TECHNISCHE UNIVERSITÄT DRESDEN

Memory Performance at Reduced CPU Clock Speeds: An Analysis of Current x86_64 Processors

HotPower'12

Robert Schöne (robert.schoene@tu-dresden.de)

Daniel Hackenberg (daniel.hackenberg@tu-dresden.de)

Daniel Molka (daniel.molka@tu-dresden.de)

ZIH
Center for Information Services & High Performance Computing

Outline

- Motivation
- Test Systems
- Benchmarks
- Results
- Conclusion

Motivation

- Energy consumption of data centers and HPC computers is increasing
- Power saving techniques are implemented into server hardware
- DVFS one of the major techniques – beside clock and power gating – to save power
- Software makes use of DVFS to increase energy efficiency
- General approach:

  "lower the core frequency during a memory phase and raise it during a computing phase" [Liv12]

Test Systems

Intel				AMD		
Xeon X5560	*Xeon X5670*	*Core i7-2600K*	*Xeon E5-2670*	Opteron 2435	Opteron 6168	*Opteron 6274*
Nehalem-EP	*Westmere-EP*	*Sandy Bridge-HE*	*Sandy Bridge-EP*	Istanbul	Magny-Cours	*Interlagos*
2x4	*2x6*	*4*	*2x8*	2x6	2x2x6	*4x2x8*
2.8 GHz	*2.933 GHz*	*3.4 GHz*	*2.6 GHz*	2.6 GHz	1.9 GHz	*2.2 GHz*
3.2 GHz	*3.333 GHz*	*3.8 GHz*	*3.3 GHz*			*2.5 GHz*
PC3-10600R	*PC3L-10600R*	*PC3-10600*	*PC3-12800R*	PC2-5300R	PC3-10600R	*PC3-12800*

Uncore Design Nehalem-EP/Westmere-EP

- Central crossbar (Global Queue)
- Cores connect to Memory, L3 and QPI via Global Queue
- Two frequency domains:
  - Core frequency
  - Uncore Frequency (2.66 GHz at Westmere-EP)

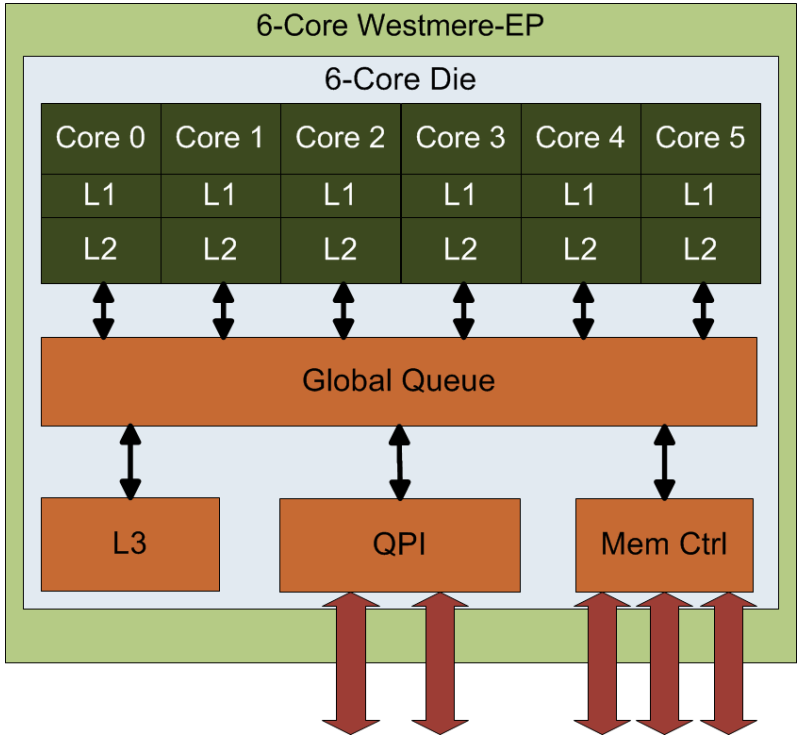

Based on [Int10]

Uncore Design Sandy Bridge-HE

- Ring setup (32 byte data)
- Cores connect to DRAM via ring bus
- One frequency domain

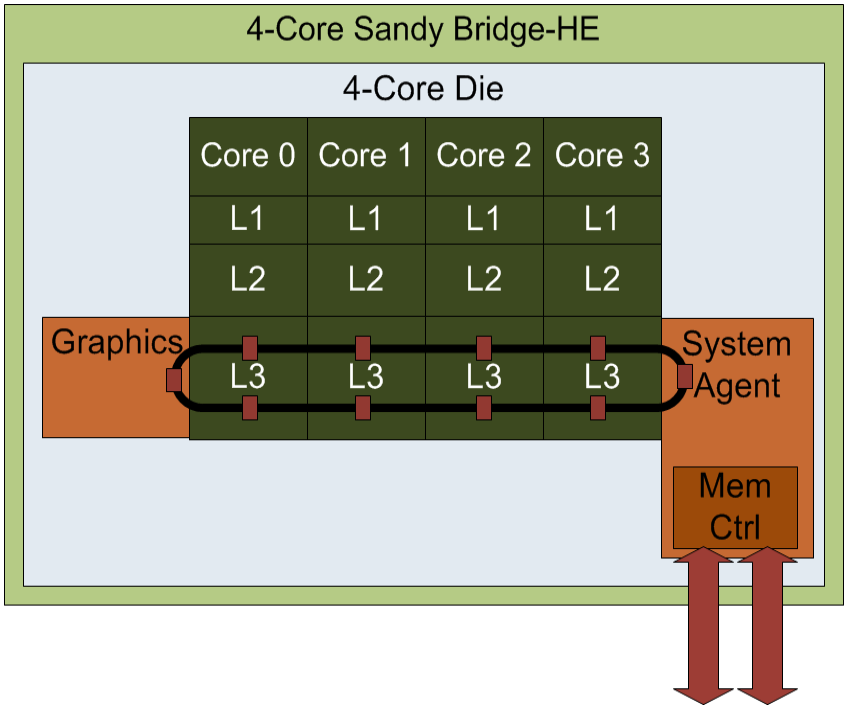

Uncore Design Sandy Bridge-EP

- Ring setup (2x 32 byte data)
- Cores connect to (most) L3, QPI and DRAM via ring bus
- One frequency domain

