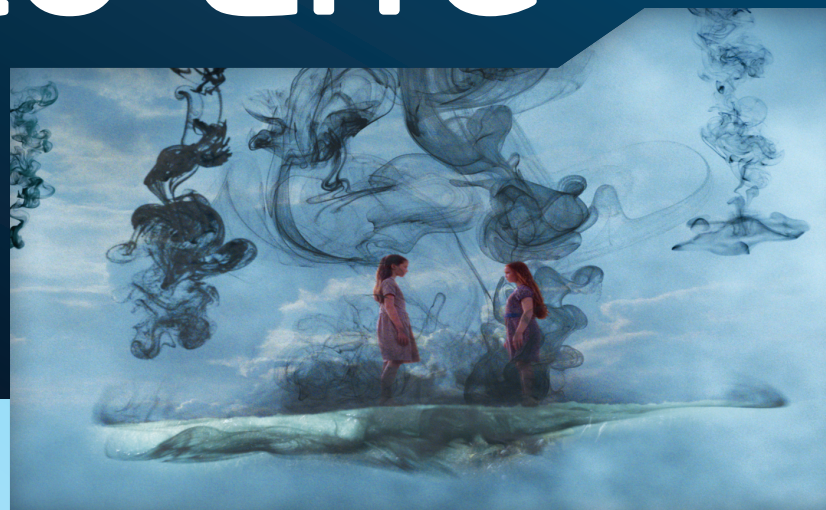# Bringing Magic to Life

EXOCORTEX AND INTEL® TECHNOLOGIES
WORK WITH VFX STUDIOS GRADIENT
EFFECTS AND WILL GARRETT TO
ACHIEVE MOVIE MAGIC

Creating Stunning Ink-in-Water Special Effects

Long-time fans of the Harry Potter film franchise are familiar with the series' ample visual effects (VFX), and this summer's box office hit *Harry Potter and the Deathly Hallows: Part 2*, the final adventure in the film series, ranks as one of the biggest VFX-driven Hollywood productions of 2011.

Among the many emotive sequences in the Harry Potter film franchise is the "memory pool" sequence created by Gradient Effects. Here is a behind-the-scenes look at how Gradient Effects, Exocortex Technology, and Intel worked together to pull off this visually stunning special effect.
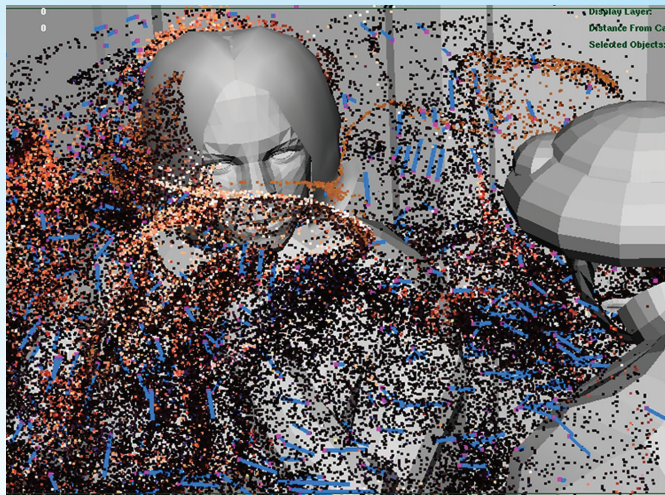
Exocortex Slipstream simulation within Autodesk Softimage*.

A Large-Scale, High-Fidelity Sequence

The memory pool sequences required an ink-in-water look where the ink would behave realistically in a large volume, then quickly assume dynamic shapes that seamlessly transition to live action footage. This complexity presented Gradient Effects with a significant challenge. Not only did it require fluid simulations of unprecedented detail and scope, but it also needed a specific "inky" look that most traditional simulators have difficulty producing, even in restricted domains. And because the sequence's timing and transitions had already been decided, there was no artistic leeway if the technology ran into limitations.

On the technical side, artists had to be able to see in real time the results of increasing the simulation resolution to more than half a billion points for final output. It was clear to Olcun Tan, co-founder and head of research and development at Gradient Effects, that traditional fluid simulators were not up to the task.

While searching for new simulation solutions, Tan discovered Exocortex's Slipstream technology. "Before committing to any technology, we did extensive research by testing all well-known technologies for the type of work required on *Harry Potter*," said Tan. "Unfortunately none were suited for the task. After seeing the first tests from Exocortex, I was instantly convinced that we had found our solution in Exocortex's technology—it was bounding-box free. The next thing I remember, I was on the phone with Ben Houston discussing how to integrate the Exocortex technology as the new engine for our proven in-house tools."

Slipstream is stable and predictable so an artist's real-time preview of a fluid simulation element is the same as the final beauty pass.

