

COSC 637 4

Parallel Computations

Introduction to CUDA: GPU Architectures

Edgar Gabriel
Fall 2010



Edgar Gabriel



References

- Intel Larrabee:
[1] L. Seiler, D. Carmean, E. Sprangle, T. Forsyth, M. Abrash, P. Dubey, S. Junkins, A. Lake, J. Sugerman, R. Cavin, R. Espasa, E. Grochowski, T. Juan, P. Hanrahan:
“Larrabee: a many-core x86 architecture for visual computing”,
ACM Trans. Graph., Vol. 27, No. 3. (August 2008), pp. 1-15.
http://softwarecommunity.intel.com/UserFiles/en-us/File/larrabee_manycore.pdf
- Nvidia GT200:
[2] David Kanter, “Nvidia GT200: Inside a Parallel Processor”,
<http://www.realworldtech.com/page.cfm?ArticleID=RWT090808195242&p=109/08/2008>
- Nvidia Fermi:
[3] David Kanter, “Inside Fermi: Nvidia’s HPC Push”,
<http://www.realworldtech.com/page.cfm?ArticleID=RWT093009110932&p=1>



Larrabee Motivation

- Comparison of two architectures with the same number of transistors
 - Half the performance of a single stream for the simplified core
 - 40x increase for multi-stream executions

	2 out-of-order cores	10 in-order cores
Instruction issue	4	2
VPU per core	4-wide SSE	16-wide
L2 cache size	4 MB	4 MB
Single stream	4 per clock	2 per clock
Vector throughput	8 per clock	160 per clock



Larrabee Overview

- Many-core visual computing architecture
- Based on x86 CPU cores
 - Extended version of the regular x86 instruction set
 - Supports subroutines and page faulting
- Number of x86 cores can vary depending on the implementation and processor version
- Fixed functional units for texture filtering
 - Other graphical operations such as rasterization or post-shader blending done in software



Larrabee Overview (II)

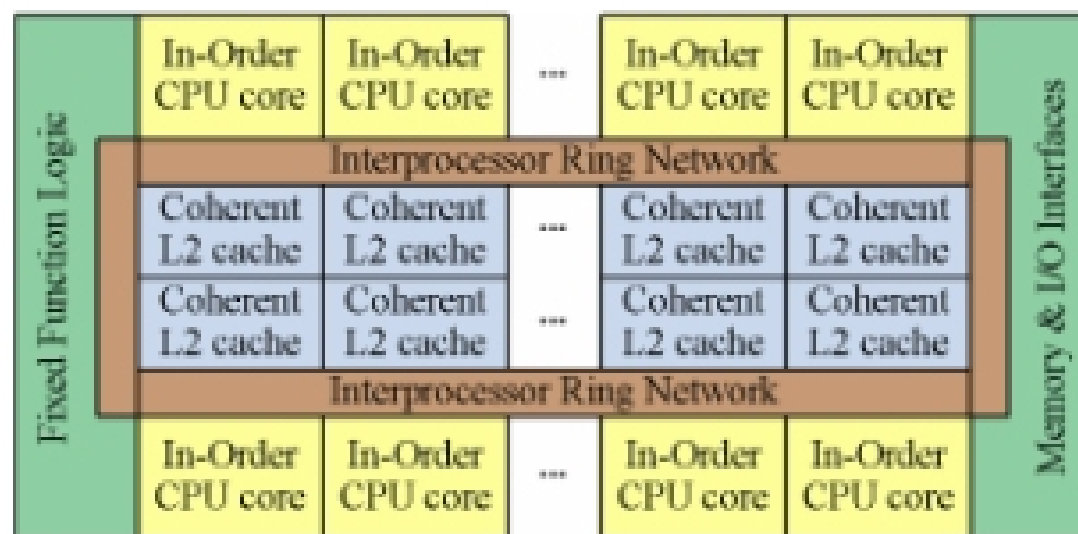


Image Source: [1]



Overview of a Larrabee Core (I)

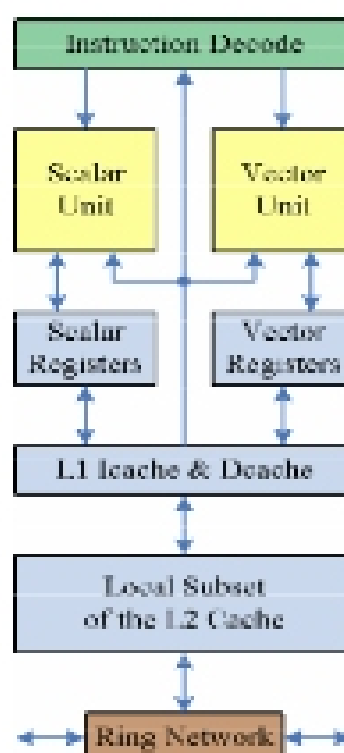


Image Source: [1]

