybrid task and data-parallel algorithms, programming models, and compiler technology.

Eric Schenk is CTO of EA's central technology group EATech, where he steers the strategic direction of EA's central technology. Eric holds a Ph.D. in computer science from the University of Toronto, where he focused on the theory of parallel and distributed computing. After completing his Ph.D. he spent time teaching and doing research at Lund University in Sweden, before leaving academia for the computer games industry in 1997. He is a past co-maintainer of the Linux kernel networking code, and started in the games industry doing research and implementations of network protocols for games. As part of his role in R&D in the games industry, he was an Adjunct Professor at Simon Fraser University from 1998 to 2001. Since

joining Electronic Arts in late 1998, he has done work in several areas including: R&D for networking protocols; leading the development of a cross platform graphics library used in many previous generation EA games; chief architect and group technical director for the central technology group; and most recently his current role. He is currently most interested in finding ways to lower the long term cost of developing games and in how to effectively write games that take advantage of highly concurrent hardware.

Kari Pulli is a Research Fellow and Member of CEO Technology Council at Nokia Research Center Palo Alto. Kari has worked in computer graphics since 1989, when he started working on his Master's project, a parallel graphics system running on transputers. Kari has a PhD in graphics and vision from University of Washington, Seattle, and worked during grad school on graphics at Microsoft / SoftImage, SGI, and Alias | Wavefront. He worked at Stanford University graphics lab 1998-99, and joined Nokia in 1999, where he has worked on research and technology of mobile graphics and imaging, contributing to many Khronos standards and M3G (Mobile 3D Graphics for Java). In 2004-06 he was a visiting scientist at MIT.