

P2S2-2012 Panel Session “Battle of Accelerator Stars”

Taisuke Boku

Professor, Deputy Director

Center for Computational Sciences

University of Tsukuba

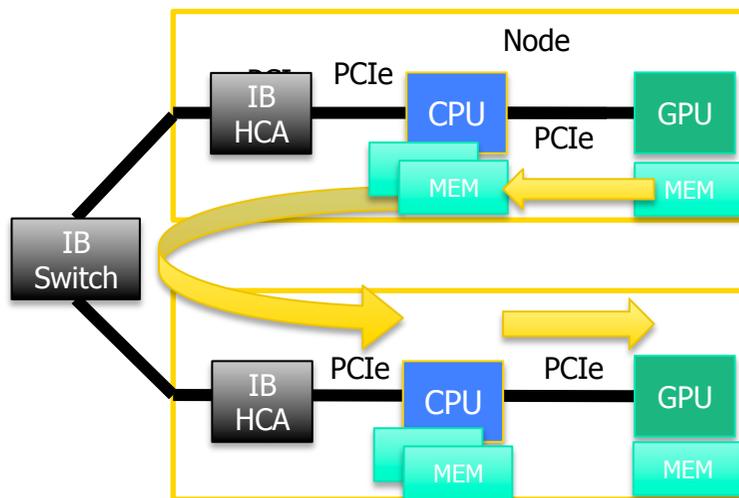


HA-PACS/TCA (Tightly Coupled Accelerator)



■ True GPU-direct

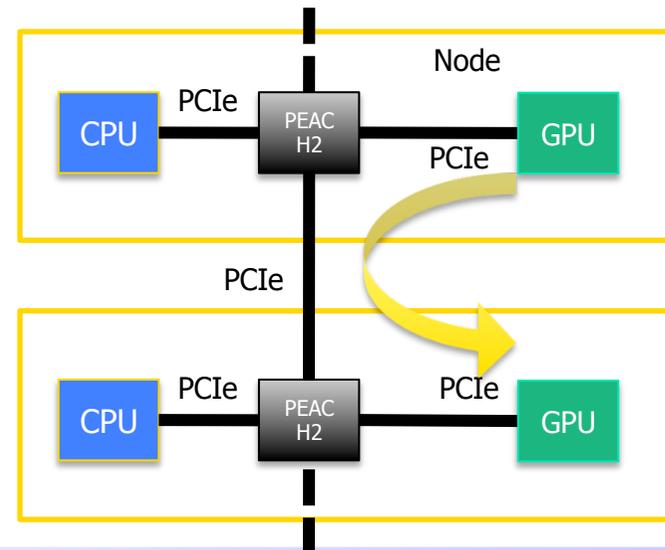
- current GPU clusters require 3-hop communication (3-5 times memory copy)
- For strong scaling, Inter-GPU direct communication protocol is needed for lower latency and higher throughput



