

Scott J. Krieder
 Department of Computer Science
 Illinois Institute of Technology
 skrieder@iit.edu

Ioan Raicu
 Department of Computer Science
 Illinois Institute of Technology
 iraicu@cs.iit.edu

Abstract

Accelerator technologies are now quite common in Supercomputers, Clusters, Grids, and personal desktops. This work will provide an overview of the current technologies that are available today, and examine future accelerator technologies. This work examines the 3 major competitors in the Accelerator market; nVIDIA, Intel, and AMD.

Why Accelerators?

Accelerators allow the machine to offload work from the CPU to the Accelerator. The Accelerator then completes the computation and returns the solution back to the host CPU.

Motivation


- Many Task Computing
- Workflow Systems
- Many Small Scientists
- Programmability
- Bridge gap between Clusters and GPU compute
- Swift + Accelerators

Advantages

- Many cores
- Power efficient

Disadvantages

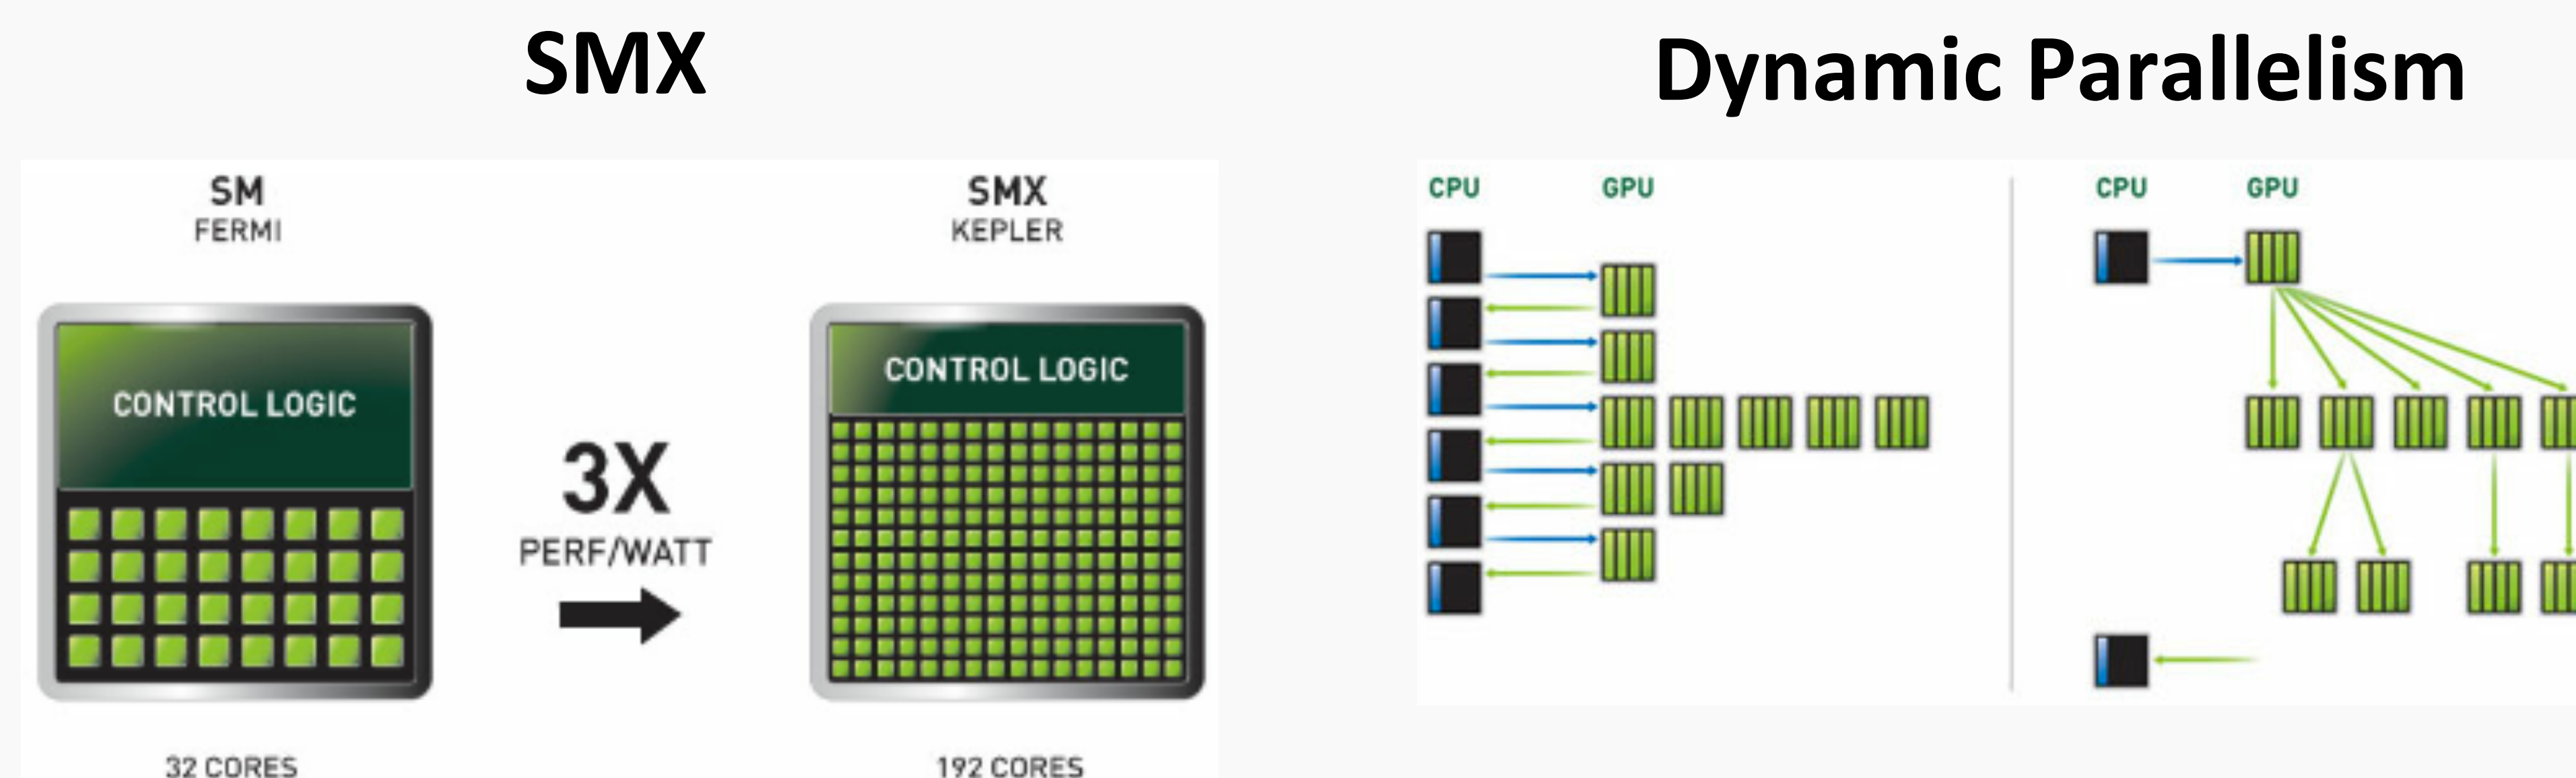
- Data Transfers
- Communication
- Programmability




References

NVIDIA

Kepler Architecture



Hyper Q



CUDA 5.0

Advanced Libraries
 GPU to GPU via. P2P
 C, C++, Fortran

Products

GeForce = Desktop Graphics
 Quadro = Professional Graphics
 Tesla = Server Edition

Intel

Many Integrate Cores (MIC)

- Architecture**
- 22 Nanometer Chip
 - 50 Intel Cores
 - X86
- Programmability**
- C Programming
 - C++ Programing
 - Fortran Programming




- Products**
- Knights Ferry
 - Research Product
 - Knights Corner
 - Future Product
 - Stampede at TACC

AMD

Programmability

- OpenCL
- JavaCL
- Fortran
- C++
- OpenCL .NET
- Pythoncl



Architectures

- GPU 40nm Firestream
- APU 32nm Fusion

Conclusions

Device	Pro	Con
NVIDIA	Maturity Raw Perf.	Programming
Intel	Programming	Availability
AMD	Openness	Adoption

Future Work

- GPU Virtualization
- Workflows with Accelerators