

# *Many Core Disruptive Technologies are Coming!\**

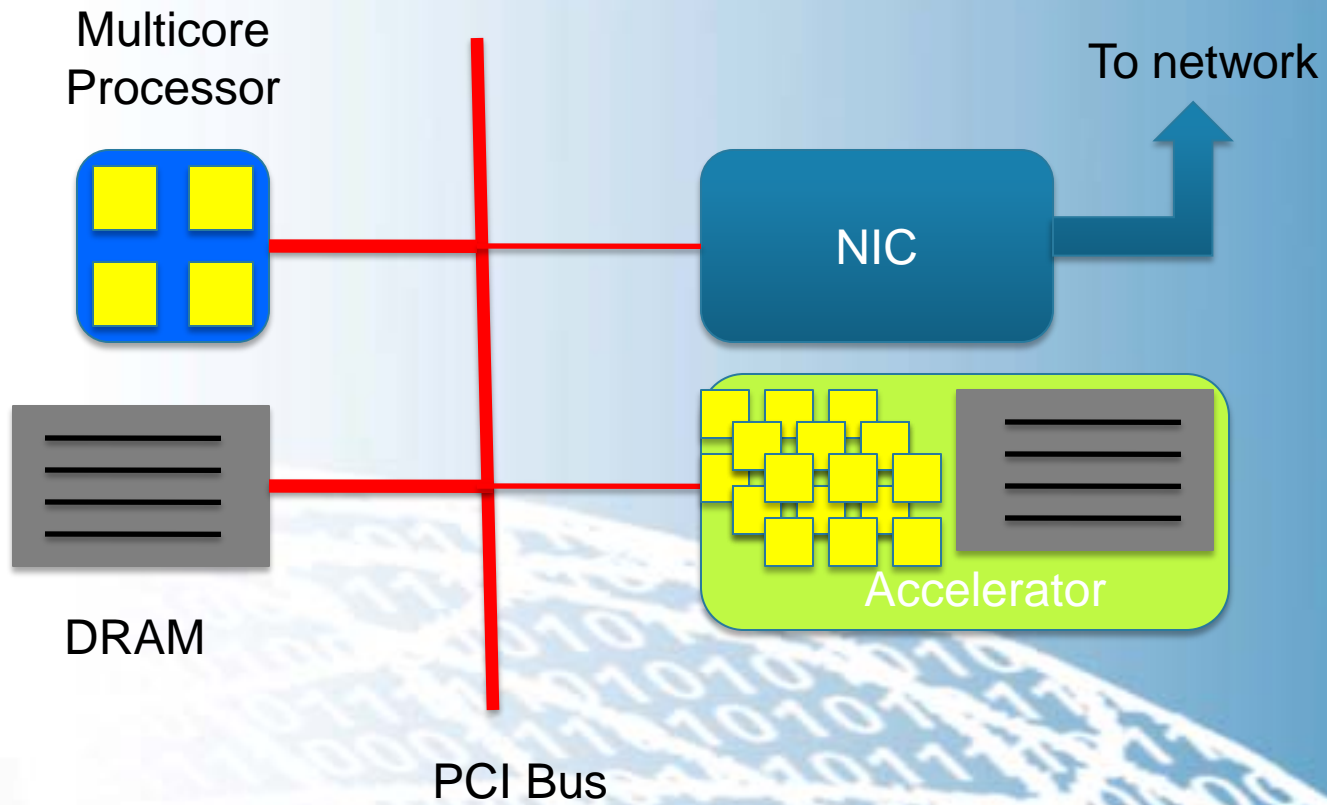
**Dr. Richard Loft  
Director, Technology Development  
Computational and Information Systems Laboratory**

**Ocean Modeling Working Group Meeting  
December 15, 2011**

# What do I mean by “Many Core”?

- 2005 processor clock speeds have stagnated
- Why? **Power consumption** of high-Hz silicon
- Design strategy: many **cores per socket clocked at low-Hz + many threads per core**
- Where are we going? processors with **hundreds of cores and thousands of threads**
- But it gets worse...

# Accelerator Node Architecture



# ***Now Moving to Many Core:***

XSEDE: TACC Stampede (2013)

10 PF

Intel MIC + Intel Sandy Bridge

DoE: ORNL Titan (2012)

~20 PF

18K NVidia Kepler GPU's + AMD Interlagos

NSF Track 1: NCSA Blue Waters (2012)

>11.5 PF

AMD Interlagos + 3000 NVidia Kepler GPU's

DoE: Argonne Mira (2012)

9.2 PF

45,152 BG/Q SoC's

# Concerns

- **Power crunch**
  - Many cores clocked @ 1.x GHz and thus use less power
  - Utility/system fit-up costs for power-hungry systems are becoming cost prohibitive.
- **Are our applications being left behind?**
  - If we can't use the large NSF and DoE systems effectively, what impact will that have on our programs?
  - We know some work is being done (WRF, HOMME) but it appears “piecemeal” and under-resourced
- **What's the right strategy?**
  - Slacker model – wait for SW/HW to “improve”...
  - Red-bull – hire a bunch of ace hackers and go for it?
  - Something in between?

# We need an integrated assessment of multiple objectives, challenges:

CESM Science Objectives

CESM Model Component Directions

Software Programming Models

Disruptive Technologies

# Possibly Many-Core Path Forward

CESM –  
Fortran+MPI+OpenM  
P  
Works well on 4-6  
core processors

Pipeline not meant to suggest that  
architectural investigations must occur  
sequentially...

Refactor code for  
higher thread  
parallelism...

Two Memory Spaces

Cuda Fortran/OpenACC  
Directives

IBM Blue Gene

BG-L – 2 cores  
BG/P – 4 cores  
BG/Q – 16x4\*

Intel MIC Processor

Kn. Ferry – 32x4\*  
Kn. Corner ~ 64x4\*  
:

Nvidia  
Graphics Card

Fermi – 16x32\*  
Kepler ~ 32x32\*  
:


**\*cores X threads/core: most cores now have some form of multithreading**

**This is hopefully the start of a broader discussion with the CESM community...**

**Thanks!**



# Blue Gene/Q: System on a Chip



## Vital Statistics

- 45 nm technology
- clocked @ 1.6 GHz
- Quad FPU's (8 flops/clock tick)
- 16 user cores + 1 system core + 1 spare
- 128 flops/clock tick
- 204.8 GFLOPS/chip
- 55W/chip
- 268 pJ/flop
- 42.7 GB/sec

Blue Gene/Q chip layout

This could be considered the Prius of modern day multi-processor systems!

# Intel MIC Processor Family

## Knights Ferry



- Software development platform
- Growing availability through 2010
- 32 cores, 1.2 GHz
- 128 threads at 4 threads / core
- 8MB shared coherent cache
- 1-2GB GDDR5
- Bundled with Intel HPC tools

Software development platform for Intel® MIC architecture



# NVidia Fermi GPU Processors

- **512 CUDA cores**
  - 32 Cores/SM
  - 16 SM
- **4x more core/SM than GT200**

