BENCHMARKING THE EVOLUTION OF PERFORMANCE AND ENERGY EFFICIENCY ACROSS RECENT GENERATIONS OF INTEL XEON PROCESSORS

Istvan Z Reguly, Balázs Dravai – PPCU ITK, Hungary
reguly.istvan@itk.ppke.hu

PMBS 24

SC24
Atlanta, GA | hpc creates.

INTEL CPU GENERATIONS

- Sapphire Rapids – 2023 Q1
  - Intel 7, first chiplet architecture - 4 compute tiles
  - Up to 56 cores, 350W
  - 8xDDR5-4800 or HBM2E
- Emerald Rapids – 2023 Q4
  - Intel 7, 2 compute tiles
  - Up to 64 cores, 350W
  - 8xDDR5-5600
- Sierra Forest – 2024 Q2
  - Intel 3, 2 compute tiles
  - Up to 192 cores, 350W
  - 8x/12x DDR5-6400
- Granite Rapids – 2024 Q3
  - Intel 3, 3 compute tiles
  - Up to 128 cores, 500W
    - 72 core 6960P tested
  - 12xDDR5-8800 MRDIMM

AMD PROCESSOR GENERATION

- AMD Milan – 2021 Q1
  - 7 nm, 8 chiplet
  - 64 cores, 280W
  - 8x DDR4-3200
- AMD Genoa – 2022 Q4
  - 5 nm, 12 chiplet
  - 96 cores, 360W
  - 12x DDR5-4800
- AMD Turin – 2024 Q3
  - 3 nm, up to 16 chiplet
  - 128/192 cores, 500W
  - 12x DDR5-6400

PERFORMANCE, EFFICIENCY

- Number of bandwidth-bound codes (mostly explicit PDE solvers)
  - CloverLeaf (low-order), Seismic (high-order), OpenSBLI (large-scale DNS), MG-CFD (FV Euler on unstructured mesh), Volna (FV NLSW on unstructured mesh)
- Compute-intensive: miniBUDE (docking proxy)
- MPI or MPI+OpenMP – many through OPS/OP2 DSLs. icpx/g++ compilers.

- Compare:
  - Runtime
  - Architectural efficiency (effective BW estimates)
  - Energy efficiency

BASELINES

	SPR HBM	SPR DDR	EMR	GNR	SRF 96c	SRF 192c	Milan	Genoa
Model	9480	8480+	8592+	6960P	6740E		7763	9B14
Cores	112	112	128	144	192	384	128	180
LLC (MB)	105	105	320	432	96	192	256	384
Cache BW	3481	4340	8149	7346	4139	7627	2454	9534
DDR BW	1475	388	542	1150	396	667	234	529
Speedup	3.8	1.0	1.39	2.96	1.04	1.23	0.6	1.36

