First Details on a Future Intel Design Codenamed 'Larrabee'

August 4 2008

Intel Corporation is presenting a paper at the SIGGRAPH 2008 industry conference in Los Angeles on Aug. 12 that describes features and capabilities of its first-ever forthcoming "many-core" blueprint or architecture codenamed "Larrabee."

Details unveiled in the SIGGRAPH paper include a new approach to the software rendering 3-D pipeline, a many-core (many processor engines in a product) programming model and performance analysis for several applications.

The first product based on Larrabee will target the personal computer graphics market and is expected in 2009 or 2010. Larrabee will be the industry's first many-core x86 Intel architecture, meaning it will be based on an array of many processors. The individual processors are similar to the Intel processors that power the Internet and the laptops, PCs and servers that access and network to it.

Larrabee is expected to kick start an industry-wide effort to create and optimize software for the dozens, hundreds and thousands of cores expected to power future computers. Intel has a number of internal teams, projects and software-related efforts underway to speed the transition, but the tera-scale research program has been the single largest investment in Intel's technology research and has partnered with more than 400 universities, DARPA and companies such as Microsoft and HP to move the industry in this direction.

Over time, the consistency of Intel architecture and thus developer freedom afforded by the Larrabee architecture will bring about massive innovation in many areas and market segments. For example, while current games keep getting more and more realistic, they do so within a rigid and limited framework. Working directly with some of the world's top 3-D graphics experts, Larrabee will give developers of games and APIs (Application Programming Interface) a blank canvas onto which they can innovate like never before.

Initial product implementations of the Larrabee architecture will target discrete graphics applications, support DirectX and OpenGL, and run existing games and programs. Additionally, a broad potential range of highly parallel applications including scientific and engineering software will benefit from the Larrabee native C/C++ programming model.

Additional details of the Larrabee architecture discussed in this paper include:

-- The Larrabee architecture has a pipeline derived from the dual-issue Intel Pentium processor, which uses a short execution pipeline with a fully coherent cache structure. The Larrabee architecture provides significant modern enhancements such as a wide vector processing unit (VPU), multi-threading, 64-bit extensions and sophisticated pre-fetching. This will enable a massive increase in available computational power combined with the familiarity and ease of programming of the Intel architecture.

-- Larrabee also includes a select few fixed function logic blocks to support graphics and other applications. These units are carefully chosen to balance strong performance per watt, yet contribute to the flexibility and programmability of the architecture.
A coherent on-die 2nd level cache allows efficient inter-processor communication and high-bandwidth local data to be access by CPU

cores, making the writing of software programs simpler.

-- The Larrabee native programming model supports a variety of highly parallel applications, including those that use irregular data structures. This enables development of graphics APIs, rapid innovation of new graphics algorithms, and true general purpose computation on the graphics processor with established PC software development tools.

-- Larrabee features task scheduling which is performed entirely with software, rather than in fixed function logic. Therefore rendering pipelines and other complex software systems can adjust their resource scheduling based each workload's unique computing demand.

-- The Larrabee architecture supports four execution threads per core with separate register sets per thread. This allows the use of a simple efficient in-order pipeline, but retains many of the latency-hiding benefits of more complex out-of-order pipelines when running highly parallel applications.

-- The Larrabee architecture uses a 1024 bits-wide, bi-directional ring network (i.e., 512 bits in each direction) to allow agents to communicate with each other in low latency manner resulting in super fast communication between cores.

-- The Larrabee architecture fully supports IEEE standards for single and double precision floating-point arithmetic. Support for these standards is a pre-requisite for many types of tasks including financial applications.

The paper's reference details: Seiler, L., Carmean D., Sprangle, E., Forsyth, T., Abrash, M., Dubey, P., Junkins, S., Lake, A., Sugerman, J., Cavin, R., Espasa, R., Grochowski, E., Juan, T., Hanrahan, P., 2008. "Larrabee: A Many-Core x86 Architecture for Visual Computing," *ACM Transactions on Graphics*, 27, 3, 2008.

The paper will be available at this Web site: doi.acm.org/10.1145/1360612.1360617 .

Provided by Intel