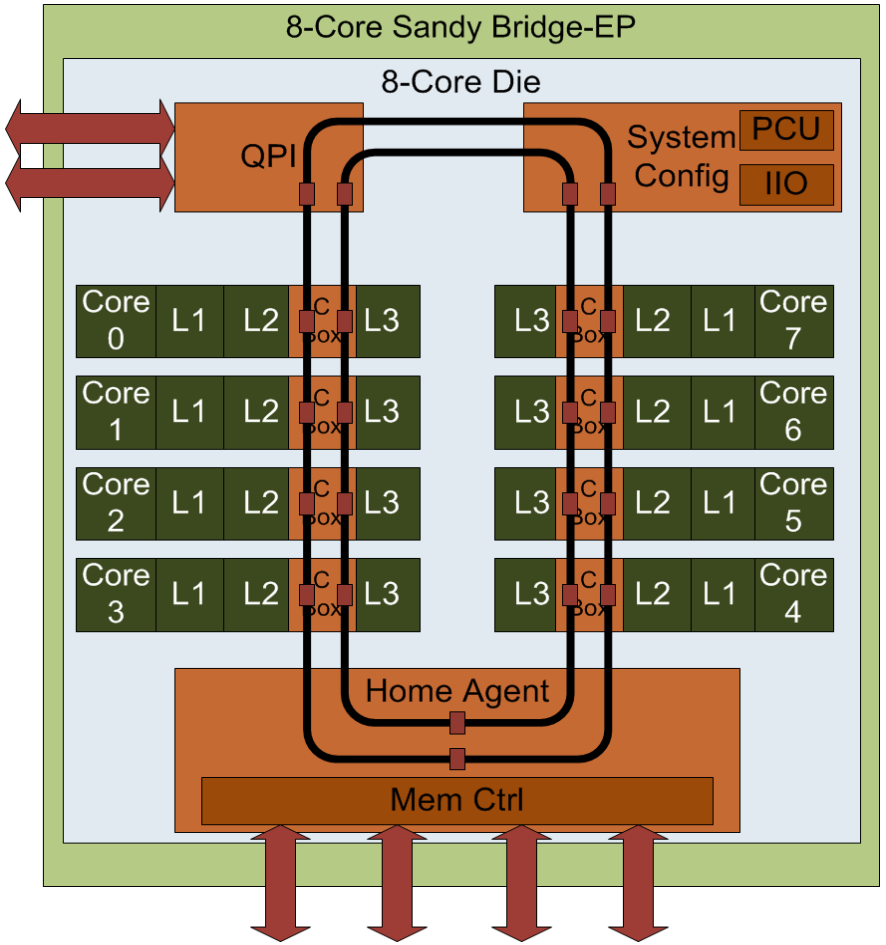

Based on [Int12]

Northbridge Design AMD Interlagos

- Fairly complex design
- Central crossbars (SRI + XBAR)
- Cores connect to L3 via SRI
- Cores connect to HT and DRAM via SRI+XBAR
- Multiple frequency domains:
  - Module frequencies
  - Northbridge frequency (2 GHz at Interlagos)

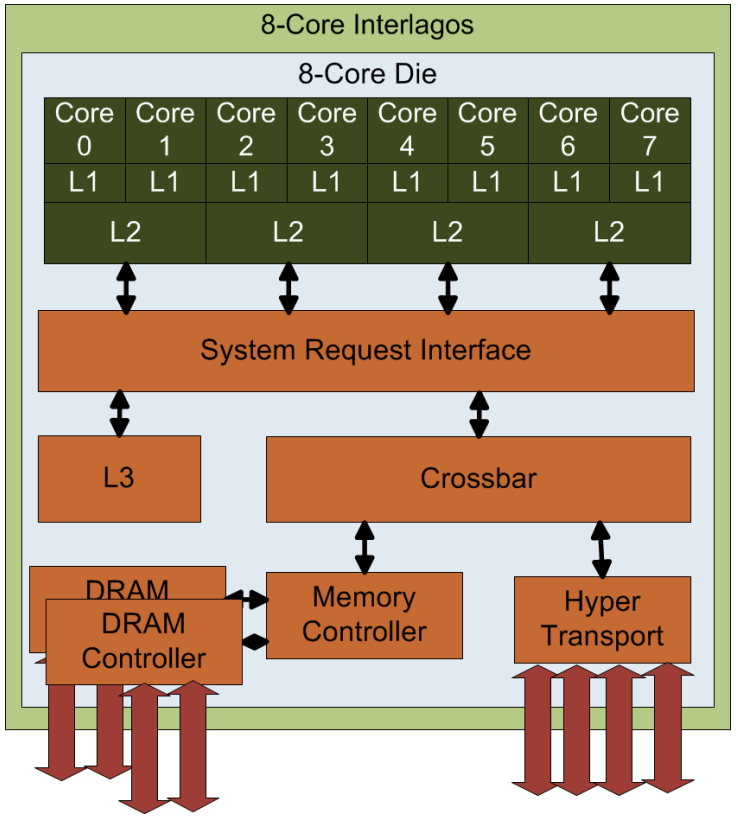

Based on [AMD12]

Microbenchmarks for 64 Bit x86 systems [Hac09]

- Memory latency and bandwidth measurements
  - Well directed placement of data in any cache or memory location
  - Coherency state control
- Implementation
  - pthreads with affinity control
  - Assembler implementation of measurement routines
  - Time measurements using Time Stamp Counter(rdtsc)
  - NUMA aware allocation

