VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

MAX

Co-design, technology exploitation & energy efficiency

Filip Vaverka

IT4Innovations národní superpočítačové centrum

Studentská 6231/1B

708 00 Ostrava-Poruba

filip.vaverka@vsb.cz

2026

MaX (Materials design at the eXascale): WP4 Objectives

- MaX flagship codes
  - BigDFT, Fleur, Quantum ESPRESSO, Siesta, **Yambo**
- **Provide data to developers** (EPI[1], EUPEX[2], EUPilot[3] and other EU initiatives)
  - Collect realistic data about HW and SW requirements of the MaX codes
  - Memory sub-system architecture and tradeoffs
- Prepare MAX codes for future HPC hardware
- Explore the use of the Modular Supercomputing Architecture (MSA)
- Study the scaling and performance of the flagship applications and workflows on hardware prototypes
  - Provide application characterization data for co-design of upcoming extreme-scale architectures
- **Evaluate the energy consumption** of the MAX codes on selected hardware platforms
  - Perform energy consumption analysis and optimization

---

[1] **EPI** European Processor Initiative (CPU)
[2] **EUPEX** European Pilot for Exascale (Platform)
[3] **EUPilot** The European PILOT

MaX Flagship Codes - Frequency Scaling (CPU)

- Quantum ESPRESSO (QE) and **Yambo**[4]
  - Expected scaling with peak power of 412 W and 557 W respectively.
- Fleur
  - Relatively low peak power (200 W of 560 W maximum) suggests low hardware utilization (potentially benchmarking artifact).
- Siesta
  - The relative runtime scaling suggests strong memory bandwidth dependency.
- Well-behaved codes show peak efficiency around 2.0 GHz.

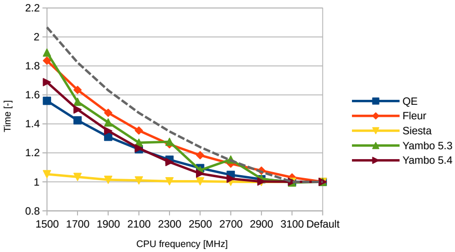

Relative Time

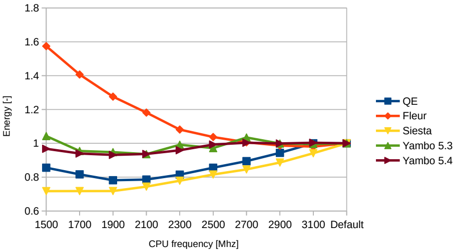

Relative Energy

[4]Karolina ACN (AMD Zen 3)

Yambo 5.3 and 5.4 Performance and Energy Efficiency

Platform	v5.3 (s (kJ))	v5.4 (s (kJ))	Speedup (-)
NVIDIA Grace	478 (257)	382 (222)	1.25 (0.86)
AMD Milan	516 (278)	436 (244)	1.18 (0.87)
AMD Milan-X	532 (282)	444 (252)	1.18 (0.89)
Intel SPR	735 (498)	632 (438)	1.16 (0.88)
IBM Power 10	1277	1127	1.13
AMPERE Altra	2011	1524	**1.32**

Yambo on Advanced HW (default)

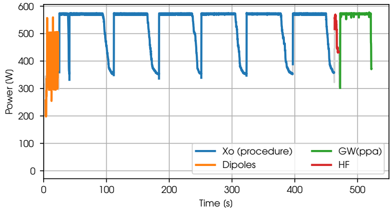

Yambo 5.3 (AMD Zen 3)

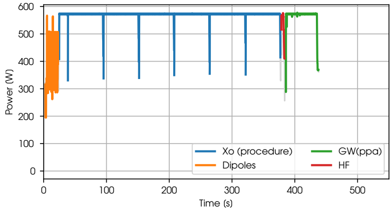

Yambo 5.4 (AMD Zen 3)

- Core-intensive benchmark[5]
  - X matrix size = 301 (0.7 MB)
- Main Yambo 5.4 optimization
  - **Xo (procedure)** orchestration for compute tradeoff
  - Workload imbalance visible in CPU power timeline

[5] GrCo: X-cutoff = **6**, num. of bands = **2000**, k-points = 7

Yambo 5.3 and 5.4 Performance and Energy Efficiency

Platform	v5.3 (s (kJ))	v5.4 (s (kJ))	Speedup (-)
NVIDIA Grace	478 (257)	382 (222)	1.25 (0.86)
AMD Milan	516 (278)	436 (244)	1.18 (0.87)
AMD Milan-X	532 (282)	444 (252)	1.18 (0.89)
Intel SPR	735 (498)	632 (438)	1.16 (0.88)
IBM Power 10	1277	1127	1.13
AMPERE Altra	2011	1524	**1.32**

Yambo on Advanced HW (default)

