



High Performance Fabrics

Comparing SRIO, PCI-e and 10Gbit Ethernet

MicroTCA Conference
Munich Nov 3rd

Irene Hahner

Product Management, Kontron

Gene Juknevicius

Senior Architect, GE Fanuc

MicroTCA Connectivity Options

- For communication within the chassis MicroTCA defines Common Options and Fat Pipe Port regions
- Common Options Port 0 and Port 1 are dedicated to 1Gbit Ethernet
- Fat Pipes (ports 4-11) are dedicated to high performance communication protocols:
 - > PCI-express
 - > SRIO
 - > Ethernet
- Focus of this presentation: Fat Pipes

PCI-express



imagination at work

PCI-express Introduction

- PCI-express is serialized version of legacy PCI interface
- PCI-express comes in multiple Generations
 - > Gen 1: 2.5GHz signaling
 - > Gen 2: 5 GHz signaling
- PCI-e comes in multiple Lane widths: x1, x2, x4, x8, x16

PCI-express Usage Model

- PCI-e is predominantly being used as Processor to I/O bus
- PCI-e has a Root Complex and End Node concept:
 - > Processor, being Root Complex enumerates PCI-e bus and assigns memory addresses to all PCI-e end nodes
 - > End Node waits for resources to be assigned and reset to be removed
- After enumeration is complete PCI-e devices can send/receive data to/from Root Complex as well as to/from each other

