



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



**UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH**

Departament d'Arquitectura de Computadors

Supercomputing with Commodity CPUs: Are Mobile SoCs Ready for HPC?

Nikola Rajovic, Paul M. Carpenter, Isaac Gelado, Nikola Puzovic,
Alex Ramirez, Mateo Valero


SC'13, November 19th 2013,
Denver, CO, USA

Commodity components drive HPC



Top500 1993, 1st edition:


Cray	vector,	41%
MasPar	SIMD,	11%
Convex/HP	vector,	5%



“ Microprocessors replaced Vector/SIMD supercomputers


- They were not faster
- They were cheaper

Vectors vs. microprocessors: transition



- ❧ In 1995, when microprocessors overtook vector/SIMD
 - Microprocessors ~10 times slower than one vector CPU (FP)
 - Performance gap closing fast
- ❧ SIMD vs. Message passing paradigms
- ❧ **Advantage: commodity volume economics**

The next step in the commodity chain



« 20M cores in Jun'13 Top500

Sold in 2012

« <10M servers


« >350M PC's

« >100M tablets

« >700M smartphones

– > 210M smartphones (1Q 2013)

History may be about to repeat itself




« In 2013, Mobile SoCs are slower
– But performance gap seems to close


« They are significantly cheaper ... in high volume

Mobile SoC vs Server – side by side

Performance



Cost



1. Leaked Tegra3 price from the Nexus 7 Bill of Materials

2. Non-discounted List Price for the 8-core Intel E5-2670, SandyBridge

Outline

- « Motivation
- « Mobile SoCs evaluation
- « Mobile SoCs in cluster environment
- « Challenges in HPC with Mobile SoCs

Platforms under study: CPU and Memory



NVIDIA Tegra 2

2 x ARM Cortex-A9 @ 1GHz
1 x 32-bit DDR2-333 channel
32KB L1 + 1MB L2



NVIDIA Tegra 3

4 x ARM Cortex-A9 @ 1.3GHz
1 x 32-bit DDR3-750 channel
32KB L1 + 1MB L2



Samsung Exynos 5 Dual


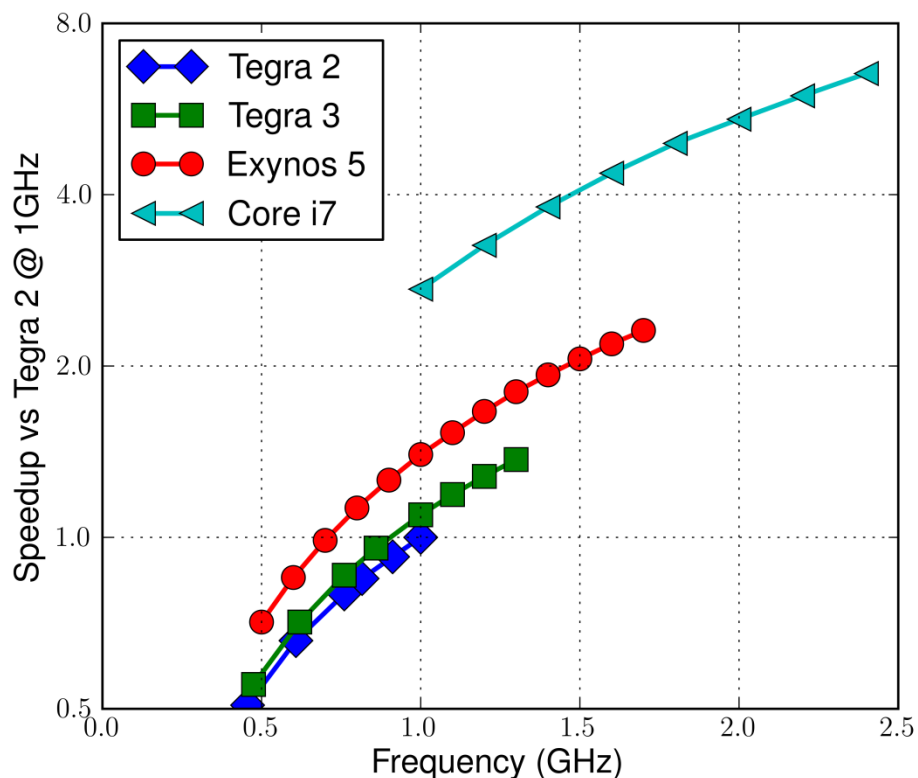
2 x ARM Cortex-A15 @ 1.7GHz
2 x 32-bit DDR3-800 channels
32KB L1 + 1MB L2



