An overview of uncore frequency impact on performance and power consumption

Rémi Dulong[1], **Amina Guermouche**[2], François Trahay[2]

1 : Université de Neuchâtel
2: Télécom SudParis

June 26, 2019

TELECOM SudParis

IP PARIS

Introduction

- Supercomputer Power consumption keeps increasing
  Summit $10MW$ (~ 8000 american households)

- Dynamic Voltage and Frequency Scaling (DVFS)
  - Set the cores frequency according to the workload

- Uncore frequency scaling (UFS)
  - Set uncore components (L3 cache, memory controllers, ...) frequency

Uncore frequency

- Uncore components frequency (L3 cache, memory controllers, . . .)
  - Sandy Bridge-EP and Ivy Brige-EP -> same frequency domain
  - Since Haswell, one domain per core

→ the processor controls the **uncore** independently from the core

→ Uncore Frequency Scaling (UFS) is used

Uncore frequency

- Uncore components frequency (L3 cache, memory controllers, . . .)
  - Sandy Bridge-EP and Ivy Brige-EP -> same frequency domain
  - Since Haswell, one domain per core
  - → the processor controls the **uncore** independently from the core
  - → Uncore Frequency Scaling (UFS) is used

- Uncore Frequency Scaling (UFS) is affected by:
  1. Stall cycles of the cores
  2. Energy and Performance and Bias Hints (EPB)
  3. c-state

Outline

1. Uncore frequency behavior and impact on memory
2. Impact on processor and DRAM power consumption
3. Impact on application performance

Comparing UFS algorithms on different architecture

HPL (NB=224), NAS SP (NPB 3.3.1)

1. nova: Intel Xeon E5-2620 v4 (Broadwell) - Grid'5000 Lyon cluster
    - 2x8 cores, 64GB of memory
    - ufreq interval: [1.2GHz, 2.7 GHz]
    - HPL N=58912 (PxQ)=(4x4)
    - SP.C.16
2. chifflet: Intel Xeon E5-2680 v4 (Broadwell) - Grid'5000 Lille cluster
    - 2x14 cores, 768GB of memory
    - ufreq interval: [1.2GHz, 2.7GHz]
    - HPL N=100996 (PxQ)=(4x7)
    - SP.C.16
3. yeti: Intel Xeon Gold 6130 (Skylake) - Grid'5000 Grenoble cluster
    - 4x16 cores, 252GB of memory
    - ufreq interval: [1.2GHz, 2.4 GHz]
    - HPL N=91840 (PxQ)=(8x8)
    - SP.D.64

UFS algorithms

Machine	HPL	NAS SP
ufreq nova (GHz)	2.7	2.7
ufreq chifflet (GHz)	2.4	2.7
ufreq yeti (GHz)	2.0	2.4

- nova: Set to the maximum as soon as at least one core is active
- chifflet: Set according to CPU frequency and memory load
- yeti: Set according to CPU frequency and memory load

Uncore frequency impact on cache and memory latency and bandwidth

- Platform: Grid5000 Lille => chifflet cluster (chifflet-1)
  - Intel Xeon E5-2680 v4
  - 2 CPUs
  - 14 cores per CPU
  - Uncore frequency: [1.2GHz,2.7GHz]
  - Memory (on 1 CPU):
    - L1 cache: 64KB
    - L2 cache: 256 KB
    - L3 cache: 35 MB
    - Memory: 378 GB

- Measurements: lmbench (version 3.0-1a)
  - latency: lat_mem_rd
  - bandwidth: bw_mem