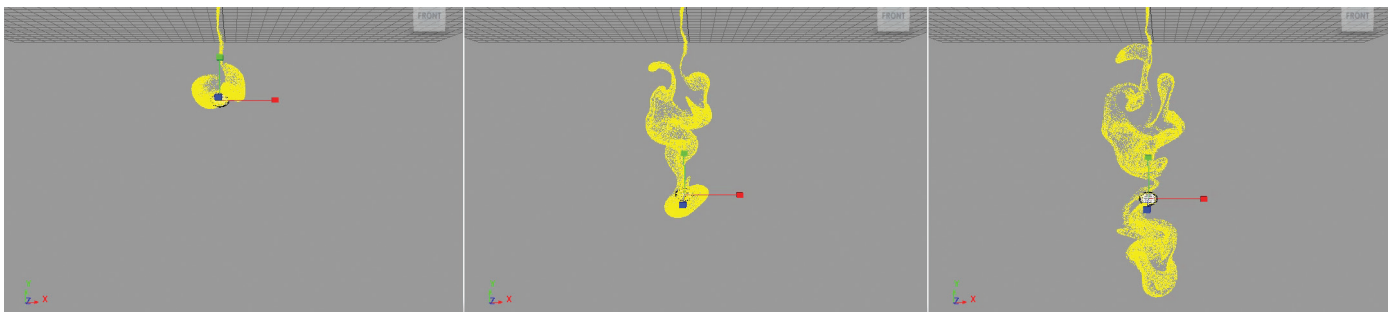## Exocortex's Slipstream Technology

Ben Houston, Exocortex's founder, has long known about the limitations of traditional simulation methods, having led the Flood fluid simulator project at Frantic Films. Exocortex was founded with the belief that the traditional limitations could be overcome, giving artists more creative freedom while controlling costs.

In the pursuit of these goals, Exocortex's team had been researching and developing its proprietary simulation technology since 2008. The result is Slipstream, a novel approach to fluid simulation that radically expands artist freedom and allows rapid accurate previews, and realistic and unbounded simulations, while cutting costs for VFX studios.

A key element of Slipstream technology is the elimination of a simulation bounding box. There is no simulation box at all, not even a dynamic moving box behind the scenes. The technology simulates freely in any environment in which one needs results at an arbitrary level of detail.

Slipstream is stable and predictable so an artist's real-time preview of a fluid simulation element is the same as the final beauty pass. Studios can handle tremendous amounts of simulation without excessive investment in hardware and personnel because Slipstream is memory and computationally efficient.

Early tests of the Harry Potter *memory pool effect using the Softimage* version of Exocortex Slipstream.*

Gradient Effects Taps Exocortex for Slipstream

Gradient Effects engaged with Exocortex after hearing about its technology. Immediately, Tan requested simulation and render tests to see if Exocortex's Slipstream technology lived up to its reputation. Exocortex's Houston quickly delivered results to Tan by modifying their pre-existing technology in combination with Autodesk's Softimage* and hardware from Intel.

Satisfied with the initial results, Gradient Effects requested further tests that demonstrated the ability to do the shape matching required in the more advanced shots of the *Harry Potter* memory pool sequences. Employing Softimage once again, Exocortex quickly designed and prototyped a set of novel extensions to its simulator to achieve the delivery results, which met with Tan's approval. Only a few days passed between Tan's initial contact with Exocortex and his decision to move forward with Exocortex's technology.

By using Exocortex's large and robust simulation-focused code base, Exocortex's Jamie Madill and Houston were able to both design and deliver an initial API that met Gradient Effects' specifications—just two weeks after the initiation of the project.

Continuing with this intense pace, three weeks later Madill and Houston delivered the final, fully multi-threaded version with advanced shape matching. The frameworks that Intel helped Exocortex develop and polish aided the rapid delivery schedule.

From Promising Research to Production Ready: Intel® Technology Makes It Happen

EARLY ON, EXOCORTEX REALIZED THAT ITS INNOVATIVE TECHNOLOGY had to be paired with a solid and performing implementation. In 2009, Exocortex began working with Intel technologist Chris Seitz.

Initially, Chris introduced Exocortex to the Intel® Threading Building Blocks library. After some fruitful initial tests that quickly produced a 2.5x speedup on consumer grade quad-core machines, Exocortex fully embraced its features. "After a couple days of integrating Intel Threading Building Blocks, the team here at Exocortex got a 2.5x performance boost," said Houston. "Our fluid simulation code is now on track to scale in a future-proof manner. The huge return we got was fantastic, and our only regret is we didn't investigate it sooner."

