E4 COMPUTER ENGINEERING

E4-ARKA: ARM64+GPU+IB is Now Here

Piero Altoè

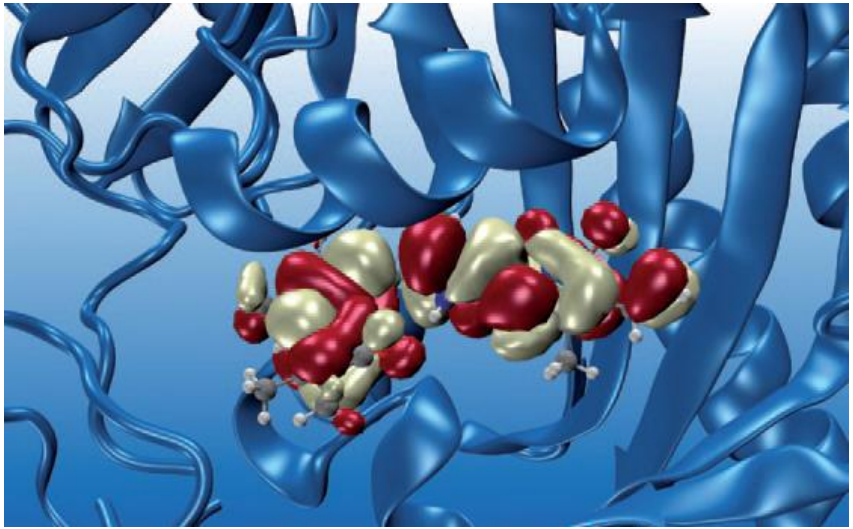

E4® Computer Engineering S.p.A. specializes in the manufacturing of high performance IT systems of medium and high range. Our products aim to accomplish both industrial and scientific research requirements and space from universities to computing centers.

Thanks to the established experience and quality in this circle, E4 has become a valued technology's supplier and it is acknowledged and appreciated by famous and prestigious organizations.

The first ARM+GPU cuda based platform 2011

Features	ARKA Blade
CPU	NVIDIA® Tegra® 3 ARM Cortex A9 Quad-Core
GPU	NVIDIA Quadro® 1000M with 96 CUDA Cores
Memory	2GB x CPU 2GB x GPU
Peak Performance	270 Single Precision GFLOPS
Network	1x Gigabit Ethernet
Storage	1x SATA 2.0 Connector
USB	3x USB 2.0
Display	3x HDMI + serial console available

CARMA Devkit
developed by **SECO**

ARKA CLUSTER SHOC BENCH 2011

The Scalable Heterogeneous Computing Benchmark Suite (SHOC) is a collection of benchmark programs testing performance and stability (using CUDA and/or openCL)

Test	ARKA	Intel+M2075	Units	ARKA/M2075
Maxspflops	263.12	1001.31	GFlops	26%
fft_sp	23.343	162.524	GFlops	14%
sgemm_n	132.482	666.272	GFlops	20%
dgemm_n	21.5152	315.110	GFlops	7%
md_sp_bw	5.5301	25.3975	GB/s	22%
Reduction	23.1895	92.7343	GB/s	25%
Sort	0.3090	1.6296	GB/s	19%
triad_bw	0.4279	6.0163	GB/s	7%

How we started 2012

Tegra 3 Q7 Module
4x ARM Cortex A9 @ 1.3 GHz
2GB DDR3

2.5'' SSD
250GB SATA 3

PCIe switch PLX 8632
32 lane Gen2
12 port

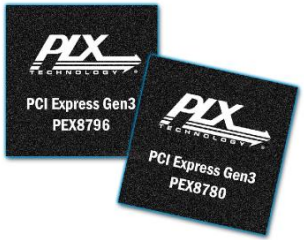

NVIDIA K20
16x PCIe Gen3
1170 GFLOPS
(peak)