Intel Core i7-2760QM



4 x Intel SandyBridge @ 2.4GHz
2 x 64-bit DDR3-800 channels
32KB L1 + 1MB L2 + 6MB L3

Single core performance and energy



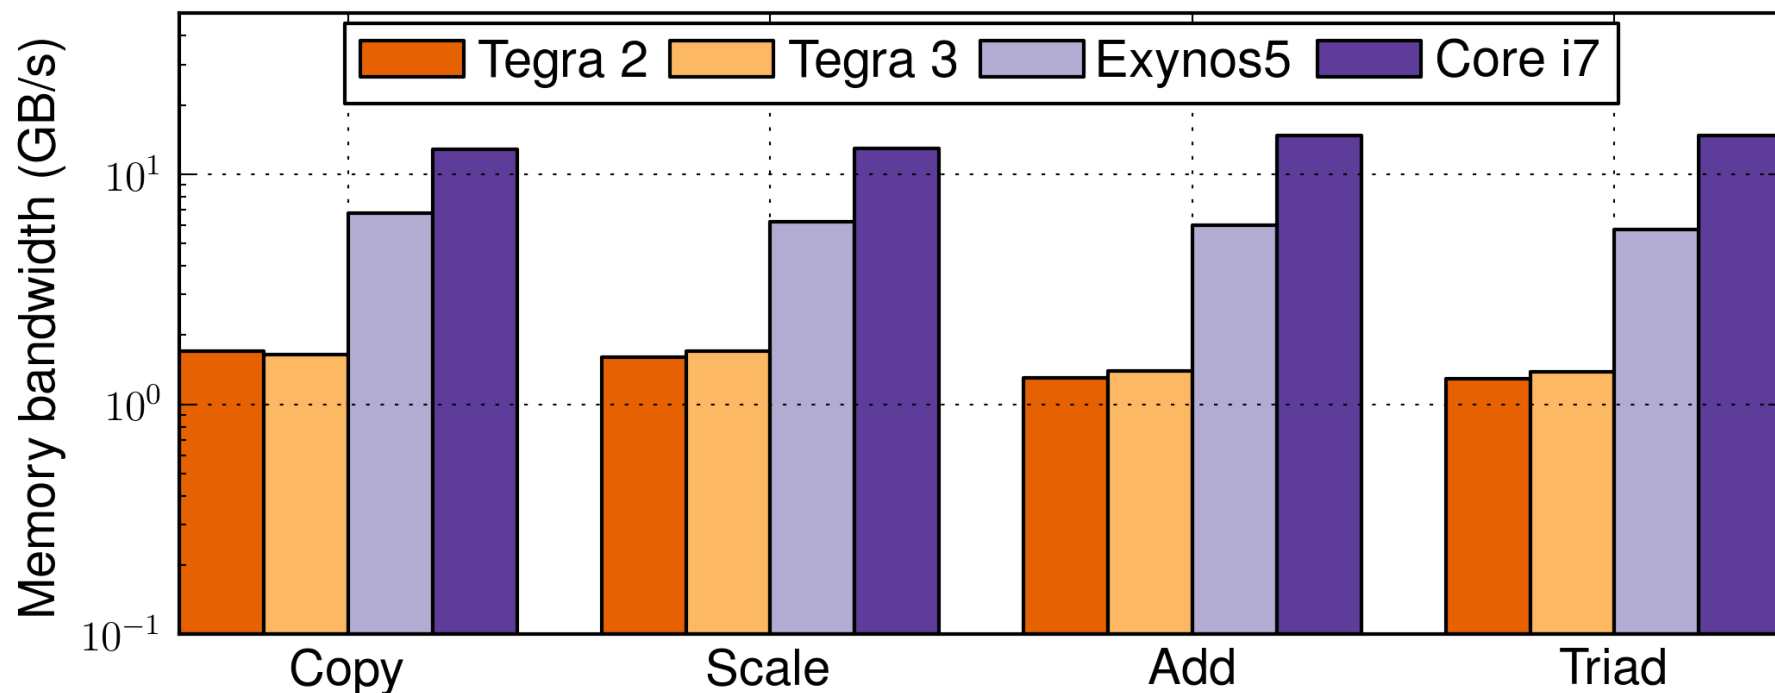
- || Cortex-A9 in Tegra3 is 1.4x faster than Tegra2 (higher clock frequency)
- || Cortex-A15 in Exynos5 is 1.7x faster than Cortex-A9 in Tegra3
 - Higher clock frequency, higher memory bandwidth, and better core microarchitecture
- || Core i7 is ~3x faster than Cortex-A15 in Exynos5 at maximum frequency
 - 2x faster at the same frequency
- || Mobile SoC platforms as efficient as Core i7 platform at their highest operating points