Point of reference

RUNTIME VS SAPPHIRE RAPIDS+DDR5

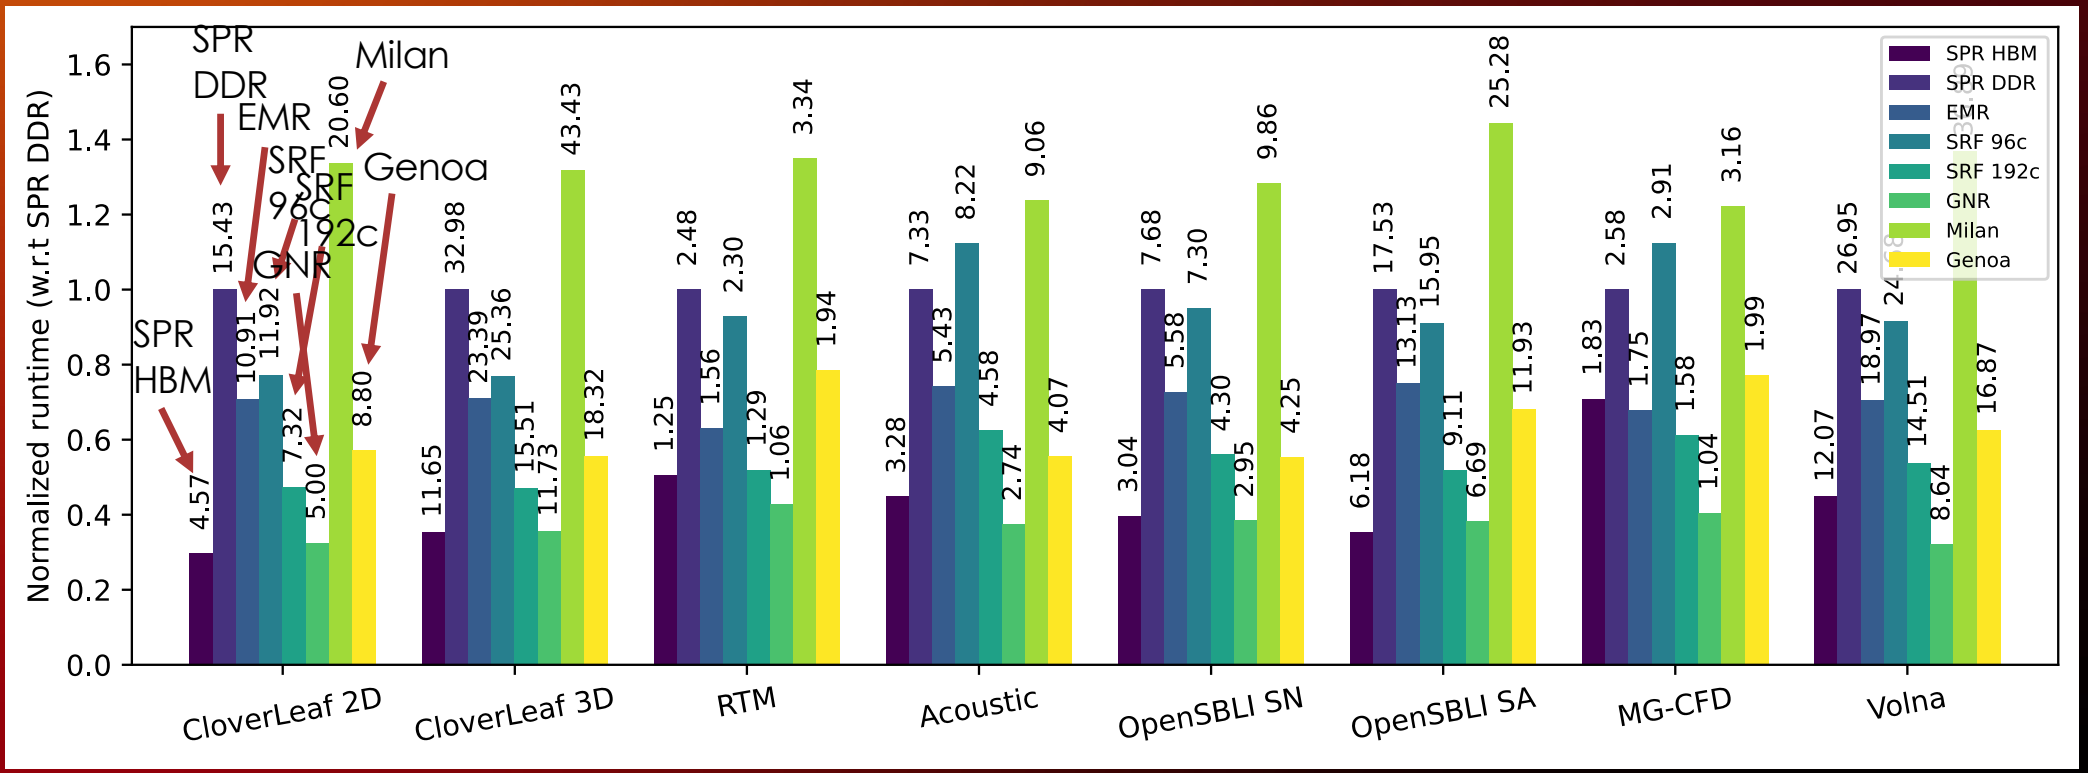

ARCHITECTURAL EFFICIENCY

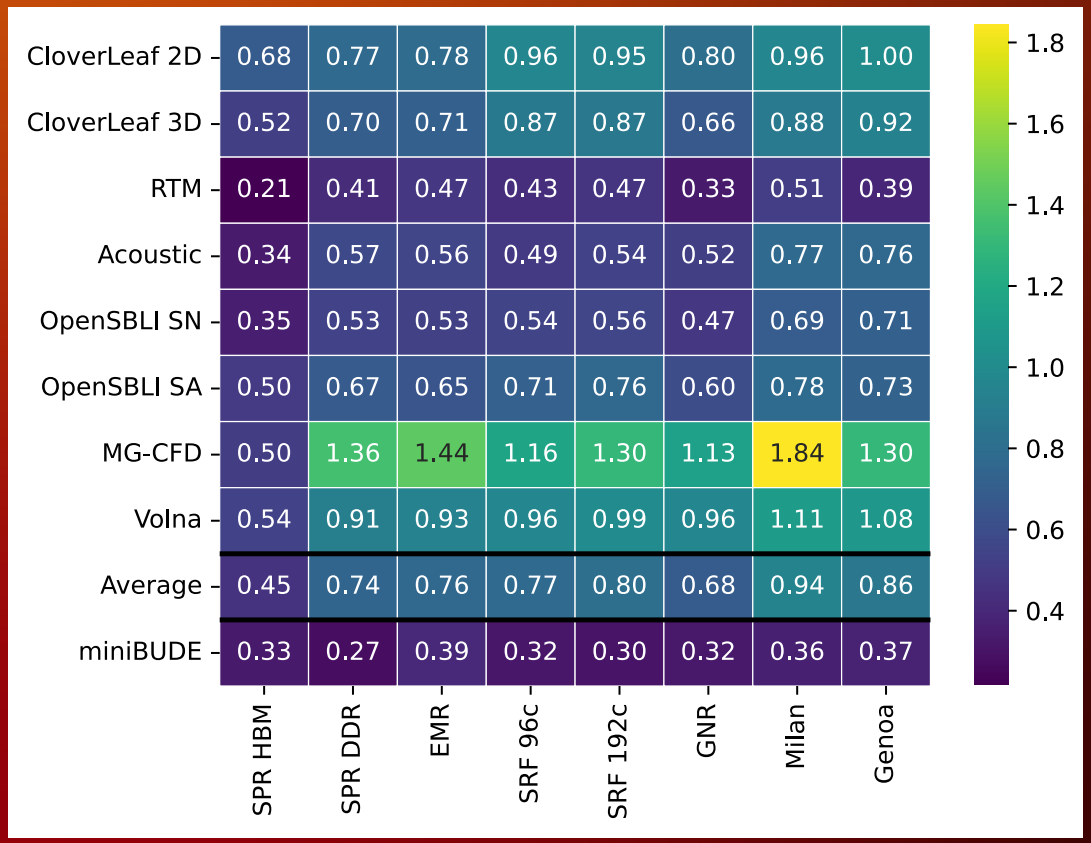

- Effective BW: array size in each loop
  - Above 1.0 if cache re-use across loops
- Significant MPI overheads on Seismic apps (30-50%)
  - Especially SPR+HBM and Genoa
- miniBUDE: fraction of peak GFLOPS/s at all-core turbo
  - Not always documented

BALANCED IMPROVEMENT VS SPR+DDR?

- How much did overall performance improve vs. the improvement in bandwidth?
  - (runtime/runtime$_{SPR}$)/(Peak BW/Peak BW$_{SPR}$)*100
- SPR+HBM: same compute - only 64%
- EMR: +8 cores, 3x cache size, 1.4x BW – 98%
- GNR: +16 cores, 4.1x cache size, 3x BW – 91%

- SRF 96 core: +40 cores, 0.85x cache size, 1.04x BW – 105%
- SRF 192 core: +136 cores, 1.71x cache size, 1.7x BW – 110%