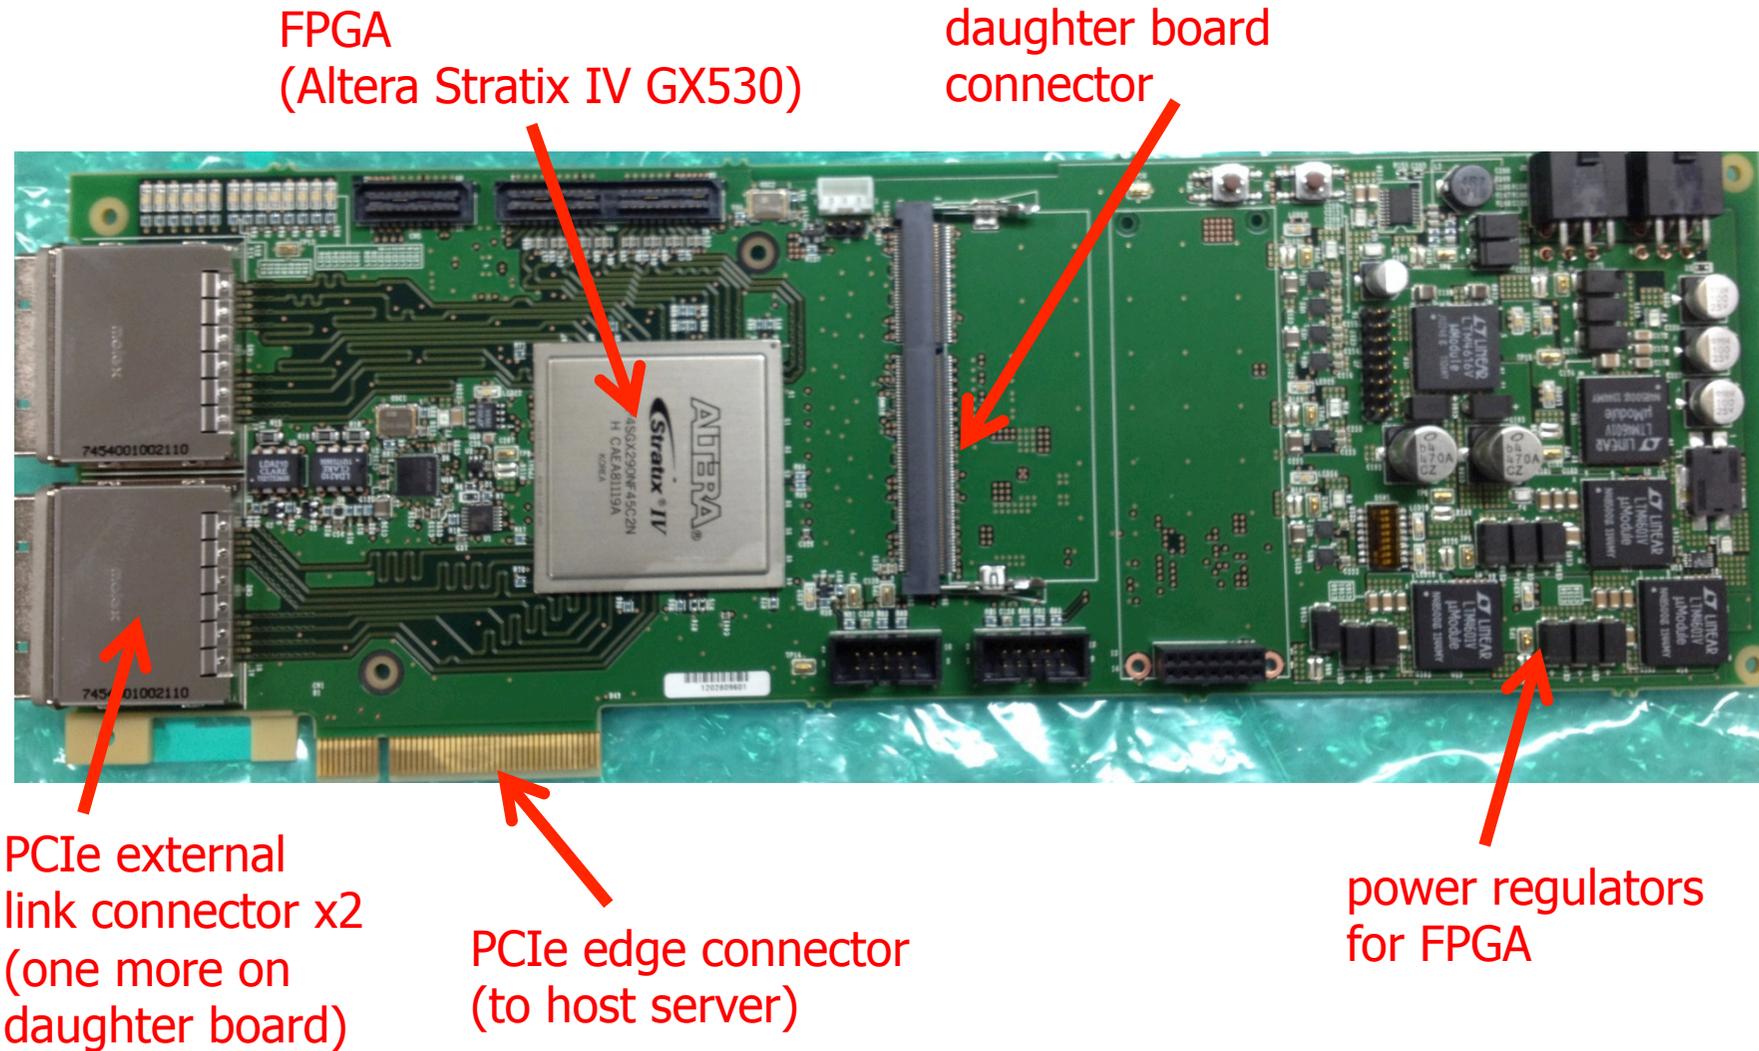
■ Enhanced version of PEACH

⇒ **PEACH2**

- x4 lanes -> x8 lanes
- hardwired on main data path and PCIe interface fabric



PEACH2 prototype board for TCA



Q1: accelerator's role in parallel processing/HPC

- ~100PF range: commodity accelerators contribute to most of “*acceleratable* applications”
 - GPU can take enough role (ex. 20PF Titan@ORNL)
 - maybe up to 1EF? (Echelon by Bill Dally)
- ~1EF range: specially embedded accelerators for ultra MPP system
 - GPU-style (= I/O device) accelerator with general CPU will be difficult to keep the requirement of performance/power limitation (1EF/20MW)
 - I/O bandwidth bottleneck
 - “very compact simple SIMD” type of accelerator should be implemented on chip as well as compact general CPU



Q2: hardware/software divergence

- **Serious problem on programmability/productivity**
 - MPI + OpenMP + CUDA/OpenCL/OpenACC
 - caused by combinations of multiple orthogonal hardware/software components
- **Divergence is not just on hard/soft**
⇒ divergence on accelerator programming
CUDA, OpenCL, OpenACC -> should be converged ??
- **Unified programming environment is strongly required to handle distributed memory + shared memory + accelerator**
 - such as XMP-dev (by U. Tsukuba)



Q3: accelerator type

- for commodity accelerator, of course GPU is the best
 - supported by general use, not just for HPC