http://www.benchit.org/wiki/index.php/X86membench

Bandwidth to DRAM at Full Concurrency

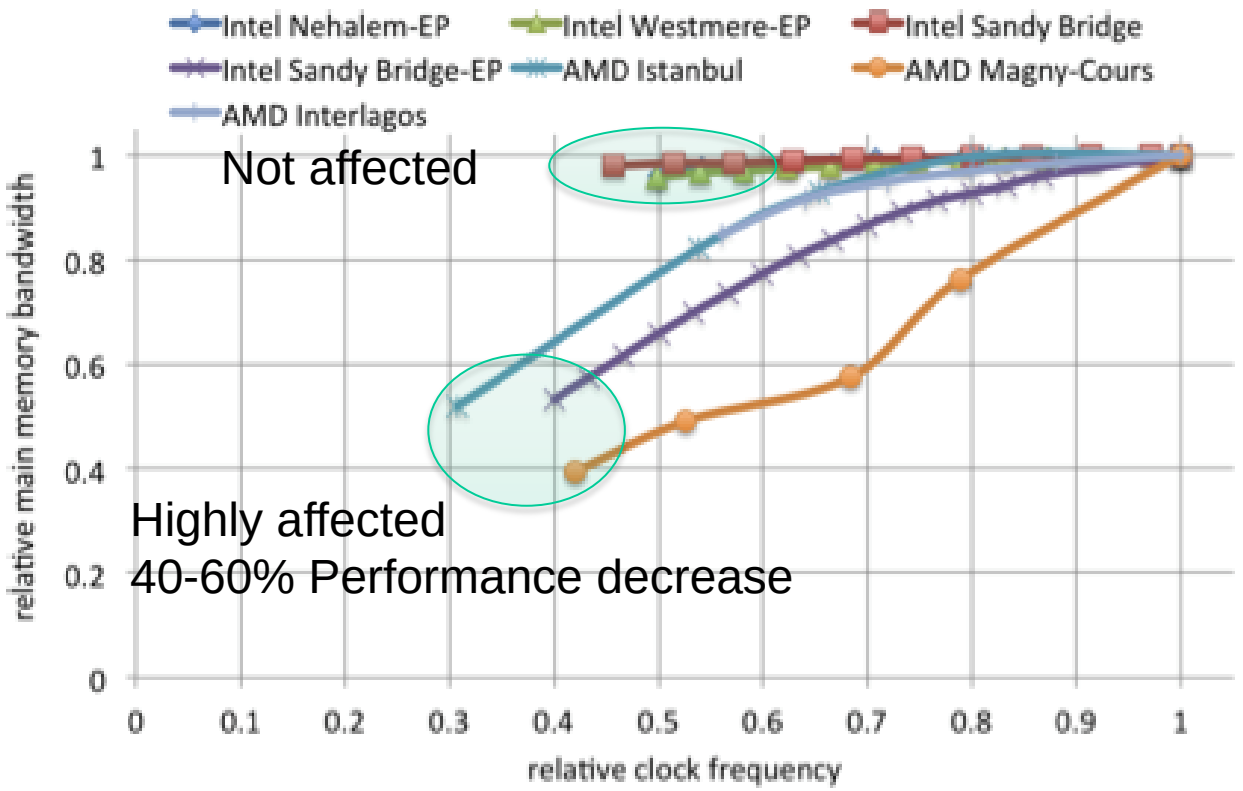

- DRAM performance at reduced clock speed
  - Affected for Intel Sandy Bridge-EP and all AMD processors

Bandwidth to DRAM – Westmere EP

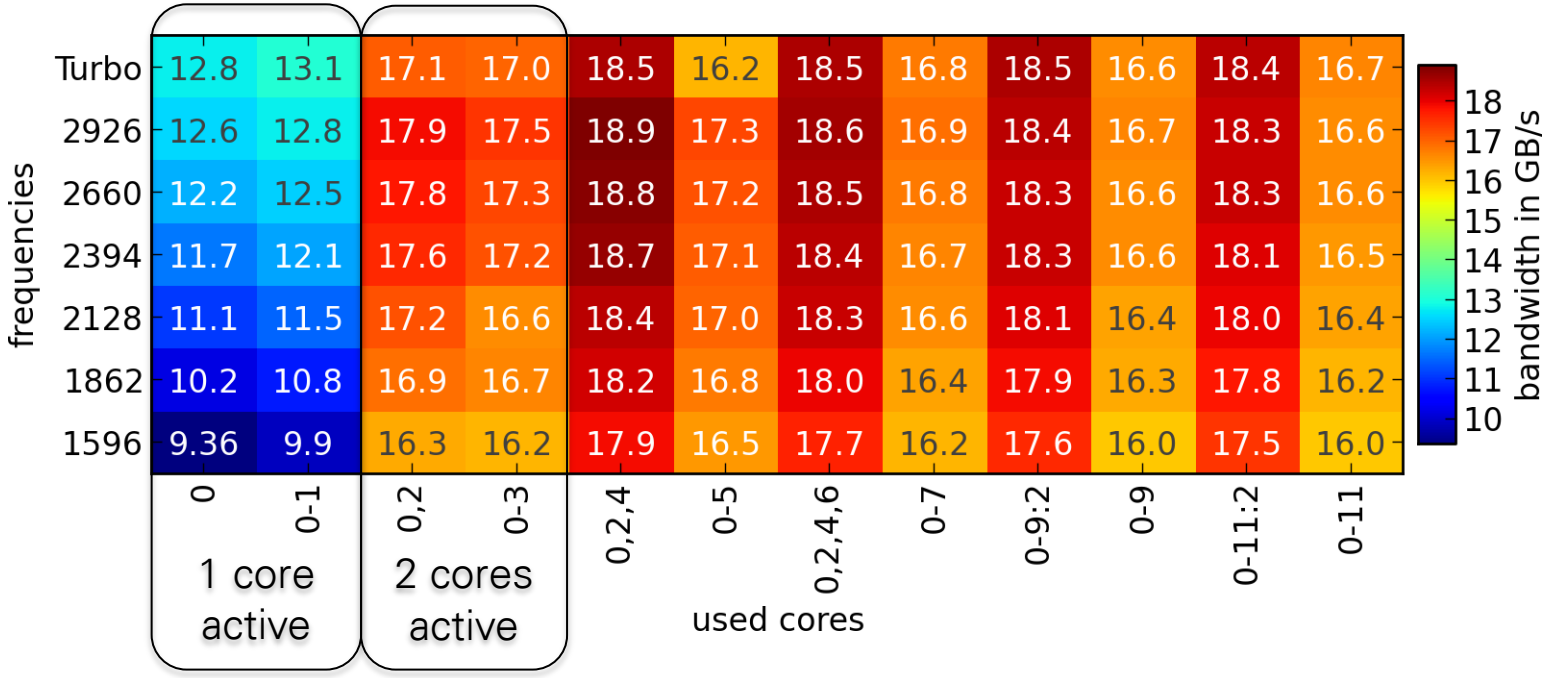

- Only minor DRAM-performance losses when using DVFS techniques for two or more active cores
- Decrease number of cores accessing DRAM (reduce concurrency)