Multicore performance and energy



- ⌘ Tegra3 platform as fast as Exynos5 platform, a bit more energy efficient
 - 4-core Cortex-A9 vs. 2-core Cortex-A15
- ⌘ Corei7 is 6x faster than Exynos5 at maximum frequency
- ⌘ Tegra3 and Exynos5 as efficient as Corei7 at the same frequency

Memory bandwidth (STREAM)



« Exynos 5 improves dramatically over Tegra (4.5x)

- Dual-channel DDR3
- ARM Cortex-A15 sustains more in-flight cache misses

« Corei7 provides ~2x more memory bandwidth than Exynos5

Outline

- « Motivation
- « Mobile SoCs evaluation
- « Mobile SoCs in cluster environment
- « Challenges in HPC with Mobile SoCs

Tibidabo: The first ARM HPC multicore cluster



Q7 Tegra 2

2 x Cortex-A9 @ 1GHz
2 GFLOPS
5 Watts (?)
0.4 GFLOPS / W



Q7 carrier board

2 x Cortex-A9
2 GFLOPS
1 GbE + 100 MbE
7 Watts
0.3 GFLOPS / W



1U Rackable blade

8 nodes
16 GFLOPS
65 Watts
0.25 GFLOPS / W




2 Racks

32 blade containers
256 nodes
512 cores
10x 48-port 1GbE switch
8x 48-port 100 MbE switch

512 GFLOPS
3.4 Kwatt
0.15 GFLOPS / W

- Cluster of developer kits, not a custom design
 - Proof of concept and insights
- Enable software stack and applications tuning

Applications scalability



Application	Description
HPL	High Performance LINPACK
PEPC	Tree code for N-body problem
HYDRO	2D Eulerian hydrodynamics
GROMACS	Molecular dynamics
SPECfem3D	3D seismic wave propagation (spectral element method)