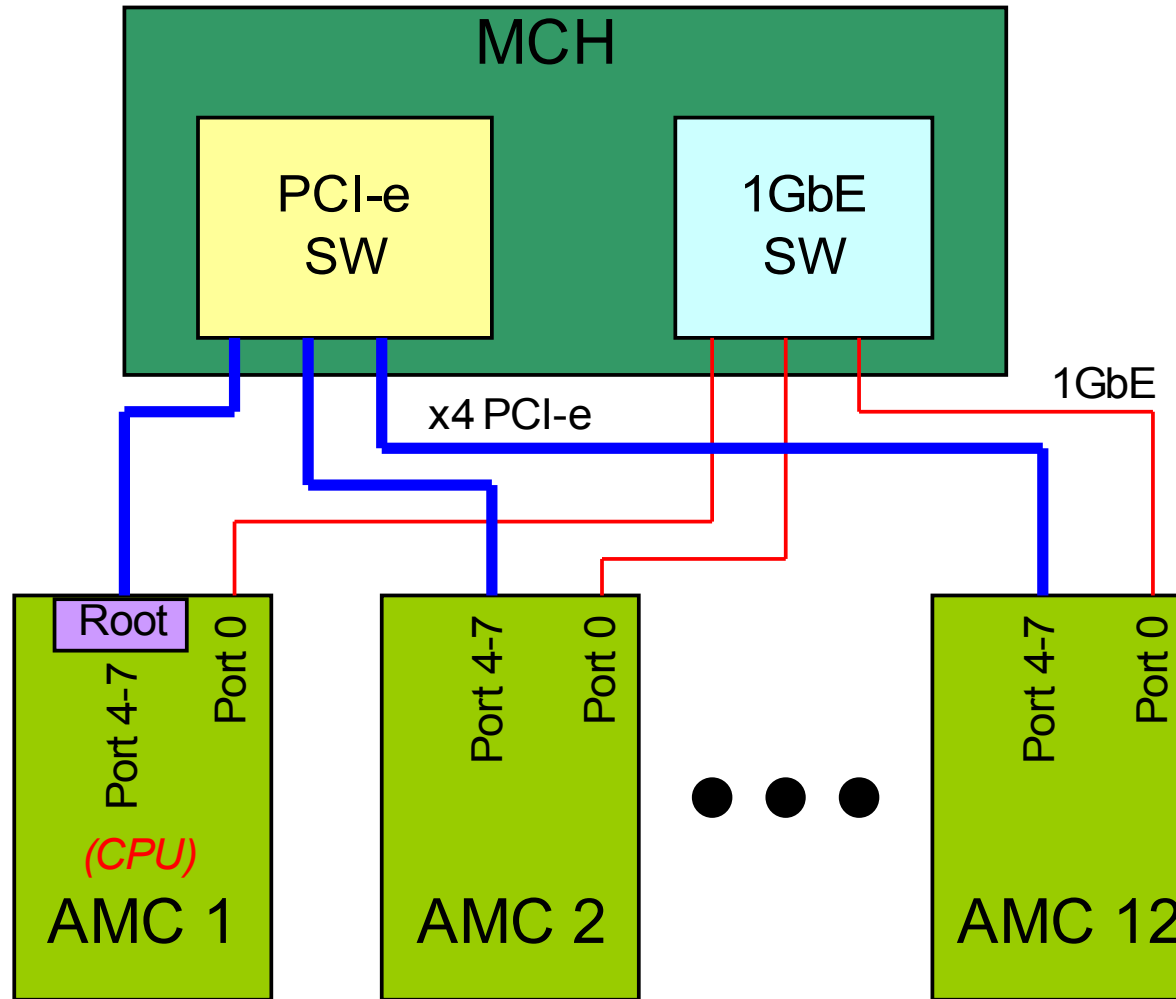
PCI-express in MicroTCA

In MicroTCA PCI-e is being used in a following way:

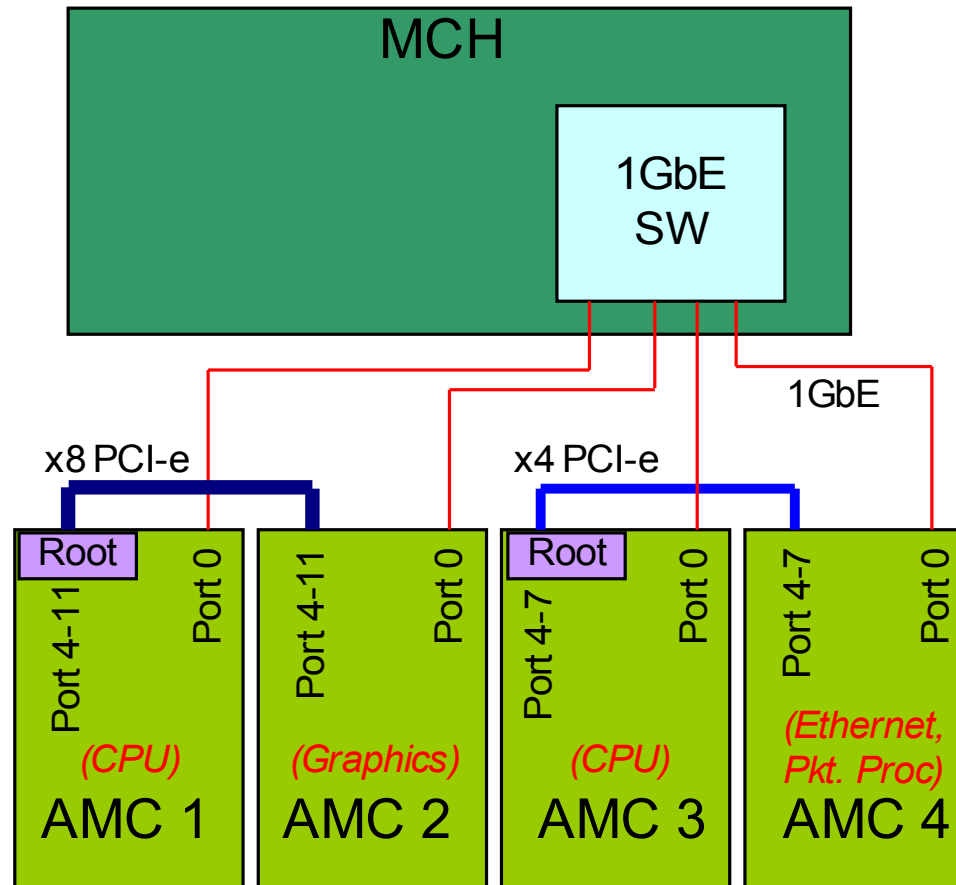
- > Fat Pipes Region: Ports 4-7

- > Note, that AMCs can have PCI-e on Ports 4-11, but MCH has at most 4 Lanes per AMC
 - x1: Port 4 2.5Gbits/sec (Gen 1)
 - x2: Ports 4-5 5Gbits/sec (Gen 1)
 - x4: Ports 4-7 10Gbits/sec (Gen 1)