Bandwidth to DRAM – Sandy Bridge-HE

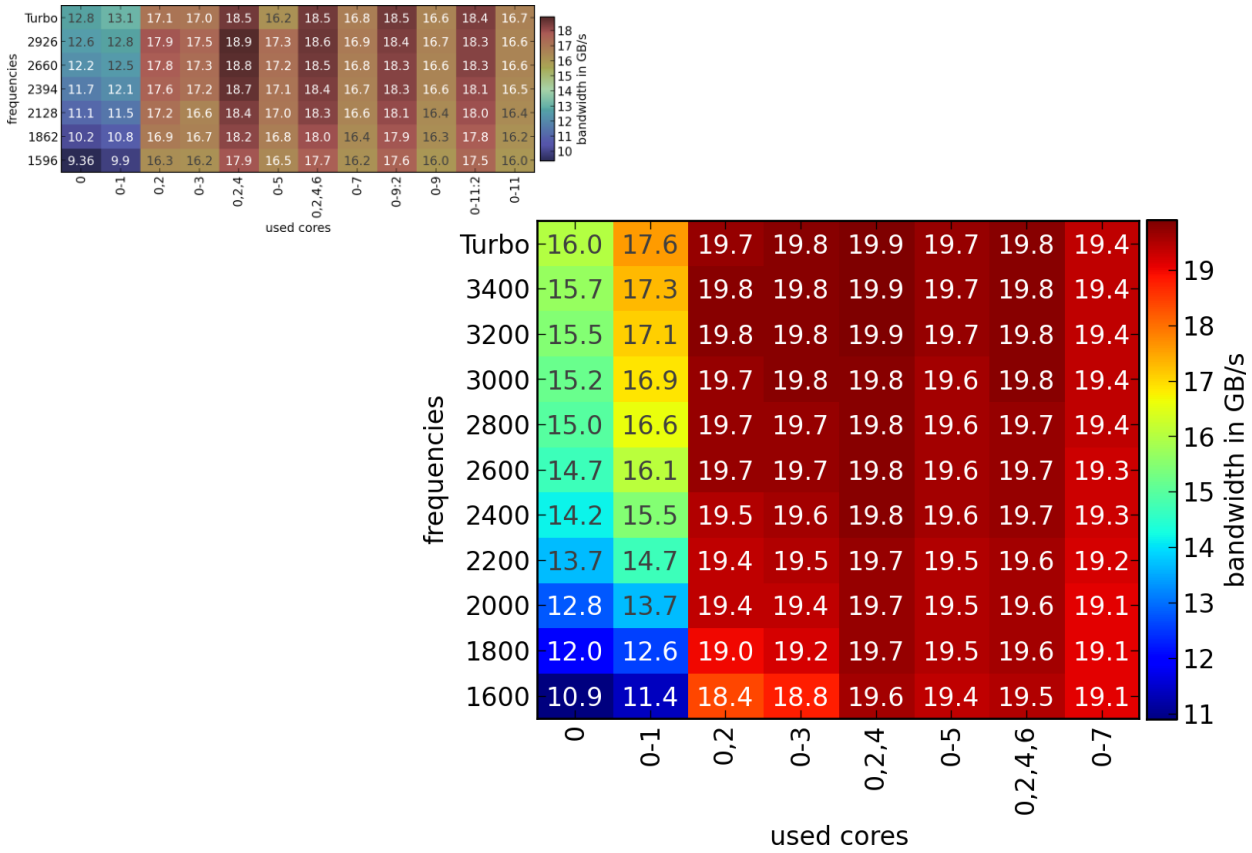

- No DRAM-performance losses when using DVFS techniques for two or more active cores
- Concurrency can be reduced additionally

Bandwidth to DRAM – Sandy Bridge EP

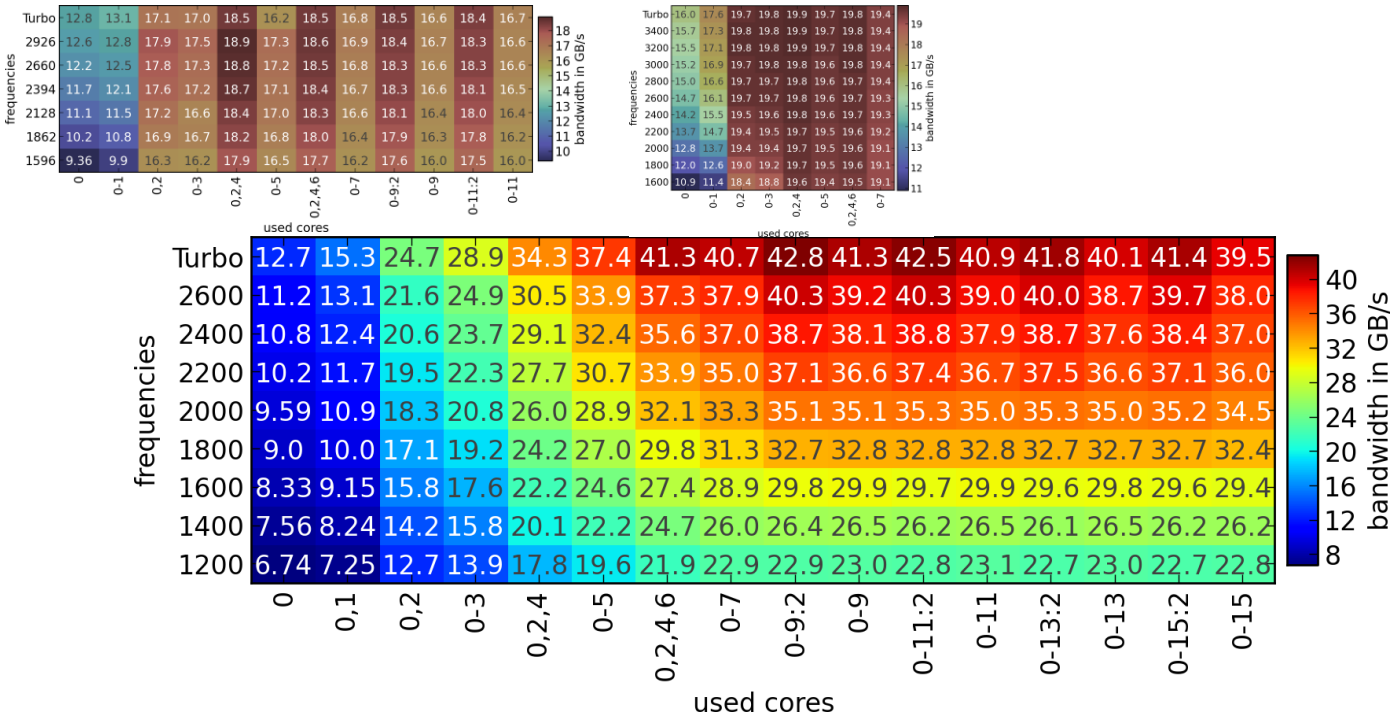

- Frequency scaling worsens performance
- Despite same architecture like Sandy Bridge-HE, totally different behavior
- Reducing concurrency as alternative