Under the leadership of Exocortex's Stefan Xenos, a cache-optimized C++ framework was created that eased the creation of optimal threading structures in a mostly transparent fashion. The entire code base now relies on this framework. Once the code base was multi-threaded, the team used the Microsoft Visual Studio*-integrated Intel® Parallel Studio, which exceeded Exocortex's expectations of a profiler tool. Intel Parallel Studio allowed them to rapidly identify bottlenecks in the code to ensure the optimization development time was spent wisely, achieving the greatest results. For further speed in critical sections, Exocortex also employed Intel® SIMD instructions to ensure maximum computational throughput.

The memory analysis tool-set feature in Intel Parallel Studio proved incredibly valuable in identifying unexpectedly dangerous memory accesses in the code base that otherwise would have gone unnoticed. The existence of these difficult-to-diagnose errors prompted Exocortex to design a fully checked debug memory model, ensuring that they were caught early on.

To ensure that optimizations were targeted for the needs of a high-end production environment, Intel provided Exocortex with professional grade hardware. "Exocortex Technologies is defining the state-of-the-art in VFX fluid simulation," said Seitz. "Intel has been pleased to work with Exocortex to ensure its simulators take full advantage of multi-threading and the Intel® Xeon® processor."

Designing a Modern Agile Pipeline

The project's compressed timeline—together with the scope of R&D and production involved—required the development of an agile pipeline capable of tying together software, creative talent, and Intel's hardware.

Tan started by combining Exocortex's simulator technology with Gradient Effects' proprietary technology running in Autodesk Maya*. In designing the process, Gradient Effects used Maya with massive data sets in excess of 40 GB, while retaining Maya's inherent flexibility. The team also designed a dynamic and high-quality emission system to achieve the all-important "inky" look.

Under Tan's direction, Exocortex modified its simulator technology to be fully deterministic: Feeding in the same inputs to the simulator would produce exactly the same results—every time. This addressed two very important elements of the production pipeline:

- It allowed artists to design the simulations in real time while assuring all their details would be realized as previewed when the simulators were re-run for the final beauty pass.
- A single simulation could be split across multiple machines with the results automatically synchronized.

To minimize the amount of data, camera-based culling and level-of-detail were employed. "The speed at which Tan adopted the Slipstream technology into his custom pipeline tools was impressive," said Houston. "It is a tremendous asset to Gradient Effects to have a VFX studio leader with such deep technical knowledge."

Many-Core Distributed Simulations

Using one machine per simulation proved inadequate because of the unprecedented detail and scope required for each shot, combined with the need to provide a workable turnaround time and the desire to minimize compositing tricks.

Because the simulations could be distributed across multiple machines, between 10 and 20 machines could be assigned to each simulation. This enabled an astounding 480 cores and 960 GB of memory to be simultaneously dedicated to a single fluid simulation. The approach proved an effective time saver with respect to simulation times, but the raw intensity of the computational power presented its own challenges.

A Storage System for High Compute Intensity

Gradient Effects' simulation and rendering farm consisted exclusively of Intel® Xeon® processor-based machines, each with 24 cores and 48 GB of RAM. Intel's processors churned through the simulation much faster than the simulation data could be transferred across the network to the storage servers. The team determined that the network was acting as a bottleneck on their ability to rapidly produce the necessary sequences.

At the same time, initially writing 1 GB per frame per computation node also created a bottleneck. To solve the problem, an innovative distributed storage system was created using the simulation nodes themselves.

Kelly Bergougnoux, IT manager and system engineer at Gradient Effects, recognized that the speed advantages of local storage greatly outweighed the minimal CPU cost per simulation node of managing a distributed storage system across the simulation nodes themselves. Intel's many-core processors already offered significant CPU power, a small portion of which could be assigned to managing the distribution store. Even assigning one core to this task would reduce the overall machine's simulation performance by only less than 5 percent, significantly less than waiting for the completion of large data transfers to a storage server.

To make the distributed store more robust, the team deployed a background data-mirroring process. This simulation slicing allowed for the use of 300–500 million elements per frame of the simulation and a turnaround time of a mere few hours.

The artists at Gradient Effects were impressed with the results. "Distributing the simulation, combined with the speed of the simulation engine, allowed us to simulate hundreds of millions of particles whereas in that same time other tools could simulate only thousands," said Dave Rindner, senior effects artist. "The biggest scenes required at most a few minutes per frame for simulation with particle counts of half a billion. Smaller scenes were only seconds per frame."

Manny Wong, lead effects artist, echoed Rindner's thoughts: "The system developed by Gradient Effects and Exocortex is pretty amazing. The amount of particles that can be manipulated is in the millions, and playblasting in Maya is in real time or close to it. And when you get a beautiful simulation of ink motion, all that's required is cranking up the amount of particles to give it that inky look."

