

# LEADING NEW DEVELOPMENTS IN VISUAL COMPUTING

## Intel's Visual Computing Program in Academia

Advancements in computing technology involve more than simply increasing speed and reducing size and power consumption. Though clearly important, enhancements to features, capabilities, and specifications must go hand-in-hand with education of the development community. Indeed, for a platform to gain long-term success, its maker must engage not only with the software developers of today—gathering requirements and collaborating on design—but also with those teaching and learning the trade.

For Intel, university outreach has been a key part of company strategy dating back nearly to its founding in 1968. At the foundation of its academic collaboration is Stanford University, which produced the team that spawned Intel's flagship microprocessor and many hundreds of engineering careers.

Overall, the Intel academic community encompasses professors, faculty, and researchers from more than 1,000 universities worldwide. Each year Intel presents its Leadership in Academia Award to some of these leaders, recognizing their outstanding contributions to computer science information technology, including creative use of Intel curriculum, training of fellow educators, review of beta curriculum, creation of case studies, blogging, and more.

In 2008, Intel and Stanford launched a new collaborative effort focused on visual computing. More than 20 universities are now part of this higher-education program and are working on projects that demonstrate how visual computing architecture can simplify and enhance the process of developing realistic simulations. This article highlights four of those projects.

## The University of North Carolina at Chapel Hill

**DINESH MANOCHA** is a distinguished professor of computer science at the University of North Carolina at Chapel Hill. Along with his co-investigator **MING LIN**, Manocha is working to develop a method of rendering sound not from sampling but from the physics of objects in motion. The projects he oversees include those exploring sound synthesis, which creates sounds from the principles of physics, and sound propagation, which explores the movement of sound once emitted.

"Geometric propagation is inspired by all the work in ray tracing," said Manocha. "And vertical simulation uses scientific methods to solve these equations. "[Recent] results are much faster and are based on our new algorithms."

Numerous practical applications for the study of sound propagation exist, and Manocha mentioned a few of them along with their associated challenges:

- Building acoustics and its effect on members of an audience: "Once you have a sound, how does it transmit? For example, do all the people in the audience have a good experience in terms of echo?"
- Virtual reality, such as flight simulators: "People develop terrain visual representations, and at the same time we want to do sound simulation to get the feeling of total immersion."
- Aircraft and automobile manufacturing: "For companies manufacturing aircraft or automobiles, engine noise is among the biggest problems. When you design a new car or airplane, what will the level of engine noise be if you change the seats from cloth to leather?"
- Urban simulations: "We have houses next to the highway and they want to build a wall to block the sound. How tall should the wall be?"

To deliver the required accuracy for their projects, Manocha's team had to solve two main problems. For the highly complex physics and extremely sophisticated mathematical models demanded by the sound synthesis project, they developed what Manocha calls "a very simple algorithm in terms of computation cost that will exploit a set of features of human perception—what we hear well and what we can't hear well." For that, Intel® processors were able to operate at rates gamers want.

The second problem, associated with sound propagation, is the expense associated with processor cycles and having to take into account the emitter, the receiver, and everything in between. "How does the sound wave, which starts from the

speaker, hits walls, gets reflected, deflected, and refracted, eventually reach the receiver?" Manocha said. "Even with some of the fastest systems in the world, we still don't get performance good enough for games."

"We have an Intel processor-based system with 16 cores— four quad-core X7350 chips—and we have to use all 16 cores for our sound propagation algorithms. In sound, we don't just shoot rays, we shoot volume. We can show a cathedral or a city with multiple traffic sounds and get only about three to four frames a second. So we are still one order of magnitude slower than what the community wants; gamers want 30 to 60 frames a second."

Manocha believes the solution rests in improvements to parallelization. "That's a challenge the industry has been trying to solve for 30 years, and it's something we're doing with sound propagation. Then even when you develop a good algorithm, how do you code and get good performance? Clearly, we have to develop algorithms to parallelize."

### University of California, Berkeley

**JAMES O'BRIEN** is an associate professor of computer science at the University of California, Berkeley, where his primary area of interest is computer animation, with an emphasis on generating realistic motion using physically based simulation and motion capture techniques.

Apple Mac Pro\* computers—each with two quad-core Intel® Xeon® processors (formerly code-named Nehalem) are at the heart of his research to render simulations of a flexible needle used in the treatment of prostate cancer. Current methods employ rigid needles and ultrasound, a difficult method that

provides very little detail for the surgeon. The new flexible needle, which he and his team is helping to develop, has a beveled tip, allowing it to travel in a circular path and avoid vital organs and bones.