Bandwidth to DRAM – Interlagos

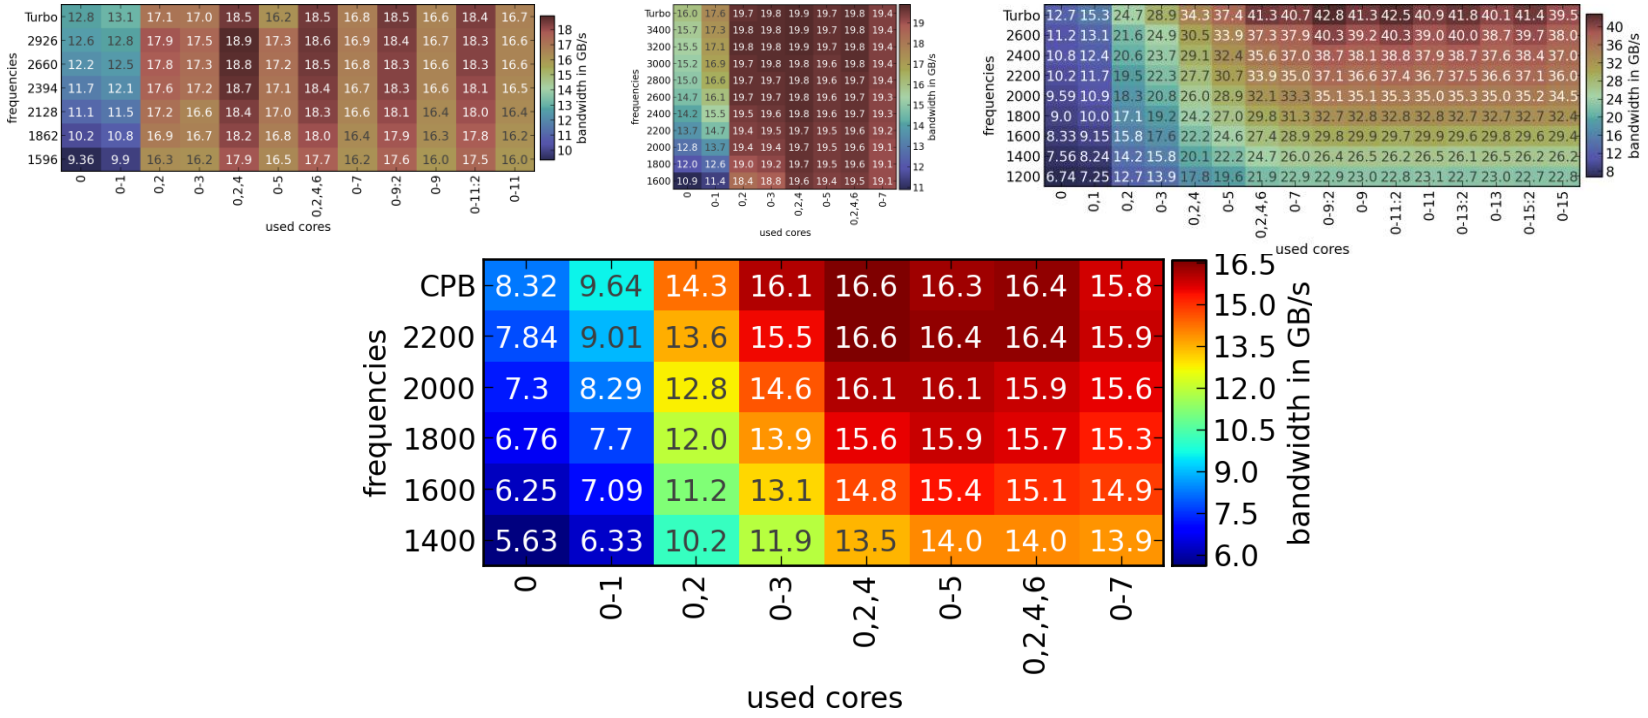

- Only moderate DRAM-performance losses when using DVFS techniques for three or four active cores
- Best performance for three active modules

Bandwidth to DRAM – Summary

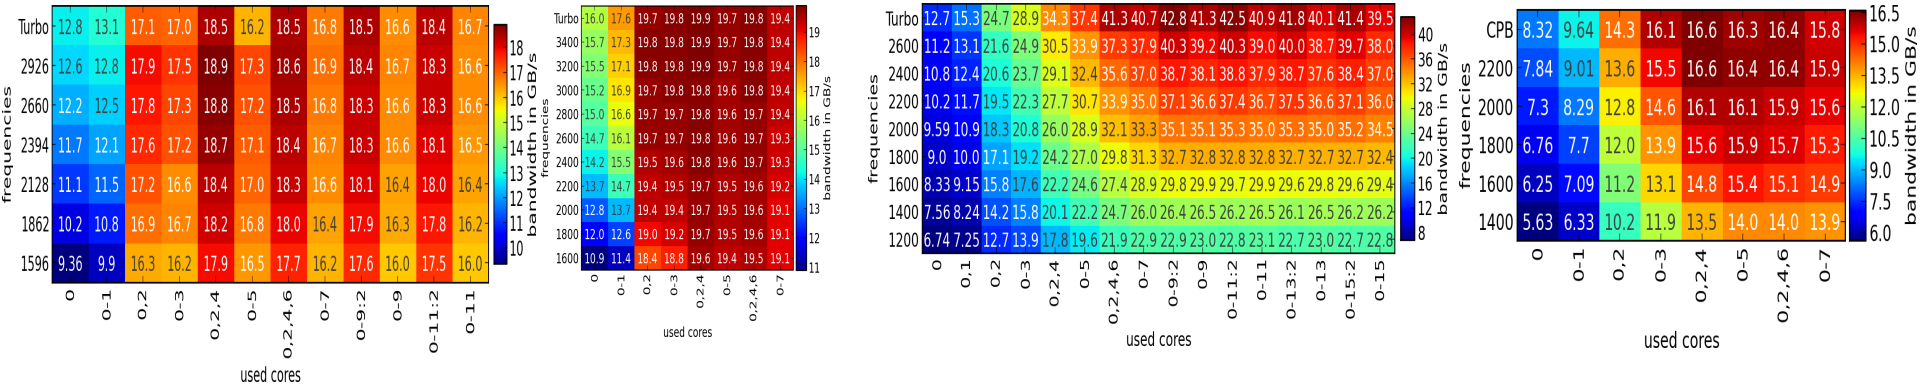

- DVFS provides no or minor impact on DRAM performance of fully occupied Intel processors – except for Sandy Bridge-EP
- Previous AMD processors (Istanbul, Magny-Cours) lower northbridge frequency with core frequency and lose memory performance
- Reduce number of tasks accessing memory