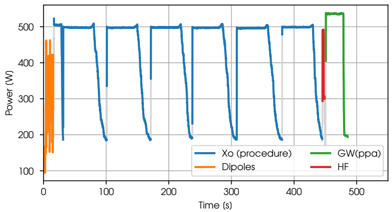

Yambo 5.3 (NVIDIA Grace)

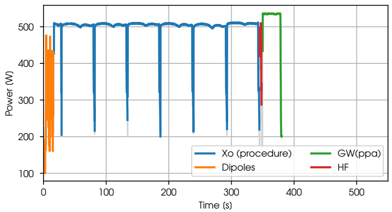

Yambo 5.4 (NVIDIA Grace)

- Core-intensive benchmark[5]
  - X matrix size = 301 (0.7 MB)
- Main Yambo 5.4 optimization
  - **Xo (procedure)** orchestration for compute tradeoff
  - Workload imbalance visible in CPU power timeline

[5]GrCo: X-cutoff = **6**, num. of bands = **2000**, k-points = 7

Intel SPR with HBM (SC4 FLAT Mode) as Rhea 1 Proxy

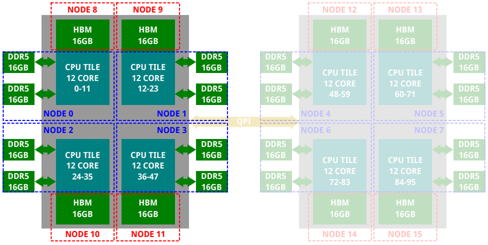

- Intel Xeon Max 9468 CPU
  - 4x 12-core tile with 16 GB of HBM
  - 105 MB of L3 Cache
  - DDR5: ~200 GB/s
  - HBM: ~700 GB/s
  - HBM has ~20 % latency penalty (for dependent loads)

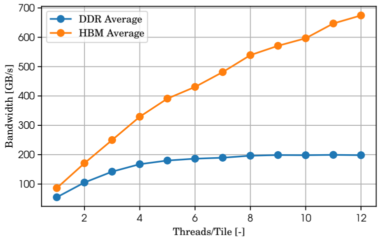

Memory Bandwidth

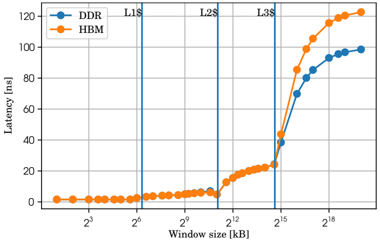

Memory Latency

Yambo 5.3 vs 5.4 on Intel SPR with HBM

Platform	v5.3 (s (kJ))	v5.4 (s (kJ))	Speedup (-)
NVIDIA Grace	327	236	1.39
Intel SPR (DDR)	737 (504)	550 (370)	1.34 (0.73)
Intel SPR (HBM)	534 (356)	378 (244)	1.41 (0.69)
DDR vs HBM	1.38 (0.71)	1.45 (0.66)	**1.95 (0.48)**

Yambo on Advanced HW (default)

- Memory-intensive benchmark[6]
  - X matrix size = 1219 (11.3 MB)
- Most of the speedup comes from Xo (procedure) (2.2x) and Xo (REDUX) (1.5x) regions.
- **Comparable speed-up from optimizations or "new" hardware!**

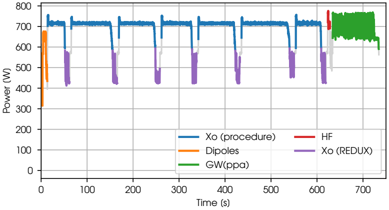

Yambo 5.3 (Intel SPR - DDR)

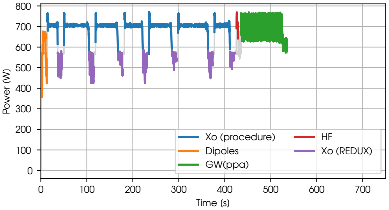

Yambo 5.3 (Intel SPR - HBM)

[6] GrCo: X-cutoff = 15, num. of bands = 200, k-points = 7

Yambo 5.3 vs 5.4 on Intel SPR with HBM

Platform	v5.3 (s (kJ))	v5.4 (s (kJ))	Speedup (-)
NVIDIA Grace	327	236	1.39
Intel SPR (DDR)	737 (504)	550 (370)	1.34 (0.73)
Intel SPR (HBM)	534 (356)	378 (244)	1.41 (0.69)
DDR vs HBM	1.38 (0.71)	1.45 (0.66)	**1.95 (0.48)**

Yambo on Advanced HW (default)

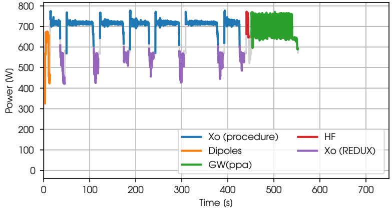

Yambo 5.4 (Intel SPR - DDR)

- Memory-intensive benchmark[6]
  - X matrix size = 1219 (11.3 MB)
- Most of the speedup comes from Xo (procedure) (2.2x) and Xo (REDUX) (1.5x) regions.
- **Comparable speed-up from optimizations or "new" hardware!**

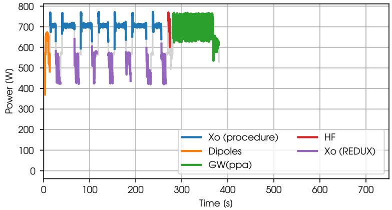

Yambo 5.4 (Intel SPR - HBM)

[6] GrCo: X-cutoff = **15**, num. of bands = **200**, k-points = 7

Object Level Data Placement

- Process level approach
  - Whole process is associated with specific memory type.
  - Potentially inefficient memory pool usage (size/bandwidth/latency).