In the future, solid-state drives (SSDs), such as the Intel® SSD 550, are likely a good investment for render farm machines and could further enhance the performance of farms under load. Exocortex is already adding Intel SSDs to its artists' machines.

Because these sequences, particularly those featuring the whale, would be critical for maintaining the movie's emotional hold on the audience, it was important to choose the right technology.

Remaking of a Classic: *Moby Dick*

In summer 2009, Canadian VFX supervisor Will Garrett was involved in pre-production on Germany's Tele Muenchen Group TV remake of *Moby Dick*. Filming was scheduled to start in early fall on location in Canada and Malta with William Hurt, Gillian Anderson, and Ethan Hawke in the lead roles.


The CG star of Moby Dick.

Exocortex Maelstrom* was used for all underwater shots of Moby Dick (left) as well as the live-action integrated chase shots on the ocean surface (below).

Simulating Open Oceans

The large production presented Garrett with a number of VFX challenges, including finding a cost-effective way to replace William Hurt's left leg with a peg for specific shots (which was achieved, by the way, by using a color-neutral gray sock with reflective markers and a lot of skilled compositing). He also needed to incorporate both live action footage and a computer-generated whale in many of the ocean sequences.

Because these sequences, particularly those featuring the whale, would be critical for maintaining the movie's emotional hold on the audience, it was important to choose the right technology while staying within the made-for-TV budget. Garrett selected Exocortex for its innovative technology, experience with liquid simulation, and ability to quickly adapt its software to the specific needs of the *Moby Dick* production.

Exocortex's Maelstrom Technology
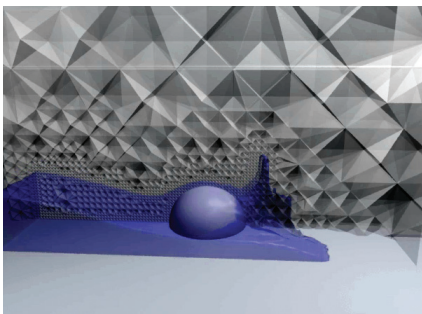
Adaptive tetrahedral cutaway.

While Gradient Effects used Exocortex's Slipstream technology for *Harry Potter and the Deathly Hallows: Part 2*, Garrett chose Exocortex's Maelstrom technology for his project. Exocortex's Maelstrom is similar to Slipstream in that both were designed to overcome the speed and scalability limitations of existing simulation approaches.

However Maelstrom uses a multi-patent-pending adaptive tetrahedra simulator core, a first in the VFX industry. The development of this unique adaptive tetrahedra simulator is the result of a three-year collaboration between Exocortex and Christopher Batty, a renowned University of British Columbia and Columbia University researcher.

Moby Dick Productions was the first production company to experience the benefits of this new approach to liquid simulation.

"Garrett was fun to work with," said Exocortex's Houston. "He has that rare and very effective combination of great interpersonal skills, an eye for detail, and the ability to leverage new technology in demanding situations."

CJK

William Hurt's leg had to be skillfully replaced with a peg to achieve the right look for Captain Ahab.

Building a Custom Solution for *Moby Dick*

The first stage of the 10-month project involved integrating the core simulator technology into Autodesk Softimage to allow rapid iteration and the ability to rerun simulations to add further detail.

Next, the team applied Batty's research innovations to achieve high-quality interaction between the liquid simulator and the intricate meshes that represented the whale and the boats. A spray, foam, and bubble system was also created to handle breaking-water situations.

Finally, Exocortex pushed the scalability of the simulator by automating and optimizing its adaptive nature. For this last task, using Intel® Parallel Studio was essential to Exocortex's development effort.

The project was a joint learning experience: Both Exocortex and the *Moby Dick* team learned how best to use the simulator and improve the software to meet the needs of the production.

The end result was a uniquely scalable, adaptive tetrahedral-based liquid simulation system tightly integrated into an Autodesk Softimage workflow. Because Exocortex relied on the Intel® Threading Building Blocks library, the software was able to take full advantage of Garrett's investment in an Intel Xeon processor-based render farm.

The extensive feature set of the simulator allowed its use on a wide variety of water shots, including whale chases and underwater sequences. In the end, over 150 shots used the Exocortex liquid simulator.

"Exocortex's custom software, combined with the Intel Xeon processor-based hardware, significantly enhanced what we were able to achieve on a made-for-TV production budget," said Garrett.

Conclusion

With the help of Intel® tools and technology, Gradient Effects and Will Garrett were able to incorporate Exocortex's technology into their tool sets, enabling their artists to create awesome VFX in a fraction of the time previously possible. Their efforts helped bring the magic to life in both *Harry Potter and the Deathly Hallows: Part 2* and *Moby Dick*. ∎