Bandwidth to L3 at Full Concurrency

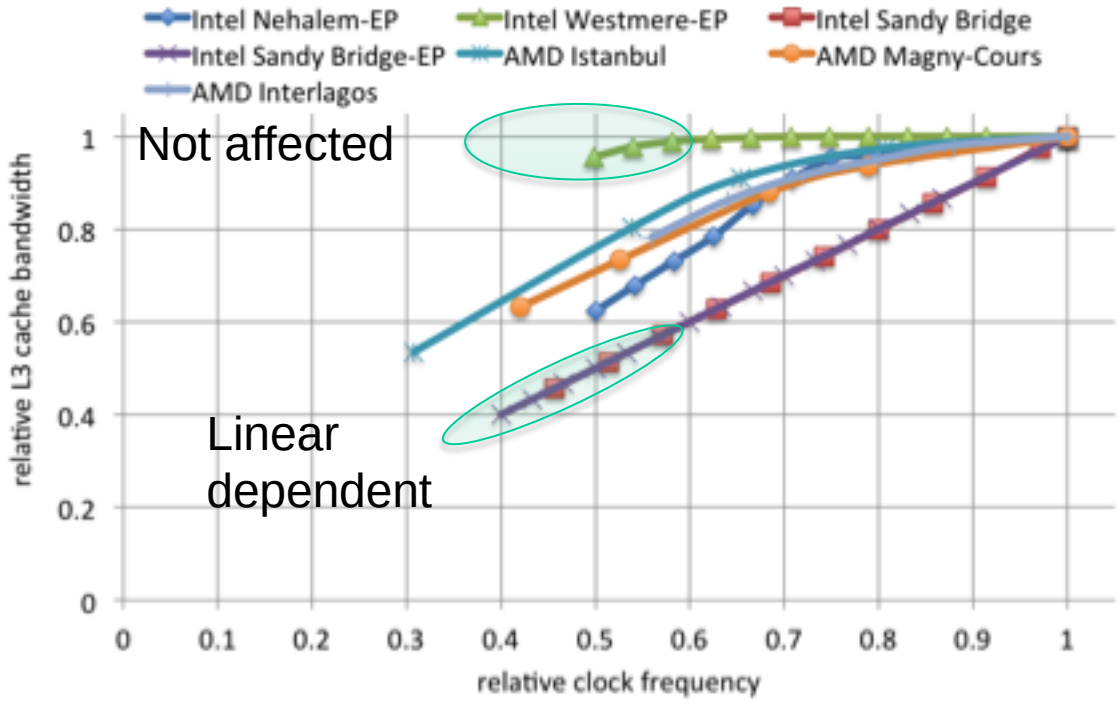

- Linear dependency on Sandy Bridge
- Westmere-EP provides high bandwidth at low frequency
- AMD processors and Nehalem-EP show similar behavior