Mini-ITX carrier
4x PCIe Gen1
SATA 2.0
1 GbE

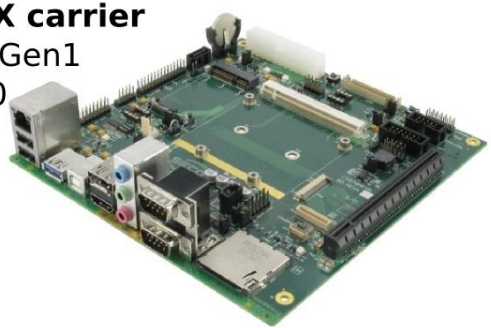

Mellanox ConnectX-3
8x PCIe Gen3 QDR

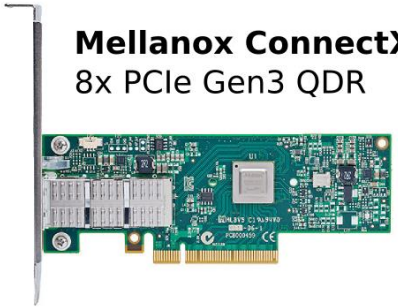

The first server ARM+K20 EK002 2013

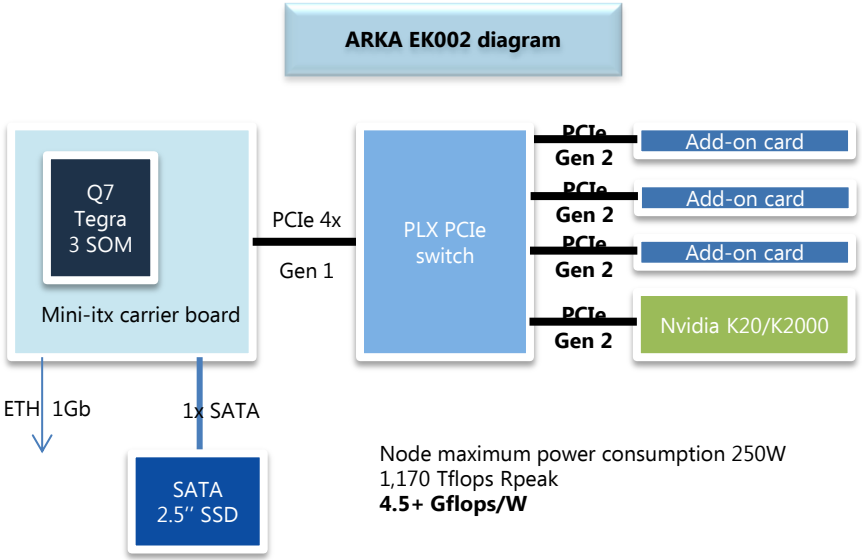

Node maximum power consumption 250W
1,170 Tflops Rpeak
4.5+ Gflops/W

EK002 Efficiency

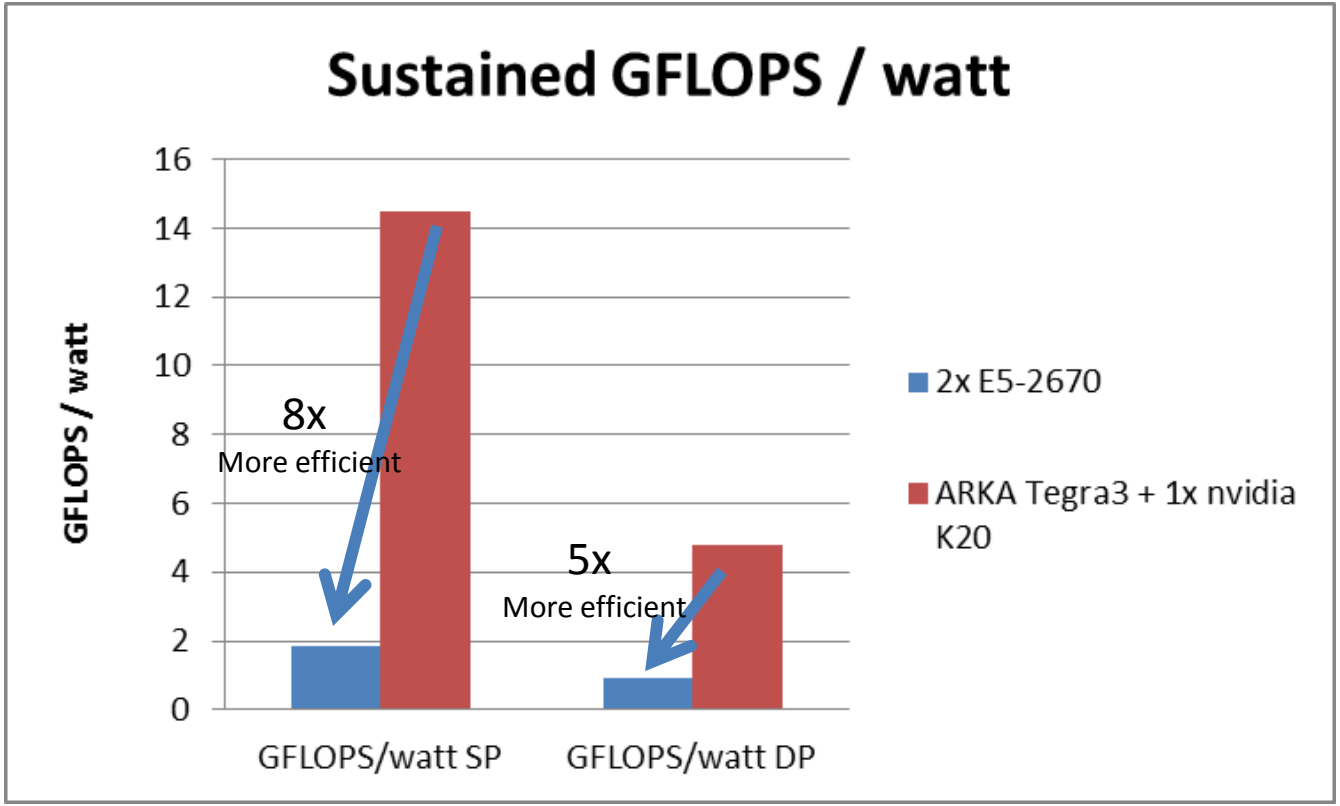

Pedraforca Cluster 2013

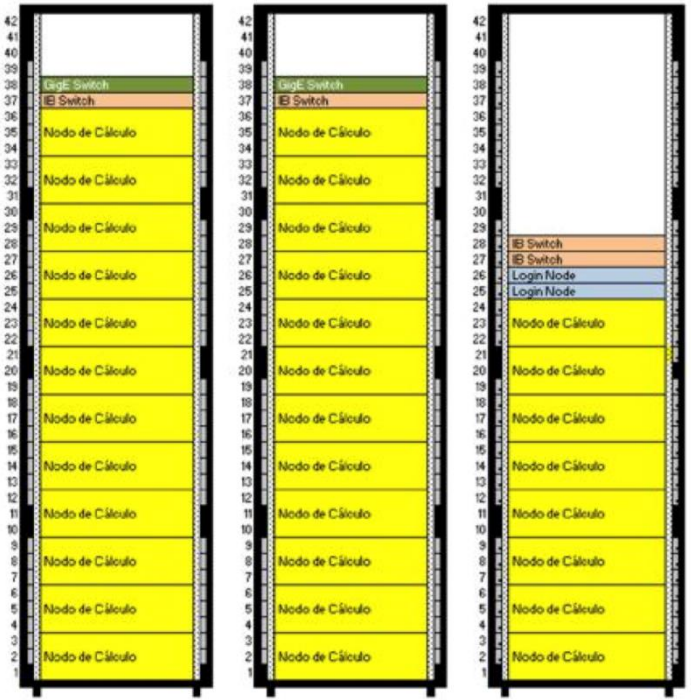

3 × rack

78 compute nodes
2 login nodes
4 36-port InfiniBand switches (MPI)
2 50-port GbE switches (storage)