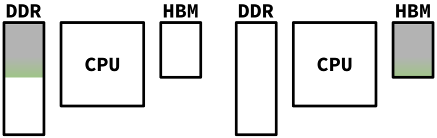

Process Memory Affinity

Object Level Data Placement

- Process level approach
  - Whole process is associated with specific memory type.
  - Potentially inefficient memory pool usage (size/bandwidth/latency).

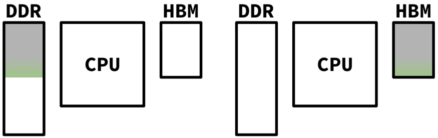

Process Memory Affinity

- Object level approach
  - Identify individual objects in the application (profiling and perf. counters).
  - Select suitable memory pool for each object to maximize performance.

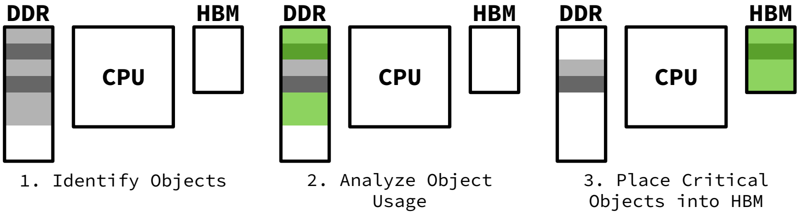

Object Memory Affinity

Intel SPR with HBM (SC4 FLAT Mode) as Rhea 1 Proxy

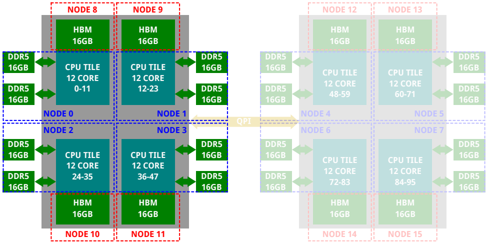

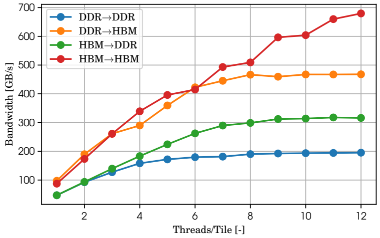

STREAM: Copy

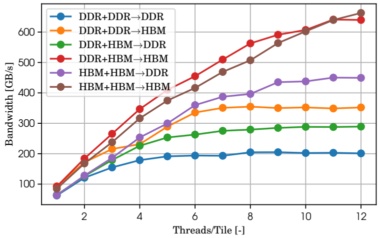

STREAM: Add

- Intel Xeon Max 9468 CPU
  - 4x 12-core tile with 16 GB of HBM
  - 105 MB of L3 Cache
  - DDR5: ~200 GB/s
  - HBM: ~700 GB/s
  - HBM has ~20 % latency penalty (for dependent loads)

YAMBO and Quantum ESPRESSO Mini-apps Results

- **YAMBO Mini-app**
  - **90%** of the data has to be placed into HBM to achieve all potential speedup (1.8×).
- **Quantum ESPRESSO Mini-app**
  - Just **50%** of the data in HBM realizes all potential speedup (1.4×).
- Other benchmarks (NPB)
  - 90 % of speed-up achieved with 60 – 70 % of the data in HBM.

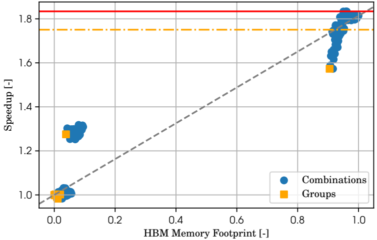

YAMBO Mini-app

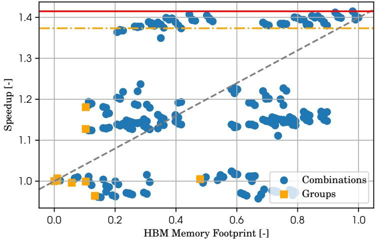

QE Mini-app (very large)

Thank you.

Filip Vaverka

IT4Innovations národní superpočítačové centrum

Studentská 6231/1B

708 00 Ostrava-Poruba

filip.vaverka@vsb.cz

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

MAX