Corporate Fact Sheet

RapidMind At-a-glance

Based in Waterloo, Canada, RapidMind is a venture-backed private company that is built on over five years of advanced research and development. The company fills a gap exists in the programming ecosystem, which is not equipped to tap the enormous performance benefits of multi-core processors. Traditional serial programs written using a typical C++ programming approach will not perform well on multi-core processors and the multithreading techniques suggested are extremely challenging and do not scale as more cores are added.

While traditional development tools are insufficient, the RapidMind platform is designed to address these gaps. The RapidMind platform makes programming these processors as easy as single-threaded, single core programming, yet takes full advantage of all available resources.

Holding	Private
Investors	BDC
Founded	2004 (as Serious Hack, Inc.)
Global headquarters	Waterloo, Ontario, Canada ('Canada's Technology Triangle')
Founders	Dr. Michael McCool and Stefanus Du Toit
Executives	Ray DePaul (President and CEO) and Matthew Monteyne (VP, Sales and Marketing)
Awards	IEEE Spectrum 2007: RapidMind Development Platform named 'One of the Five best technologies of the year'

History

For over five years, Prof. Michael McCool at the renowned University of Waterloo has been leading research into advanced programming interfaces for the graphics processors embedded in video cards. This research, funded by the CITO, resulted in a novel programming system called Sh. The Sh system harnessed the power of these high-performance co-processors, which in a PC can have up to thirty times the computational power of the host processor, and applied it to both graphics and non-graphics applications. On June 8th, 2004, Prof. McCool and Stefanus Du Toit co-founded Serious Hack Inc. to commercialize this technology.

Since that time, the company has been renamed from Serious Hack to RapidMind. The product has also been expanded to target additional processors (such as the IBM Cell), and to enable a broader set of applications. True to its roots, RapidMind continues its commitment to creating innovative and groundbreaking solutions to complex problems.

Product

The RapidMind software platform enables software developers to deliver high performance on multi-core and stream processors. Developers of HPC and enterprise software are using RapidMind today to accelerate their applications on the GPU and the Cell Broadband Engine™. With RapidMind, software developers are able to program in standard C++, using their existing compilers and IDEs, and the resulting applications are portable and currently run on either GPUs or the Cell BE.

470 Weber ST N • Suite 203 • Waterloo • Ontario • Canada • N2L 6J2
T +1 519.885.5455 • F +1 519.885.1463 • info@rapidmind.net • www.rapidmind.net

Executive bios

Ray DePaul, President and CEO

Ray has two decades of experience bringing successful high technology products to market. Most notably, Ray spent nearly five years with Research In Motion where he was responsible for all product management activities for the award winning BlackBerry product line. Ray has also played key roles in industries as diverse as investment banking, media servers, telecommunications, networking, and operating systems. Ray holds a Bachelors of Mathematics degree in Computer Science as well as a Masters of Business Administration.

Ray also serves on the Board of Directors for RapidMind Inc.

Dr. Michael McCool, Founder and Chief Scientist

Michael is an Associate Professor at the University of Waterloo and co-founder of RapidMind. He continues to perform research within the Computer Graphics Lab at the University of Waterloo. Professor McCool has a diverse set of published papers, and his research interests include high-quality real-time rendering, global and local illumination, hardware algorithms, parallel computing, reconfigurable computing, interval and Monte Carlo methods and applications, end-user programming and metaprogramming, image and signal processing, and sampling. Michael has degrees in Computer Engineering and Computer Science.

Michael also serves on the Board of Directors for RapidMind Inc.

Stefanus Du Toit, Founder and Chief Architect

Stefanus is a co-founder of RapidMind. He has led the development and evolution of the RapidMind platform since 2003. Stefanus has extensive experience in the areas of graphics, GPGPU, systems programming, and compilers. Stefanus holds a Bachelors of Mathematics degree in Computer Science.

Matthew Monteyne, VP, Marketing

Matthew most recently held the position of Senior Product Manager for the Enterprise Business Unit at Research in Motion (RIM). Matthew was responsible for defining and bringing to market wireless solutions serving both the enterprise and small/medium business markets as part of the BlackBerry offering. Prior to joining RIM, he held a product management position at PixStream, a successful start-up purchased by Cisco Systems. Matthew has an Engineering degree and a Masters of Business Administration.