ENERGY EFFICIENCY VS SAPPHIRE RAPIDS

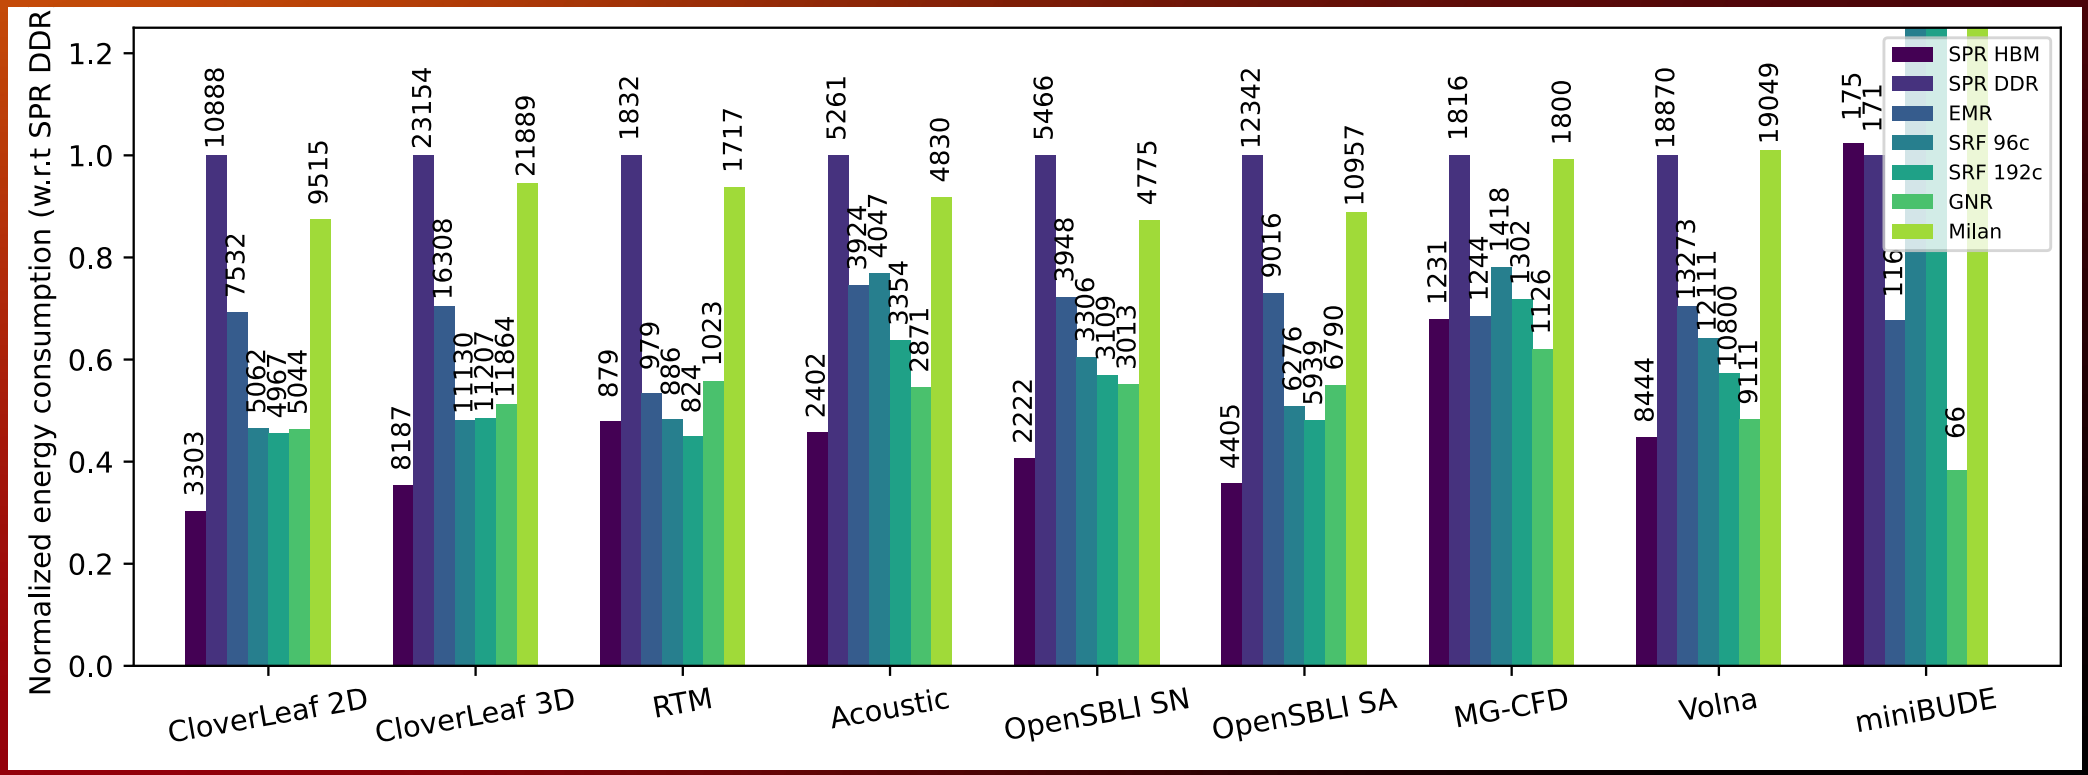

POWER

- 250W TDP for SRF 96c
- 280W TDP for Milan
- 350W TDP for SPR, EMR, SRF 192c
- 360W TDP for Genoa (no RAPL)
- 500W TDP for GNR
- Plus RAM