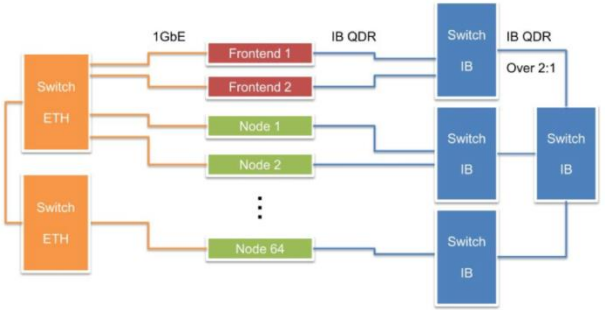

X-Gene: *Highly Integrated Server on a Chip ™*

applied micro

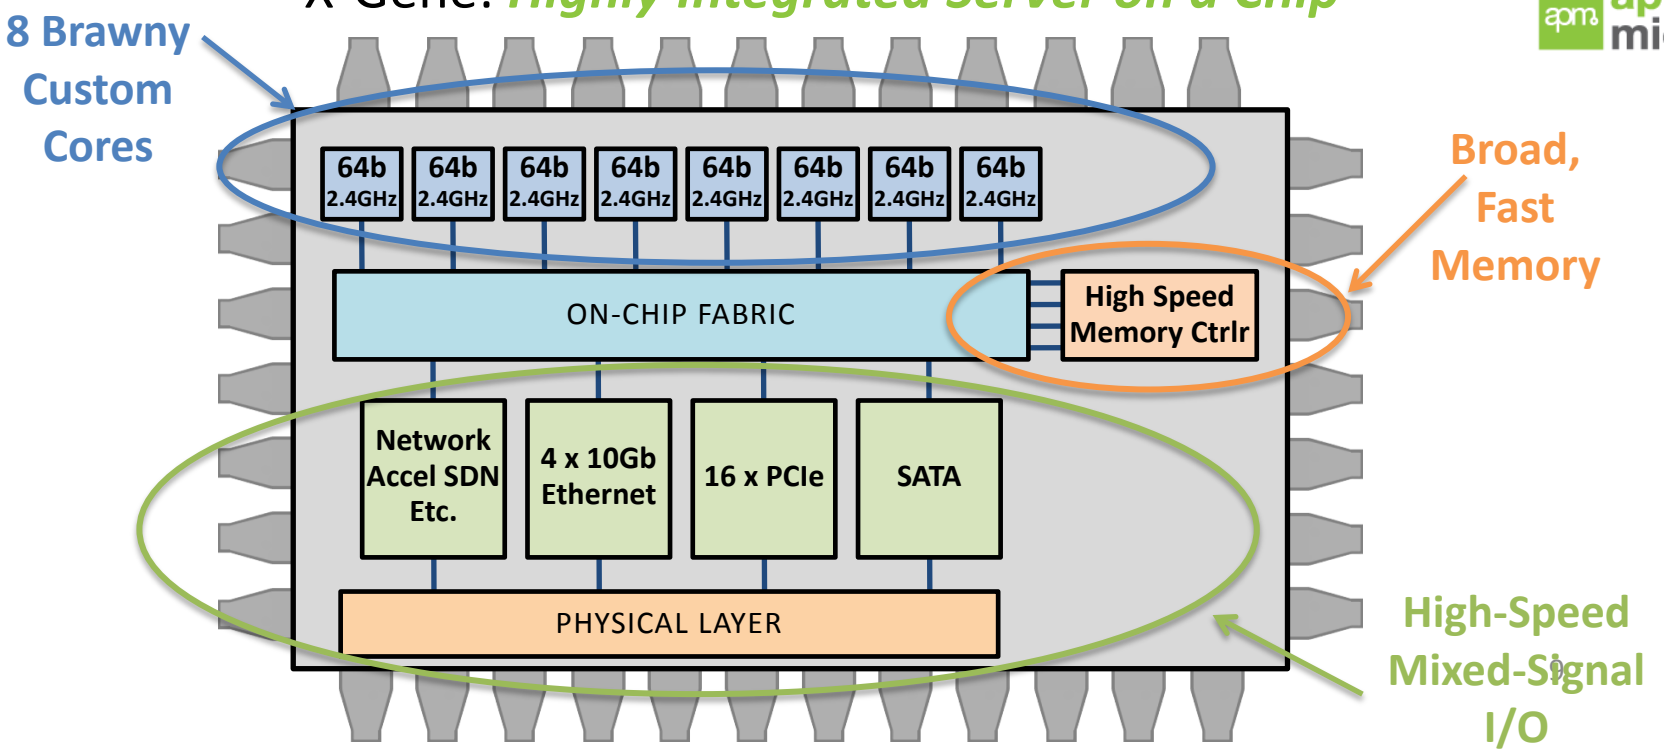

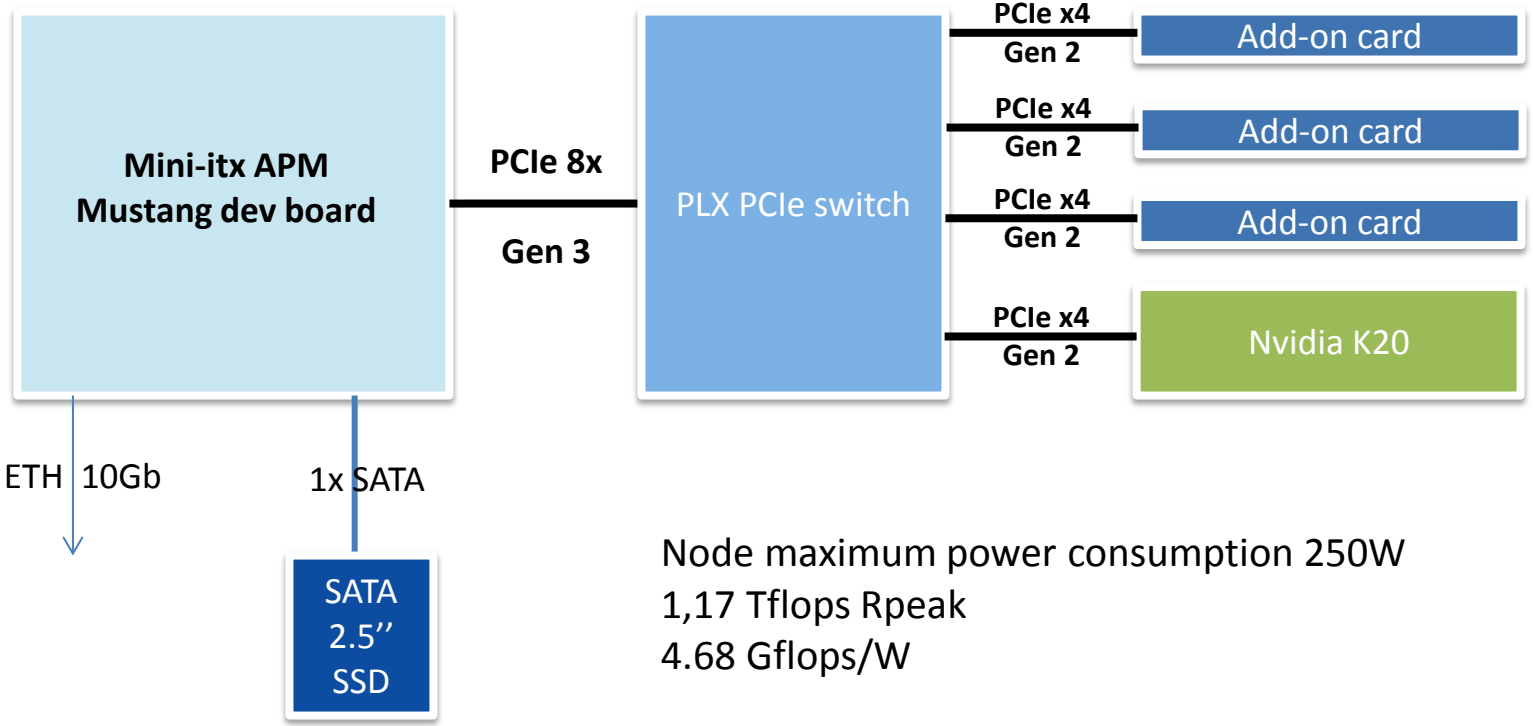

Node maximum power consumption 250W
1,17 Tflops Rpeak
4.68 Gflops/W

```
[root@apm2 bandwidthTest]# ./bandwidthTest
[CUDA Bandwidth Test] - Starting...
Running on...

Device 0: Tesla K20m
Quick Mode

Host to Device Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)  Bandwidth(MB/s)
  33554432                     1430.7

Device to Host Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)  Bandwidth(MB/s)
  33554432                     1639.0

Device to Device Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)  Bandwidth(MB/s)
  33554432                     146841.6
Result = PASS
```