Bandwidth to L3 – Details

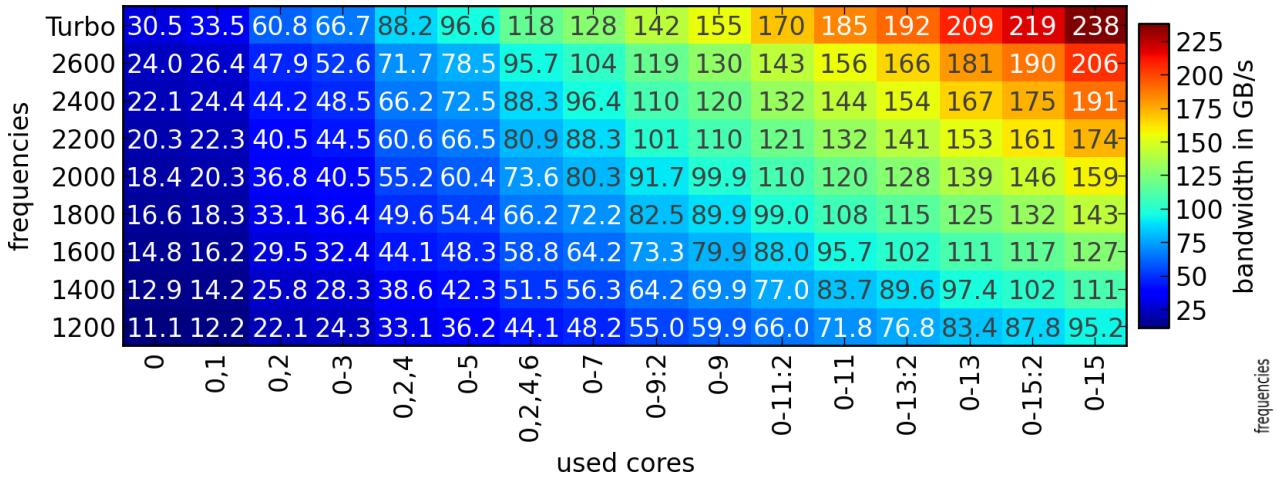

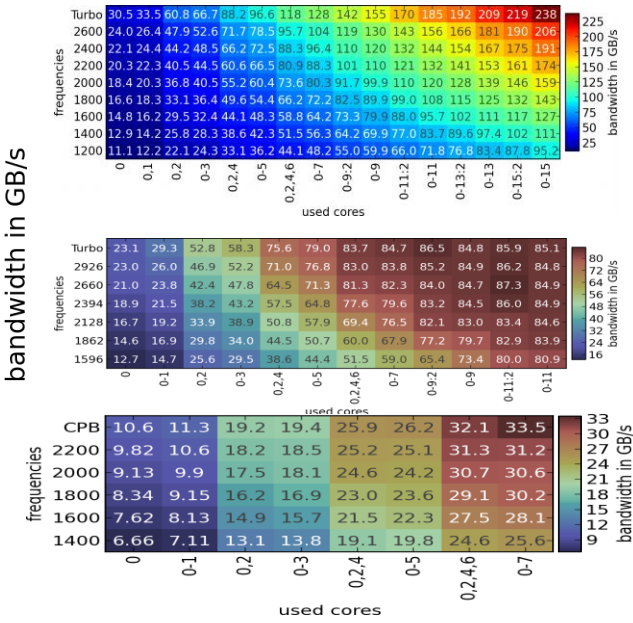

- Sandy Bridge L3 bandwidth depends on frequency

Bandwidth to L3 – Details

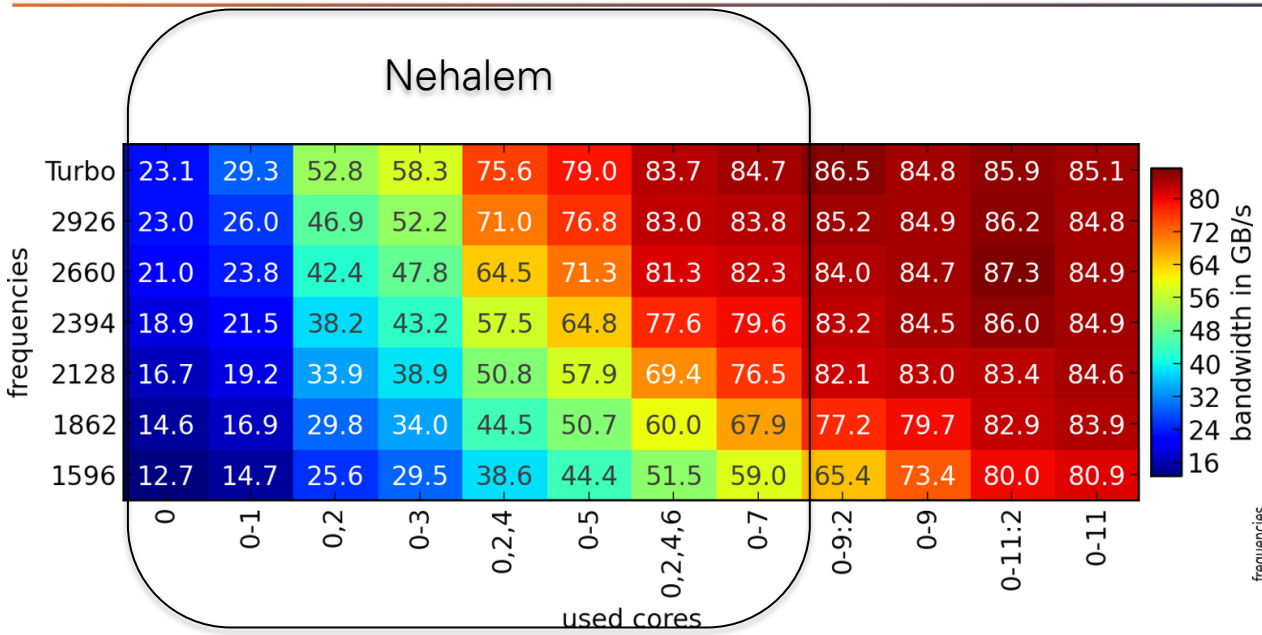

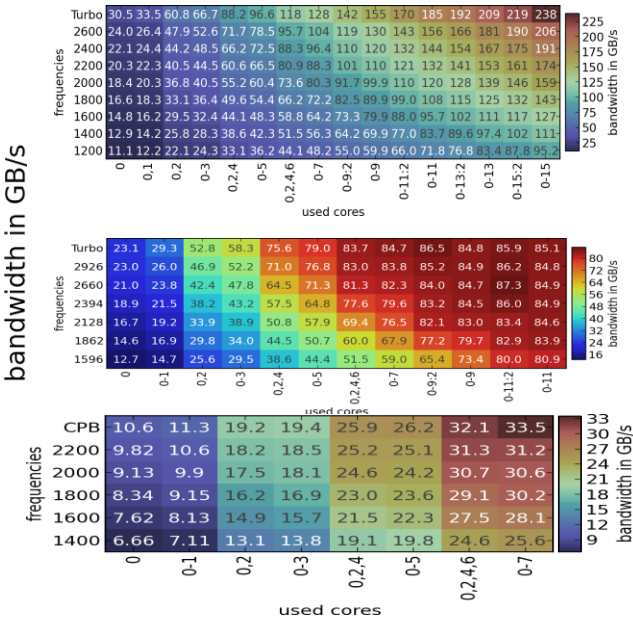

- Sandy Bridge L3 bandwidth depends on frequency
- Westmere-EP needs all cores for full bandwidth at lowest frequency

Bandwidth to L3 – Details

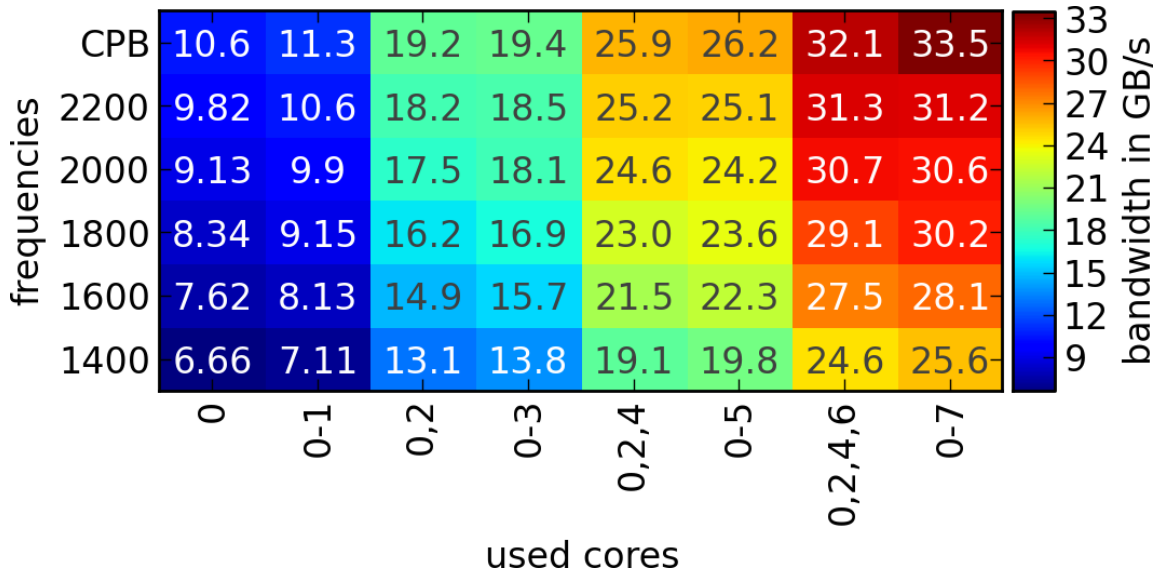

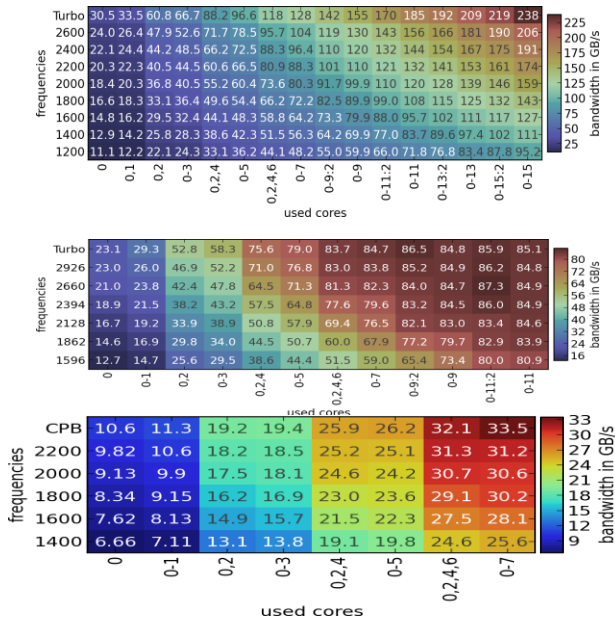

- Sandy Bridge L3 bandwidth depends on frequency
- Westmere-EP needs all cores for full bandwidth at lowest frequency
- AMD L3 bandwidth depends on core frequency – despite separate northbridge frequency domain