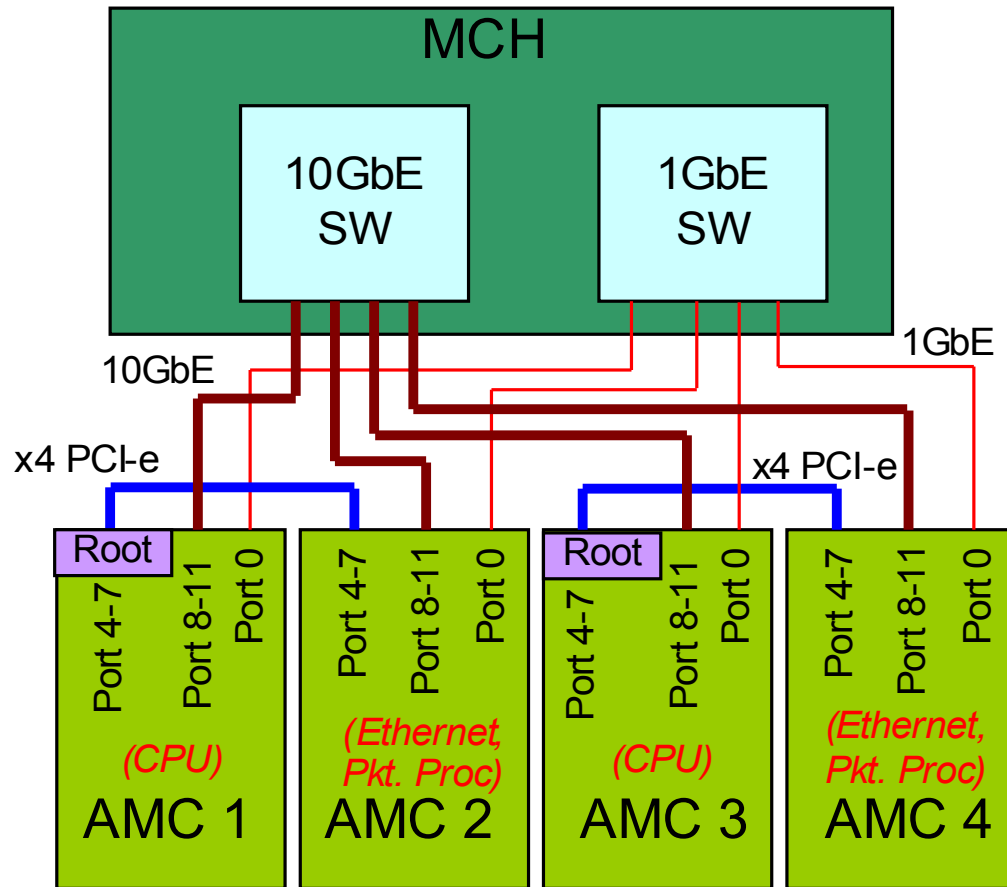
PCI-e Usage Example



PCI-e Point-to-Point Usage Example



PCI-e Coexistence with Ethernet and SRIO Example



PCI-e Performance Estimates

Type	Theoretical Max Throughput	Practical Max Throughput	CPU Utilization	Latency
x1 Gen 1	2.5Gbit/sec	~1.5Gbit/sec	Low	Low
x4 Gen 1	10Gbit/sec	~6Gbit/sec	Low	Low
x4 Gen 2	20Gbit/sec	~12Gbit/sec	Low	Low

Usage Cases

- Predominantly applications where CPU connects to number of I/O devices (Ethernet NICs, Fiber Channel HBAs, VGA cards, Analog/Digital I/Os, etc)
- Applications where devices, in addition to communicating with CPU, need to also communicate directly with each other
 - > Graphics Processing Units, GPUs, (using Crossfire mode, etc)
 - > Packet Processors (passing data directly between them)

Ethernet



imagination at work

Ethernet Introduction

- Ethernet is the most prevailing communication technology today
- It supports multiple speeds, most common today being 1Gbit and 10Gbit
- It supports multiple media options:
 - > backplane
 - > copper cable
 - > fiber cable

Ethernet Usage Model

- Most popular way to communicate via Ethernet is by using standard Socket APIs, supported by most operating systems
- Typically either TCP (for guaranteed delivery) or UDP (for best effort delivery) protocols are being used on top of Ethernet Layer-2 protocol