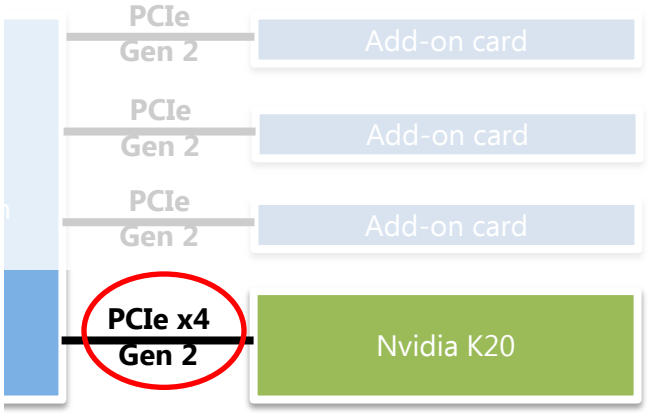

Ibv_rc_pingpong : 19548.12 Mbit/s
3.69 µs

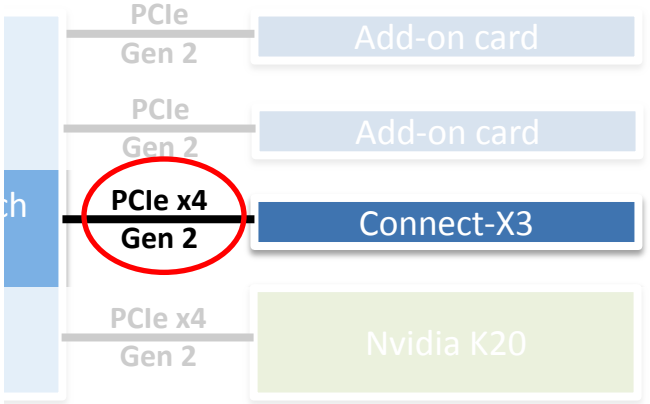

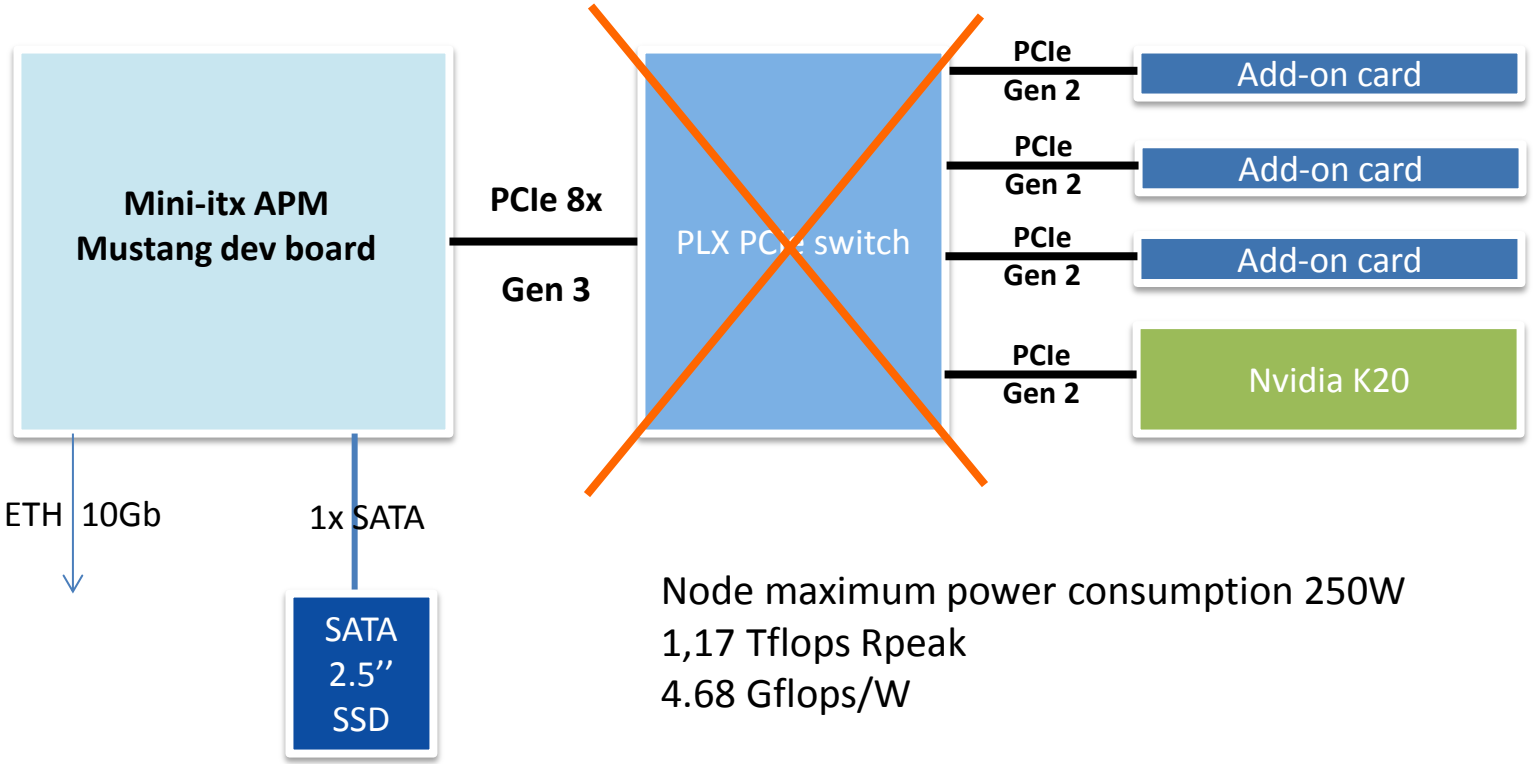

Node maximum power consumption 250W
1,17 Tflops Rpeak
4.68 Gflops/W

FEATURES
Form Factor: 2U
CPU 1x APM X-Gene 8 cores
GPU up to 2x NVIDIA Kepler® K20, K40, K80
Memory Up to 128 GB RAM DDR3
Network 2x 10 GbE SFP+, 1x Infiniband FDR QSFP
Storage 4x SATA 3.0