Weak scalability test with HPL

- 97 GFLOPS on 96 nodes (51% efficiency, linear scaling)
- 0.12 GFLOPS/W


Strong scalability with the rest

- Very small input set, 1GB DRAM per node

Interconnect evaluation: SoCs under study

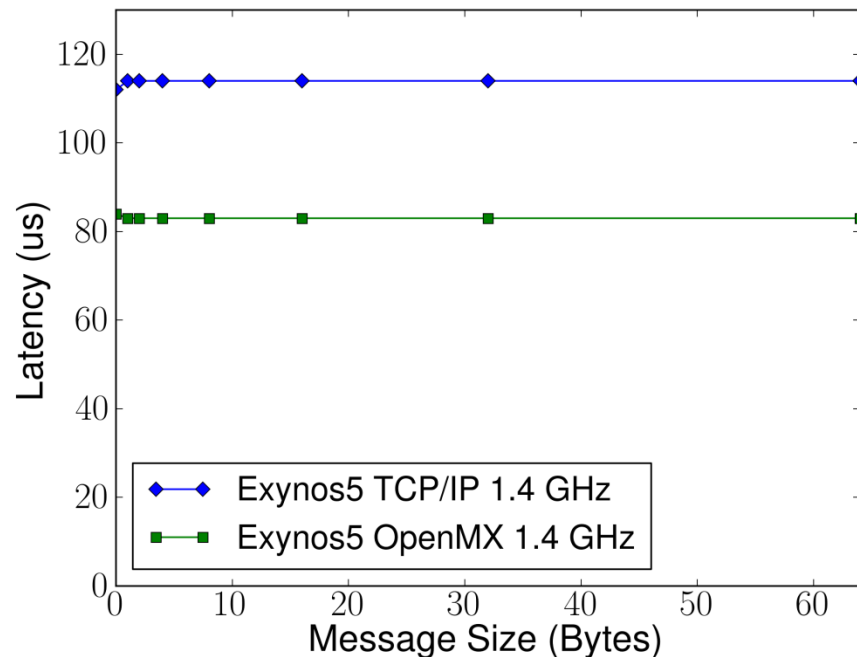



NVIDIA Tegra 2
1 GbE (on PCIe)
100 Mbit (on USB 2.0)



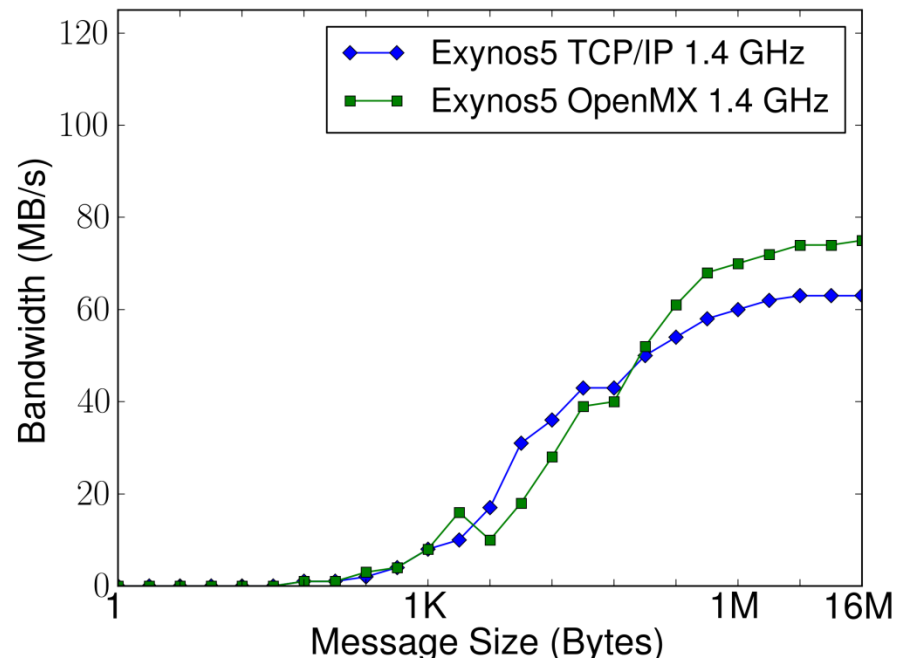
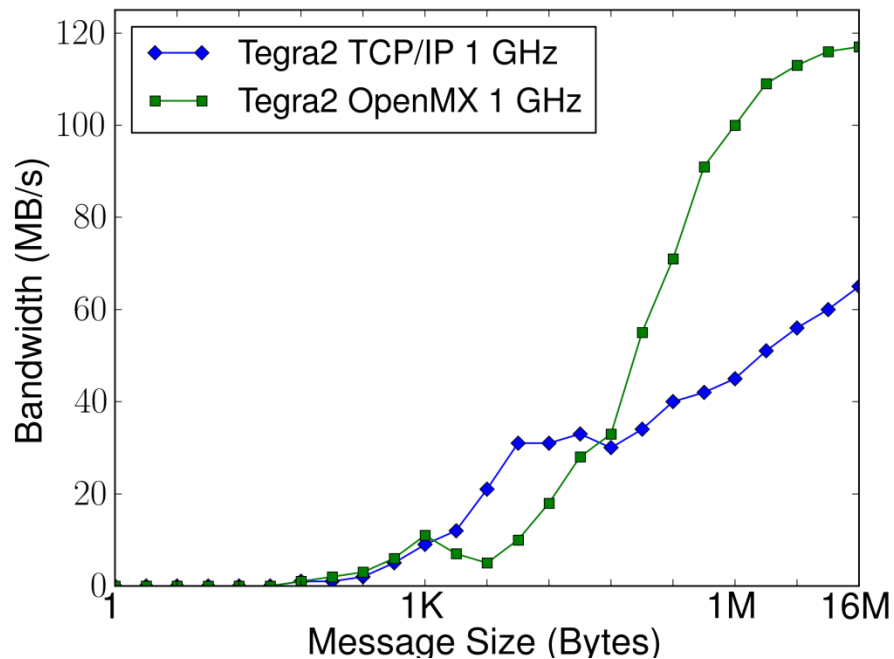
Samsung Exynos 5 Dual
1 GbE (on USB3.0)
100 Mbit (on USB 2.0)