 - but, if the required performance range by HPC overcomes what is required for general graphic processing ... ?
- Many core (MIC)
 - at this moment, a kind of accelerator
 - very general purpose with high bandwidth memory by GDR
 - from KNL, must not be used as accelerator
- FPGA
 - if just for computation power, not so powerful anymore
 - good for very specialized purpose, but memory is a serious problem
 - additional feature such as binding with communication ?
- **Hybrid chip with CPU + GPU**
 - Most of CPU & GPU vendors are considering
 - Maybe the best style keeping compatibility with traditional CPU & GPU



Q4: programming model/library

- **Unified programming model for**
 - distributed memory
 - shared memory
 - accelerator

is desired for productivity
⇒ same as Q2



Q5: critical research challenge

- **Memory capacity vs. bandwidth balance**
 - **DDR vs. GDR**
ex. K20 or K20X of NVIDIA – memory size is relatively small compared with traditional series of GPGPUs
- **I/O bottleneck is very serious issue when GPU (or any PCIe based accelerator) performance is getting more powerful**
- **CPU and its DRAM will be just a bottleneck, and how to make accelerators communication with each other directly**
- **Solution: embedded accelerator architecture coupled with a communication feature**
 - **Transputer: CPU + communication**
 - **CPU + accelerator + communication = future**