Uncore frequency impact on cache and memory latency and bandwidth

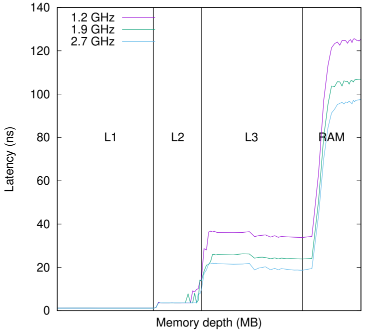

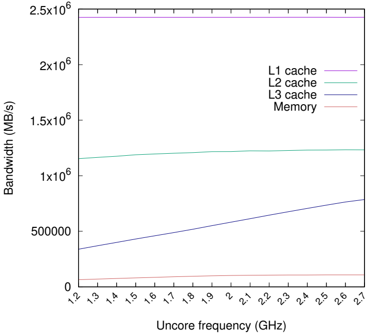

Uncore frequency impact on cache and memory latency and bandwidth

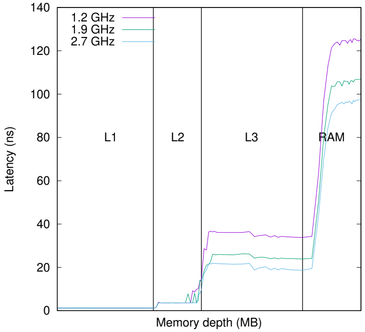

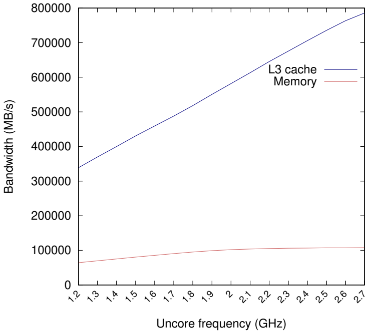

Uncore frequency impact on cache and memory latency and bandwidth

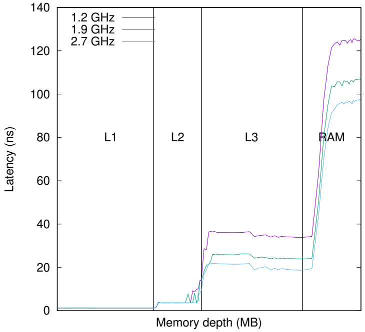

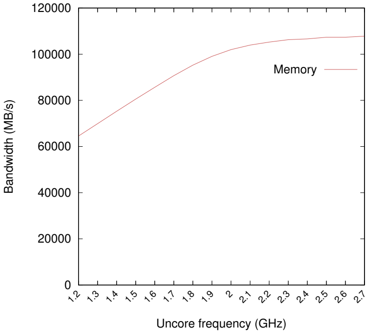

Outline

1. Uncore frequency behavior and impact on memory
2. Impact on processor and DRAM power consumption
3. Impact on application performance

Settings and used applications

- Platform: Grid5000 Lille => chifflet cluster (chifflet-1)
  - Intel Xeon E5-2680 v4
  - 2 CPUs
  - 14 cores per CPU
  - Turboboost Frequency 2.8GHz
  - Power measurement: LIKWID (likwid-perfctr) every 1s
  - Uncore frequency interval: [1.2 GHz, 2.7 GHz]
  - Turboboost activated
  - Hyperthreading deactivated

- Applications
  - HPL: CPU-intensive - N=100996, NB=224, (PxQ) = (4x7)
  - NAS Parallel Benchmarks (NPB 3.3.1):
    - lu(C), ep(C): CPU-intensive
    - bt (C), cg(C), ft(C), sp(C), mg(D): memory-intensive

Uncore frequency impact on processor power consumption

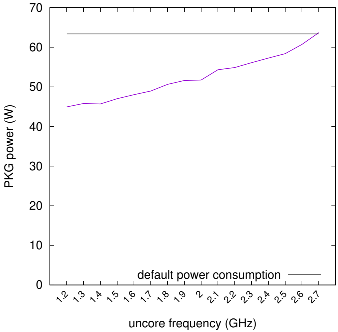

CG.C.16 (only socket 0)

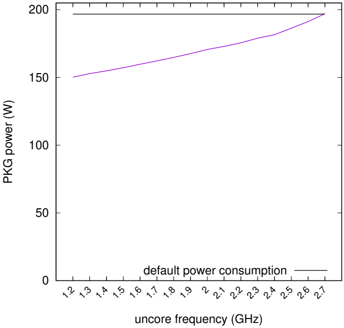

LU.C.28 (socket 0 + socket 1)