Expansion slots 2x PCI-E 3.0 x8 (in x16)

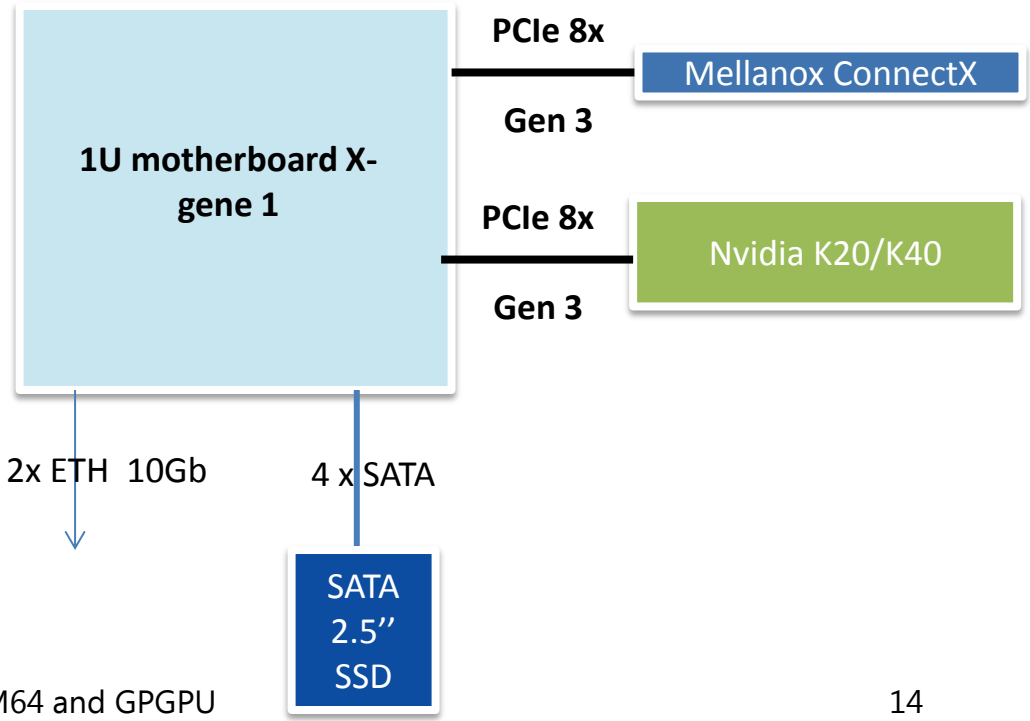

OS:
Ubuntu derivative for ARM 14.04 LTS

DEVELOPMENT TOOL:
NVIDA CUDA 6.5 (compilers, libraries, SDK)
MPI 2.0 Libraries
GNU compilers
Scalable Heterogeneous Computing (SHOC) Benchmark Suite

CLUSTER, MONITORING, MANAGEMENT TOOLS (optional):
Resource/Queue manager
Monitoring tools for cluster
Parallel shell for cluster-wide commands
Bare-metal restore

The ARKA RK003 Server (3)