Technology primer

Processor vendors can no longer significantly scale performance by increasing clock rate. Instead, future processors will scale by including an increasing number of cores on a single chip. Unfortunately, multi-core processors are extremely challenging to program. Task-based thread programming methodologies introduce issues of synchronization, deadlock, load balancing, and non-determinism that need to be dealt with explicitly. Multi-core architectures also put additional pressure on the memory system, since the gap between off-chip bandwidth and on-chip processing power will grow exponentially. Current programming models do not address this issue at all. Finally, heterogeneous multi-core processors have been introduced that include specialized cores. This can provide enormous performance benefits, but also introduces problems with portability.

Application developers require a solution to programming that allows them to rapidly build reliable, maintainable, scalable, debuggable, portable, and future proof code. Since future processors will scale in performance by adding cores, any application built today should be designed with scalable parallelism in mind. One approach to scalable performance is data parallelism. Data parallelism is a general programming model that can scale to a large number of cores, since increasing the problem size naturally increases the amount of available parallelism. In addition, programming frameworks based on data parallelism can directly address the problems of data management and optimized memory access. Finally, data-parallel programming models are deterministic and do not suffer from synchronization and deadlock problems.

470 Weber ST N • Suite 203 • Waterloo • Ontario • Canada • N2L 6J2
T +1 519.885.5455 • F +1 519.885.1463 • info@rapidmind.net • www.rapidmind.net

RapidMind provides a software development platform that allows software vendors to deliver high performance on multi-core and stream processors, including the GPU and the Cell BE. Without sacrificing development simplicity, maintainability, or portability, the RapidMind platform allows an application developer to target both multi-core and heterogeneous multi-core processors from a conceptually single-threaded C++ program. The RapidMind interface also completely eliminates the overhead of C++, so the modularity constructs of this language can be used to structure large applications without sacrificing performance.

The RapidMind platform acts as an embedded programming language inside C++. It is built around a small set of types that can be used to capture and specify arbitrary computations. Arbitrary functions, including control flow, can be specified dynamically. Parallel execution is primarily invoked by applying these functions to arrays, generating new arrays. Access patterns on arrays allow data to be collected and redistributed. Collective operations, such as scatter, gather, and programmable reduction, support other standard parallel communication patterns and complete the programming model.

Terminology

GPU	Graphics Processing Unit or GPU (also occasionally called Visual Processing Unit or VPU) is a dedicated graphics-rendering device for a computer or gaming console. Designed to be efficient at manipulating and displaying computer graphics, GPUs incorporate a highly parallel structure that makes them more effective than typical CPUs for a range of complex algorithms. Example manufacturers of GPUs include ATI and Nvidia.
Cell BE	Cell is the abbreviation of Cell Broadband Engine Architecture. Jointly developed by the STI alliance (Sony, Toshiba and IBM), Cell is a microprocessor that combines a general-purpose core with streamlined co-processing elements that greatly accelerate multimedia and vector processing applications, along with other forms of dedicated computation. The first major commercial application of Cell was in Sony's PlayStation 3 game console. Mercury Computer Systems offers a dual Cell server, and Toshiba has announced plans to incorporate Cell in HD televisions.
Multi-core CPU	A Multi-core CPU (Central Processing Unit) combines two or more independent processors into a single unit or integrated circuit (IC). A dual-core device contains two independent microprocessors. As CPUs reached their maximum clock speed, manufacturers such as AMD and Intel have adopted a multi-core strategy as a means of driving continued performance increases.
Stream Processor	Stream processors process a series of processor-intensive operations in parallel, achieving increased performance. GPUs (Graphics Processing Unit) are an example of commercial stream processors.

470 Weber ST N • Suite 203 • Waterloo • Ontario • Canada • N2L 6J2
T +1 519.885.5455 • F +1 519.885.1463 • info@rapidmind.net • www.rapidmind.net