Interconnect evaluation: latency



- ❧ TCP/IP adds significant CPU overhead
- ❧ OpenMX driver interfaces “directly” to the Ethernet NIC
- ❧ USB in Exynos5 adds extra latency on top of network stack

Interconnect evaluation: bandwidth



- ❧ TCP/IP overhead prevents Tegra2 from achieving full bandwidth
 - OpenMX does achieve peak bandwidth
- ❧ USB overheads prevent Exynos 5 from achieving full bandwidth, even with OpenMX

Outline

- « Motivation
- « Mobile SoCs evaluation
- « Mobile SoCs in cluster environment
- « Challenges in HPC with Mobile SoCs

Challenges in HPC with Mobile SoCs

Our experience: Hardware

- ❧ Mobile SoCs still do not target HPC
 - Platforms are available only as developer kits
- ❧ Developer boards are not designed for continued high-performance operation
 - Usually, no cooling infrastructure
 - Very hard for packaging (also low density)
- ❧ PCIe not reliable in Tegra2 and Tegra3
 - Could fail to initialize during boot
 - Stop responding during heavy workloads

Challenges in HPC with Mobile SoCs

Our experience: System software and applications

- ❧ Software ecosystem is built for high-compatibility across diverse ARM-powered Mobile SoCs
 - Entire Linux distributions compiled with lowest level of optimizations
 - Softfp ABI is still mainstream
- ❧ Linux distributions and kernels are not ‘turn key’ solutions
 - Rely on vendors to provide with correct set of kernels and Linux distribution images

Mobile SoC limitations for HPC

- ❧ 32-bit memory controller
 - Even though ARM Cortex-A15 offers 40-bit address space
- ❧ No ECC protection in memory
 - Limiting factor for scalability after certain number of nodes
- ❧ No standard server I/O interfaces
 - Provide USB 3.0, SATA and (minimal) PCIe
- ❧ No protocol offload engines
 - e.g. TCP/IP runs on CPU
- ❧ Low grade thermal package
- ❧ **These are only design decisions, not really unsolvable problems**
 - ARM server SoCs don't have any of these restrictions

Are Mobile SoCs ready for HPC?

- ⌘ Mobile SoCs enjoy aggressive roadmaps and fast innovations
 - driven by commodity components business dynamics and market
- ⌘ They have to address their limitations before entering HPC
 - ECC, interconnect, 32-bit address space, low-grade thermal package ...
- ⌘ Mobile SoCs may introduce a new class of supercomputers:
 - Faster, cheaper and more energy efficient
 - If vendors decide to include a minimum set of required features