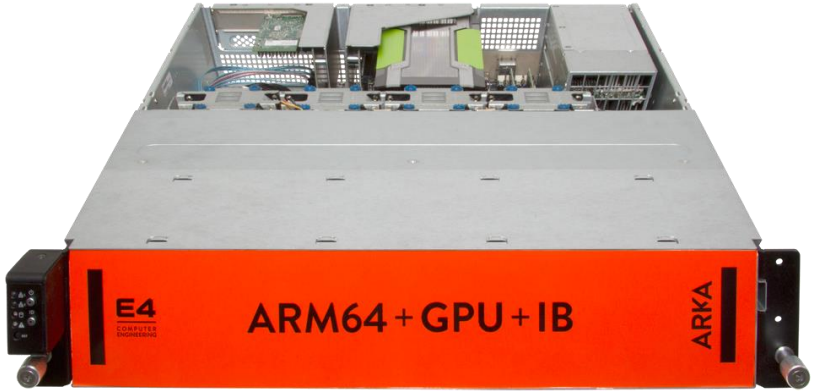

```
Device 0: 'Tesla K20m'
Device selection not specified: defaulting to device #0.
Using size class: 1

--- Starting Benchmarks ---
Running benchmark BusSpeedDownload
    result for bspeed_download:                  3.1592 GB/sec
Running benchmark BusSpeedReadback
    result for bspeed_readback:                  3.3575 GB/sec
Running benchmark MaxFlops
    result for maxspflops:                    3095.8700 GFLOPS
    result for maxdpflops:                    1165.9700 GFLOPS
Running benchmark DeviceMemory
    result for gmem_readbw:                    147.9020 GB/s
    result for gmem_writebw:                   139.9250 GB/s
```

```
Device 0: 'Tesla K20m'
Device selection not specified: defaulting to device #0.
Using size class: 1

--- Starting Benchmarks ---
Running benchmark BusSpeedDownload
    result for bspeed_download:                  3.1592 GB/sec
Running benchmark BusSpeedReadback
    result for bspeed_readback:                  3.3575 GB/sec
Running benchmark MaxFlops
    result for maxspflops:                    3095.8700 GFLOPS
    result for maxdpflops:                    1165.9700 GFLOPS
Running benchmark DeviceMemory
    result for gmem_readbw:                    147.9020 GB/s
    result for gmem_writebw:                   139.9250 GB/s
```

Ibv_rc_pingpong:

Latency
1.57 usec

Bandwidth
39.2Gb/s

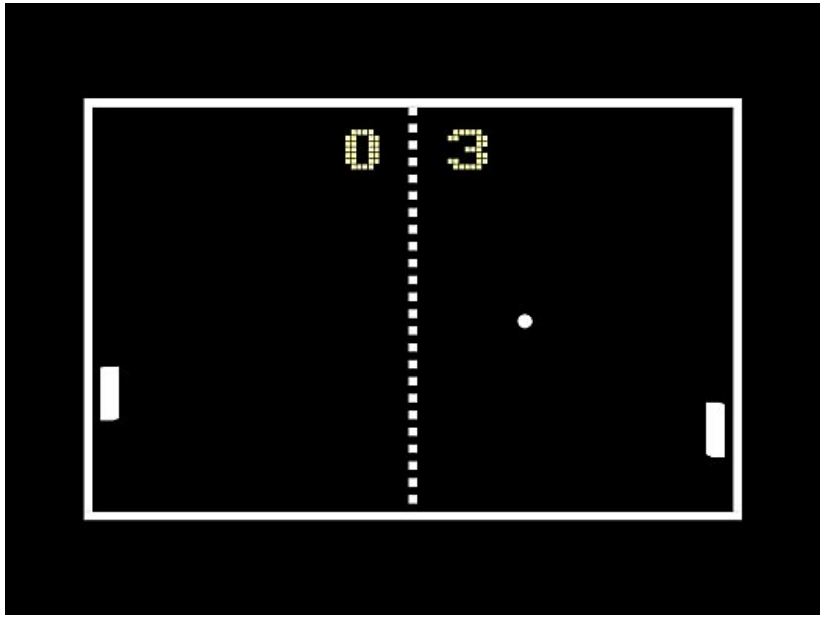

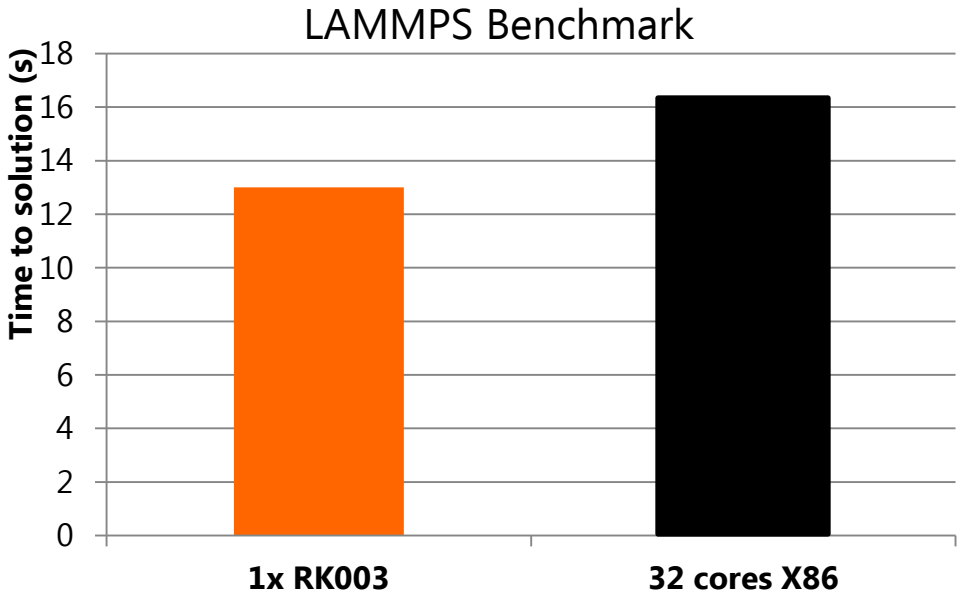

SETUP

COMPILER

-Intel compilers on E5-2650v2

-MKL

-GCC on ARM

-FFTW 3.2.2

HARDWARE

-Dual E5-2650v2, Infiniband QDR

-ARM X-gene1+ 1 GPU Nvidia K20m

INPUT FILES

-WdV 1M particle

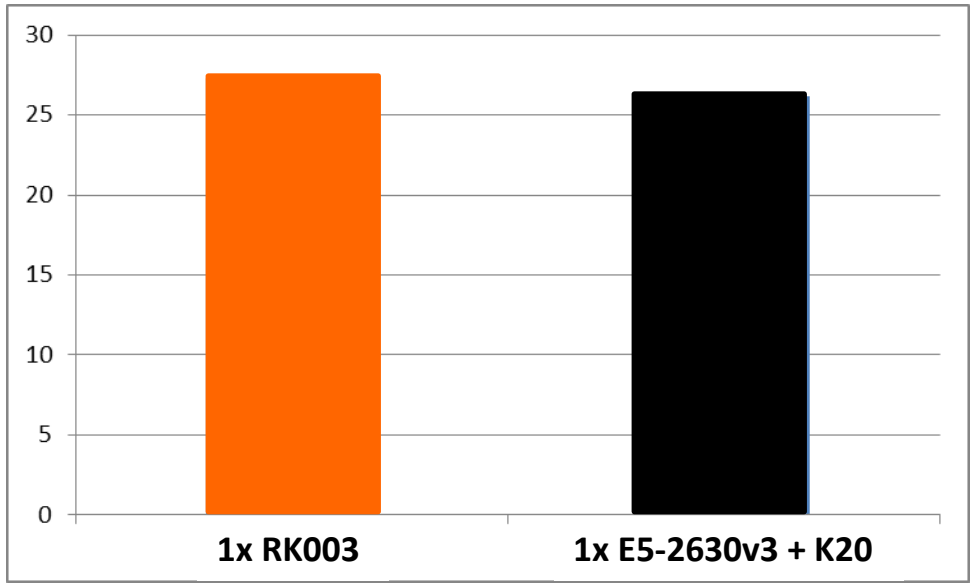

SETUP

COMPILER GNU

Cuda 6.5

Cublas and cufft

HARDWARE

-Dual E5-2630v3 + 1 GPU Nvidia K20m

-ARM X-gene1+ 1 GPU Nvidia K20m

INPUT FILES

-WdV 1M particle

- 1000 steps

Power Consumption

SETUP
COMPILER GNU
Cuda 6.5
Cublas and cufft

HARDWARE
-Dual E5-2630v3 + 1 GPU Nvidia K20m
-ARM X-gene1+ 1 GPU Nvidia K20m

GPU performances:

RK003 – avg load power consumption: 229w
MaxFlops: 1165 GFLOPS
Benchmark: 5,1 GFLOP/w
power consumption in idle: 92 W

Xeon E5 - avg load power consumption: 410w
MaxFlops: 1165 GFLOPS