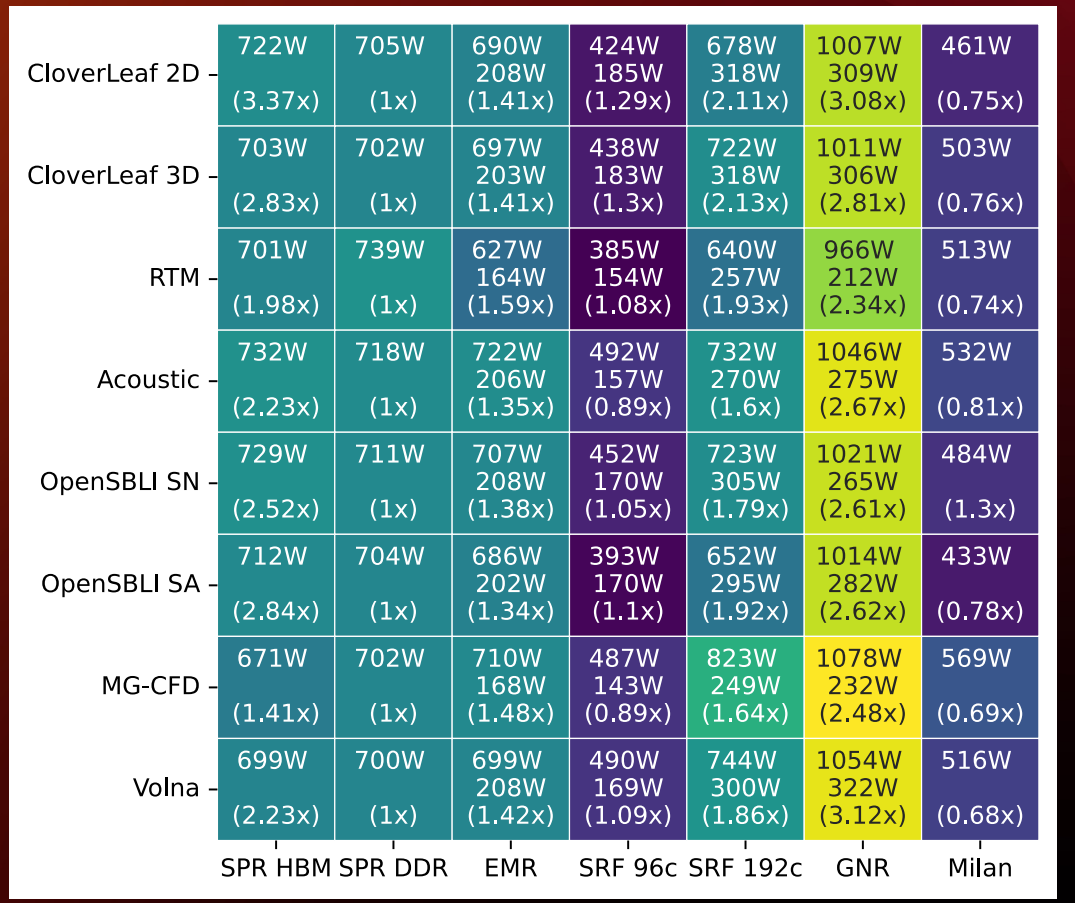

CONCLUSIONS

- Rapid evolution with interesting trade-offs
- Intel's differentiated product lines – evolution vs. SPR
  - SRF 96 core: Slightly higher performance (1.09×) but at a 1.4× lower power
  - SRF 192 core: Significantly higher performance(1.87×), at the same power
  - GNR: Even higher performance (2.72×), but at 1.44× more power
- MRDIMM may be a critical differentiating factor vs. AMD
  - Big advantage for memory-bound codes
- Still behind with energy efficiency – process issues