Uncore frequency impact on memory power consumption

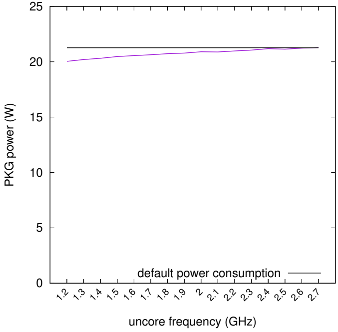

CG.C.16 (only socket 0)

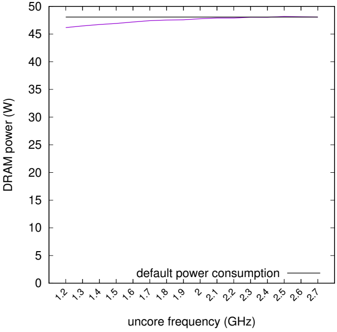

LU.C.28 (socket 0 + socket 1)

Uncore frequency impact on applications power consumption

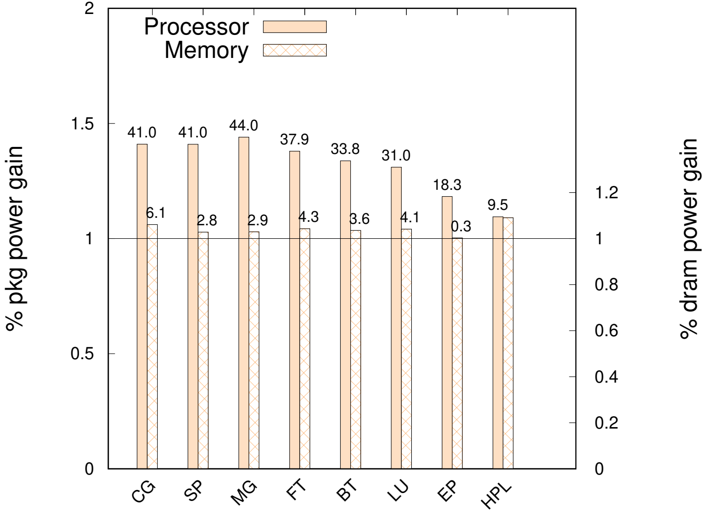

Outline

1. Uncore frequency behavior and impact on memory
2. Impact on processor and DRAM power consumption
3. Impact on application performance

Uncore frequency impact on applications performance

application	Max. PKG power gain	Max. DRAM power gain	Max. perf. loss
CG	41%	6.11%	36.28%
SP	40.95 %	2.79%	11.6%
MG	44%	2.93%	14.5%
FT	38%	4.28%	22.66%
BT	33.77%	3.55%	19%
LU	31.37%	4.12%	16%
EP	18.2%	negligible	0.4%
HPL	9.5%	9%	30.7%