"This is surgical simulation," O'Brien said. "So the basic idea is to realistically model the process of inserting a needle into living tissue for the purposes of modeling biopsies—a technique called 'brachytherapy'—where the physician wants to kill cancerous tissue without hurting healthy tissue. By rotating the base of the needle you can control what direction it curves in and steer it around parts of the body you don't want to penetrate. It's hard enough to do that with a rigid needle, imagine doing it with a flexible one."

The simulator was designed to model the interaction between needle and tissue—the motion of the needle as it tries to pierce the tissue—and do so in real time or faster. Achieving the necessary accuracy required the development of new mathematical techniques for coupling the two simulations—the needle simulation and what's called the "elastic rod and the tissue finite element deformable simulation."

"Because there are forces that are interchanged, the needle pushes on the tissue and the tissue pushes on the needle," O'Brien said. "We had to make sure that was done (1) accurately, otherwise the whole exercise would be pointless, and (2) fast and in a stable fashion so the simulation can still be real time."

O'Brien said he finds working with finite element simulations interesting in part because they're very parallel. "The work being done on one element is being done in every other element in the simulation, and there are thousands

of them. On the other hand, it's an unstructured mesh that's being dynamically restructured. The structure changes as the needle goes through it. So a lot of the algorithms that assume a regular layout of the data or grid-like structure are easy to parallelize. And when we're trying to parallelize code that has these unstructured data sets, it is a little more challenging but it does map nicely to multi-core."

Are there non-surgical applications for his work? "Oh yeah," O'Brien said. "One technique is a needle going into a human being, another is a Jedi warrior going around smashing things. The two use very similar underlying simulation methods."

### St. Petersburg State Polytechnical University

Driver-Inter Ltd. is a Russian 3D game developer affiliated with the St. Petersburg State Polytechnical University. **VLADIMIR BELYAEV** was a technical lead at Driver-Inter, where he worked on shallow water rendering and more recently on ways to improve the realism of simulated grass. He does most of his development on quad-core systems.

According to Belyaev, the main goal for development of realism for his grass simulation was to focus on bodies moving through the grass. "Imagine waves across a grass field. And with the waves you can see [figures] standing on a hill and looking down. Then you see somebody in the grass, and it [divides] and you can see the trail behind the person. If you go down and see the trail close up, you'll see each grass blade, and some laying flat where they were stepped upon."

Belyaev admits that handling all the variables, whether having to do with storage, control, or manipulation, presented a challenge. "There is a lot of

information you have to handle because if you want to make this truly real, you have billions of values. If you want to make a field of grass, you need a level of unique detail. If you stand on the shore, you need a level of detail around you. This is the same for grass—you only see blades of grass within 40 or 50 meters around you. After that, you see a field of grass and the trail of the person who walked through the grass.”

Another difficulty arose in rendering the varied levels of detail that inevitably coexist in a scene. “You have to design two physical models,” Belyaev said. “One is for the high detail of one blade of grass with some additions to make it easy to handle and get less screens with the same quality. So the first challenge was to handle all the information and a large amount of data. The second was to handle the collisions with the body and the right design of the physical models. It was extremely challenging to make this all look seamless.”

### NHTV University

As part of its ongoing collaboration with academia around the world, Intel also works with universities on developing the next-generation visual computing curricula. For example, at NHTV Breda University in the Netherlands industry veterans run International Game Architecture and Design (IGAD), a four-year program that twice has been voted the best vocational university program in the Netherlands.

Program modules include C++ development of a ray tracer and software rasterizer, plus pixel shaders and hardware APIs. NHTV expects to achieve around 300 million rays per second, benefiting game developers and visual artisans alike.

**JACCO BIKKER** is an instructor at IGAD and is also a PhD student and senior lecturer for the International Architecture and Design program at the NHTV University of Applied Sciences in the Netherlands, where he teaches C/C++ and graphics programming.

A 10-year game industry veteran, Bikker is currently working on a real-time ray tracer called Arauna, which will render realistic 3D images of virtual worlds. The use of fast multi-core processor technology is playing a significant role in helping to accelerate the engine.

“A sufficient frame rate for games is about 30 frames per second,” Bikker said. “For multi-player, you want to see a higher rate. For resolution, we started with a low 400 pixels by 300 pixels because we anticipated we would have students developing on laptops. We have eight-core machines and we get

## INTEL RESEARCH PROJECTS WITH UNIVERSITIES AROUND THE WORLD

### FOUNDATIONS

**University of Colorado: Dirk Grunwald**  
Signal Processing Algorithms

**Carnegie Mellon: Franz Franchetti**  
Generating IPP Library Functionality Using SPIRAL

**Georgia Tech: Hyesoon Kim**  
Thread Fairness in Many-Core Architectures

### GRAPHICS: CURVES & SURFACES

**Columbia: Eitan Grinspun**  
Simulation and Control of Curves and Surfaces

**UCLA: Joseph Teran**  
Multi-Grid Algorithms for the Simulation of Nonlinear Deformable Solids, Designed and Optimized for the Intel® Multi-Core Platforms