Benchmark: 2,8 GFLOP/w
power consumption in idle: 151

THUNDERX

- Up to 48 custom ARMv8 cores @ 2.5GHz
- Single & Dual socket configs
- Integrated PCIe, SATA, 10/40GbE , Ethernet Fabric
- VirtSOC™: Low latency end to end virtualization from VM to virtual port
- Four Workload Optimized Families
  - Compute, Networking, Storage and Security
- Shipping today in Single and Dual Socket Configurations

FEATURES
Form Factor: 2U
CPU 1x CAVIUM Thunder-X 48 cores
GPU up to 1x NVIDIA Kepler® K20, K40, K80
Memory Up to 128 GB RAM DDR3
Network 2x 10 GbE SFP+, 1x IB FDR QSFP, 1x 40 GbE QSFP
Storage 4x SATA 3.0
Expansion slots 1x PCI-E 3.0 x8 (in x16), 1x PCI-E 3.0 x8

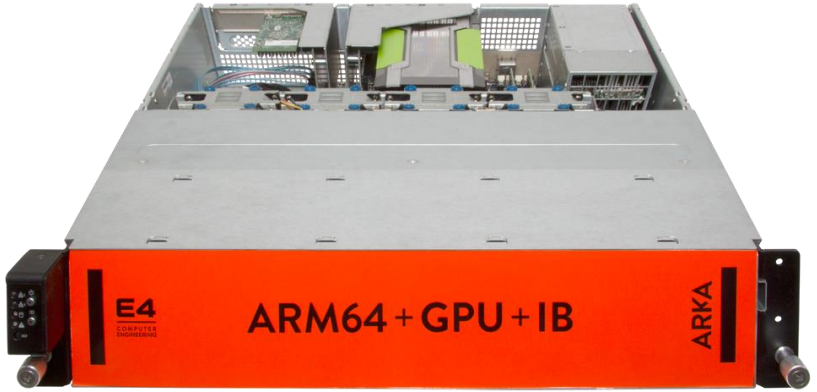

BOOTH 430

Cosimo Gainfreda
Simone Tinti
Wissam Abu-Ahmmad
Francesca Tartaglione (now at OIST)
Marco Ciacala
Alessandro De Filippo

Filippo Mantovavi

E4

COMPUTER
ENGINEERING

BOOTH 430