Uncore frequency impact on Memory-intensive applications performance

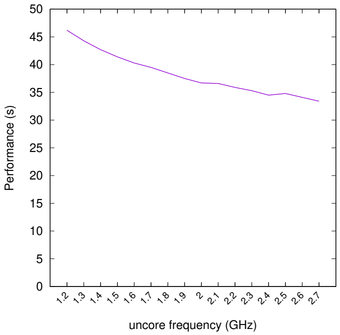

CG.C.16 performance

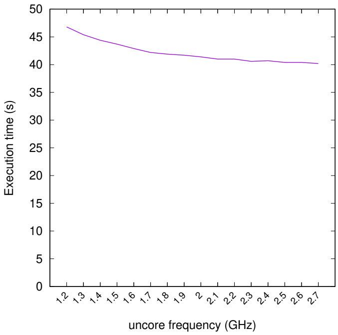

LU.C.28 performance

Power gain VS eperformance loss for memory intensive applications

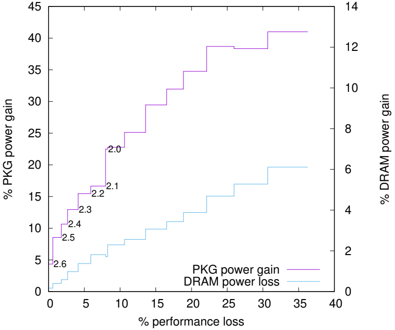

CG.C.16
8% energy gain

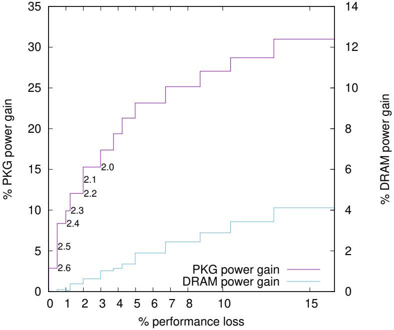

LU.C.28
10.35% energy gain

Power gain with performance loss upper bound (2%)

application	PKG power gain	DRAM power gain	perf. loss	ufreq (GHz)	default ufreq (GHZ)
CG	10.65%	negligible	1.76	2.4	2.7
SP	12.1%	negligible	1.19	2.4	2.7
MG	10.47%	negligible	1.47	2.4	2.7
FT	10.59%	negligible	1.78%	2.4	2.7
BT	9.11 %	negligible	1.87	2.4	2.7
LU	13.7%	negligible	1.99%	2.1	2.7
EP	18.2%	negligible	0.4%	1.2	2.7
HPL	negligible	negligible	-2.14%	2.1	2.4

Power gain with performance loss upper bound (2%)

application	PKG power gain	DRAM power gain	perf. loss	ufreq (GHz)	default ufreq (GHZ)
CG	10.65%	negligible	1.76	2.4	2.7
SP	12.1%	negligible	1.19	2.4	2.7
MG	10.47%	negligible	1.47	2.4	2.7
FT	10.59%	negligible	1.78%	2.4	2.7
BT	9.11 %	negligible	1.87	2.4	2.7
LU	13.7%	negligible	1.99%	2.1	2.7
EP	18.2%	negligible	0.4%	1.2	2.7
HPL	negligible	negligible	-2.14%	2.1	2.4

Using uncore frequency to improve application performance

- Context:
  - Processors are designed to stay within a limit of power consumption (Thermal Design power)
  - When processors reach this limit, core frequencies are decreased
- HPL behavior:
  - HPL power consumption reaches TDP
  - The average frequency when running HPL is 2.4GHz (while the default frequency is 2.8GHz)
- Uncore frequency as a power leverage
  - Decreasing uncore frequency reduces power consumption
  - CPU frequency is automatically increased
  - ☺ Performance is improved while power consumption is the same

Using uncore frequency to improve application performance

- Platform: Grid5000 Grenoble => yeti cluster (yeti-2)
  - Intel Xeon Gold 6130
  - 4 CPUs
  - 16 cores per CPU
  - Turboboost Frequency 2.4GHz (1.9GHz with AVX512)
  - Power measurement: LIKWID (likwid-perfctr) every 1s
  - Uncore frequency interval: [1.2 GHz, 2.4 GHz]
  - Turboboost activated
  - Hyperthreading deactivated
- Powercap (powercap-set tool)
  - BT (124W), CG (124W), MG (124W), LU (124W), SP(124W): 100W
  - EP (115W): 98W
  - FT (123W): 118W

Using uncore frequency to improve application performance

application	default ufreq (GHz)	default freq(GHz)	perf. gain	ufreq (GHz)	freq (GHz)
BT	1.8	2.3	3.74%	1.6	2.4
EP	2.2	2.3	18.6%	1.2	2.7
LU	1.8	2.1	15%	1.5	2.3

Outline

1. Uncore frequency behavior and impact on memory
2. Impact on processor and DRAM power consumption
3. Impact on application performance

Conclusion and current

- Conclusion
  - Applications can benefit from reducing uncore frequency for power and energy consumption
  - CPU-intensive applications can show a gain in performance when decreasing uncore frequency
  - Processors tend to reduce CPU frequency instead of uncore frequency to stay within a powercap
- Things to do
  - Study more applications (ideas are welcome)
  - Study correlation with L3 cache
- Current work
  - Combine uncore frequency scaling and task placement
  - Combine UFS, DVFS and power capping