### GRAPHICS: FLUID INTERACTION

**Technical University of Munich: Ruediger Westermann**  
Real-Time Fluid Interaction for Games

**St. Petersburg State Polytechnical: Vladimir Belyaev**  
Real-time Interactive Shallow Water Simulation

### GRAPHICS: GLOBAL ILLUMINATION

**Oregon State University: Mike Bailey**  
Real-Time Radiosity

### GRAPHICS: IMAGE PROCESSING

**Purdue: Chris Hoffmann**  
Real-Time Image Processing and Field Coherence

### GRAPHICS: IMAGE RECONSTRUCTION

**University of Washington: Steve Seitz**  
Build Your Own World

### GRAPHICS: LIGHTING

**Dartmouth: Fabio Pellacini**  
Interactive Realistic Lighting on Many-Core Architectures

### GRAPHICS: PHYSICS SIMULATION

#### (FRACTURE/DESTRUCTION)

**Stanford University: Ron Fedkiw**  
Real-Time Destruction

### GRAPHICS: RAY TRACING

**NHTV Breda (Netherlands): Jacco Bikker**  
Real-Time Ray Tracing for Games on Consumer Hardware



## INTEL RESEARCH PROJECTS WITH UNIVERSITIES AROUND THE WORLD

### GRAPHICS: RENDERING

**Cornell: Kavita Bala**

Interactive Scalable Complex Illumination for Realistic Rendering

**MIT: Fredo Durand**

Decoupled Sampling for Many-Core Graphics Hardware

**St. Petersburg State Polytechnical: Vladimir Belyaev**

Grass Simulation & Rendering

**Stanford University: Pat Hanrahan**

A Hardware-Accelerated Generalized Graphics Pipeline

**University of Illinois: John Hart**

Programmable Rasterization

**University of Texas, Austin: Don Fussell**

Software-based Rendering Techniques

**University of Victoria: Bruce and Amy Gooch**

Novel Applications and Data Structures for Reconfigurable Graphics Pipelines

### HIGH-PERFORMANCE COMPUTING (HPC)

**University of Victoria: Nigel Horspool**

Multi-Core Just-In-Time Compilation

### HPC: CIBR, AUDIO, FINANCIAL MODELLING

**UC Berkeley: Kurt Keutzer**

Application Research for Many-Core Platforms at the Berkeley PALLAS Project

### HPC: CROWD SIM / GRAPHICS: PHYSICS

**University of North Carolina, Chapel Hill: Dinesh Manocha and Ming Lin**

Real-Time Sound Simulation from Motion Physics

### HPC: SPEECH RECOGNITION

**UC Davis: John Owens**

Speech Recognition on Commodity Data-Parallel Processor

### INTERACTIVE MEDIA

**UC San Diego: Sheldon Brown**

The Scalable City

### MEDICAL: IMAGING

**Friedrich-Alexander University Erlangen-Nuremberg:**

**Joachim Hornegger**

Implementation and Evaluation of Medical Image Processing Algorithms

**University of Michigan: Satish Narayanasamy**

Efficient Execution of Medical Imaging Applications

### MEDICAL: SIMULATION / PHYSICS

**UC Berkeley: James O'Brien**

Real-Time Simulation for Surgical Planning and Training

higher frame rates at 800 by 600 and even higher. More recently with [Intel® Core™] i7 processors, we get 1024 [by 768] resolutions without frame rate problems.

The Arauna ray tracer was specifically built with games and performance in mind. Like many real-time ray tracers, it is highly optimized and will try to keep rays together for higher throughput.

“The rays per second are more or less stable regardless of resolution,” Bikker said, “but you need more rays for higher resolution. The number of rays per second increases a little if you have higher resolutions. In general, more rays end up on the same triangles so you have better coherence, and that improves performance. The Intel Core i7 [processor] can achieve approximately 100 million rays per second.”

Bikker added that using multi-core processors compares favorably to the use of GPUs from say NVIDIA or ATI. “If you look at what the CPU can do these days, performance numbers are approaching what a GPU can do. Currently a single Intel Core i7 processor will always beat a single high-end GPU. And that’s good. And if you look at raw compute performance of a GPU, a CPU will make far better use of performance. So you have less teraflops on an [Intel] Core i7 processor than on a high-end GPU and still it outperforms the GPU.

### University Outreach

Through these relationships and dozens of others around the world, Intel is helping to build computing platforms for generations to come. By working closely with educators, recognizing their achievements, and rewarding efforts to advance Intel education programs, the company not only gains insight into academic requirements, but also helps guarantee that future developers will have the expertise to best use the power of Intel® platforms. ■