Ethernet in MicroTCA

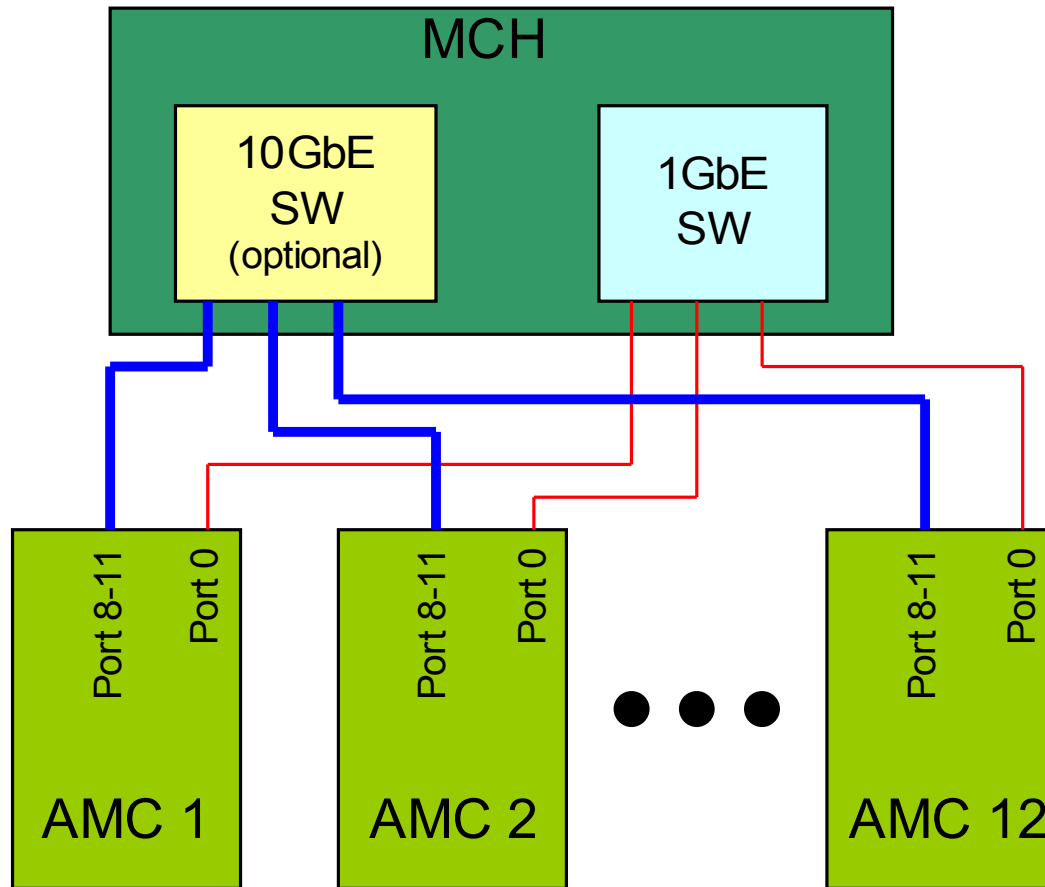
In MicroTCA Ethernet is being used in two places:

- > Common Options Region: Port 0 and Port 1:
 - 1Gbit (1000BASE-BX)
- > Fat Pipes Region: Ports 8-11, extending into Ports 20-17:
 - 4-port 1Gbit (1000BASE-BX)
 - 1-port 10Gbit (10GBASE-BX4)
 - 2nd 10Gbit on Ports 20-17 (10GBASE-BX4)


Ethernet in MicroTCA, Next Step

- > Today 10Gbit Ethernet on MTCA backplane uses 10GBASE-BX4 interface, also known as XAUI:
 - 4 Lanes at 3.125GHz signaling
- > Going forward Ethernet in MicroTCA will also support 40Gbit connectivity:
 - 4 Lanes of 10GBASE-KR (10.3125 GHz signaling)
- > Current AMC specification already supports 10GHz signaling rates and 40Gbit Ethernet has been demonstrated in MicroTCA implementations

1GbE and 10GbE Usage Example



10GbE Point-to-Point Usage Example



Ethernet Performance Estimates

Type	Theoretical Max Throughput	Practical Max Throughput	CPU Utilization	Latency
1 Gbit	1Gbit/sec	0.5Gbit: (64B pkt, TCP) 0.97Gbit: (1500B pkt, TCP)	Moderate	High (driven by Operating System)
10 Gbit	10Gbit/sec	5Gbit: (64B pkt, TCP) 970Gbit: (1500B pkt, TCP)	High	High (driven by Operating System)

Usage Cases

- Most applications where higher latency and higher CPU utilization is acceptable
- Applications where data entering/leaving the MicroTCA system is natively in Ethernet Packet Format (Deep Packet Inspection devices, Gateways, Routers, etc)
- Applications where system interconnect needs to be extended outside the box and over long distances

SRIO



imagination at work