Conclusion

- Before using DVFS, check performance at different frequencies
- No common explanation for all system, highly hardware dependent
- Reducing the number of memory bound tasks often an alternative, can be used additionally on some systems

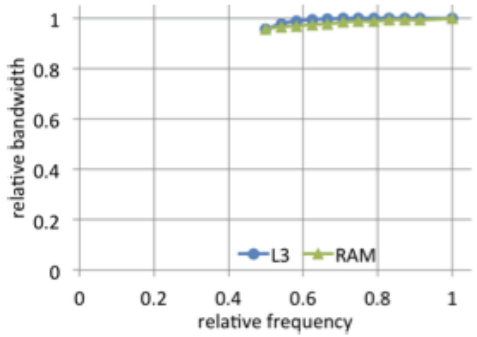

(a) Westmere-EP

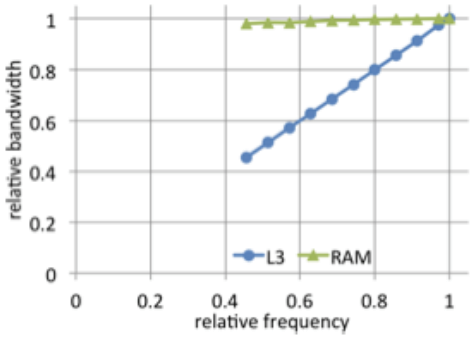

(b) Sandy Bridge-HE

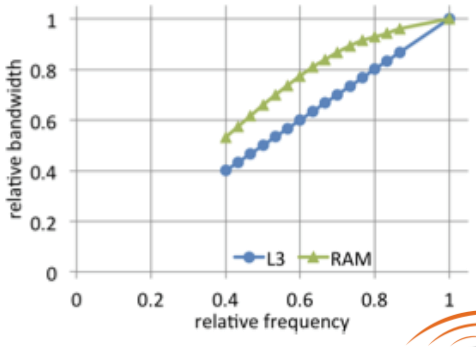

(c) Sandy Bridge-EP

COOL Computing

Questions?

Fourth International Conference on

Energy-Aware High Performance Computing

2nd and 3rd of September 2013, Dresden, Germany

Acknowledgement: Cool Silicon Spitzencluster

Sustainability through energy efficiency. The region "Silicon Saxony" acted on this maxim and founded the Leading Edge Cluster "Cool Silicon - Energy Efficiency Innovations from Silicon Saxony".

The Leading Edge Cluster Cool Silicon

- is aiming at providing a technological basis for increasing the energy efficiency in the information and communications technology sector.
- is assuring Germany and Europe a leading role in the key technology of "energy efficiency in the information technology sector".

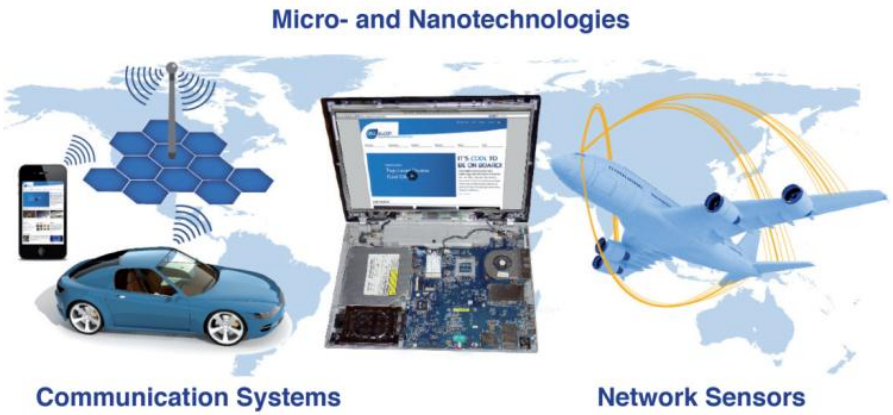

Helping to solve a global conflict: The worldwide desire for free communication collides with global climate protection targets, in case the energy consumption of communications technology keeps expanding at its current pace.

Cool Silicon - Areas

Area 1
Micro- und Nanotechnologies

The core objective of the Area 1 project partners is the development of basis technologies, analysis and production methods for the production of energy efficient electronics and their application in order to decrease the energy consumption of computer systems.

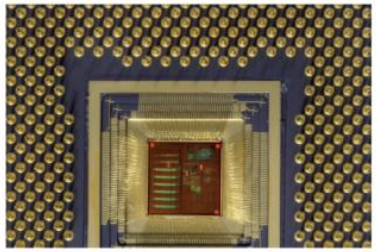

Area 2
Communication Systems

In Area 2, the research and development projects are focussing on the improvement of energy efficiency in communications infrastructures and mobile devices.

Area 3
Network Sensors

The project CoolSensorNet is the Lead Project of Area 3. It conducts research on the whole electronic chain's specific requirements, including sensors, analog electronics, A-D converters, processor systems and the telemetry unit.

Funded via BMBF 13N10186

References

[AMD12] BIOS and Kernel Developer's Guide (BKDG) for AMD Family 15h Models 00-0Fh Processors; Advanced Micro Devices; Rev. 3.08; March 12

[Hac09] Hackenberg, Molka, Nagel; Comparing Cache Architectures and Coherency Protocols on x86-64 Multicore SMP Systems;MICRO 42; 2009

[Int10] The Uncore: A Modular Approach to Feeding the High-Performance Cores; Hill, Bachand, Bilgin, Greiner, Hammarlund, Haff, Kulick, Safranek; Intel Technology Journal; Volume 4; Issue 3; 2010

[Int11] 2nd Generation Intel Core Processor Family: Intel core i7, i5 and i3; Oded Lempel; HotChips 23; 2011

[Int12] Intel® Xeon® Processor E5-2600 Product Family Uncore Performance

Monitoring Guide, Intel; March 2012

[Liv12] Computer using too much power? Give it a REST (Runtime Energy Saving Technology); Livingston, Triquenaux, Fighiera, Beyler, Jalby; Computer Science - Research and Development; 2012