



AMBER AND JUPITER: FIRST EXPERIENCES

AMBER/GROMACS DEVELOPERS MEETING 2025

16 June 2025 | Andreas Herten | Forschungszentrum Jülich, Jülich Supercomputing Centre

Content

JUPITER

- System Overview

- Benchmark Suite

- JUPITER Benchmark Suite

 - Infrastructure

 - Benchmark Details

 - TCO Results

JUREAP

- First Results
- Conclusion


JUPITER

⚡ About JUPITER

- **JUPITER:** First European exascale supercomputer (HPL: 1 EFLOP/s *soon*)
 - Procured by EuroHPC JU, BMFTR (Federal Ministry of Research, Technology, and Space), MKW (NRW Ministry of Culture and Science)
 - Hosting entity: Forschungszentrum Jülich / Jülich Supercomputing Centre for Gauß Center for Supercomputing (GCS)

⚡ About JUPITER

- **JUPITER:** First European exascale supercomputer (HPL: 1 EFLOP/s *soon*)
 - Procured by EuroHPC JU, BMFTR (Federal Ministry of Research, Technology, and Space), MKW (NRW Ministry of Culture and Science)
 - Hosting entity: Forschungszentrum Jülich / Jülich Supercomputing Centre for Gauß Center for Supercomputing (GCS)



JUPITER System Overview

- ParTec/Eviden consortium
- Implementing Modular Supercomputing Architecture



JUPITER System Overview

- ParTec/Eviden consortium
- Implementing Modular Supercomputing Architecture
- JUPITER **Booster**: High scalability, 1 EFLOP/s HPL, > 40 EFLOP/s FP8
≈ 6000 nodes: 4× Grace-Hopper superchip, 4× network
- JUPITER **Cluster**: High versatility, high memory-to-compute ratio
≈ 1300 nodes: 2× SiPearl Rhea1 (HBM), 1× network




JUPITER System Overview

- ParTec/Eviden consortium
- Implementing Modular Supercomputing Architecture
- JUPITER **Booster**: High scalability, 1 EFLOP/s HPL, > 40 EFLOP/s FP8
≈ 6000 nodes: 4× Grace-Hopper superchip, 4× network
- JUPITER **Cluster**: High versatility, high memory-to-compute ratio
≈ 1300 nodes: 2× SiPearl Rhea1 (HBM), 1× network
- Network: 200/400 Gbit/s NVIDIA InfiniBand NDR (DragonFly+)
- Storage: 29 PB flash, 310 PB HDD, 370 PB tape
- Energy: 17 MW limit (HPL); direct liquid-cooled, energy re-use
- Modular data center of containers 



JUPITER Booster Node Design



4 × GH200 superchip (CG4)


- Hopper: 132 SMs
- Grace: 72 Arm Neoverse-V2 cores
- 900 GB/s NVLink C2C
- 120 GB LPDDR5X memory, 500 GB/s (CPU)
- 96 GB HBM3 memory, 4000 GB/s (GPU)

TDP superchip: 680 W

- GPU-to-GPU: NVLink 4, 300 GB/s per pair
- CPU-to-CPU: cNVLink, 200 GB/s per pair
- Cache coherence between all GPUs and CPUs
- 4 × NDR200 InfiniBand (200 Gbit/s)


JUPITER Cluster Node Design

- 2 × SiPearl Rhea1 CPU per node, each
 - 80 Neoverse-V1 cores
 - 256 GB DDR5 memory
 - 64 GB HBM memory
- Origins in European HPC ecosystem: EPI, EUPEX, and further projects
- To be deployed soon




JUPITER Network

- NVIDIA Mellanox InfiniBand NDR
- Dragonfly+ topology
 - 27 Dragonfly groups
 - Inside each group: full fat tree
- Adaptive routing, SHARPV3



JUPITER Network



- In group
 - 240 nodes
 - 5 racks
 - 15 L1 (lower), 16 L2 (upper) switches
 - Per rack: 48 nodes, 192 GH200

MDC Concrete Foundation; Andreas for scale



Delivery of first entry hall containers



Delivery of data hall



Delivery of
IT Module 2



Photo by Herwig Zilkner

*Installation of
Compute Blades
in IT Module 1*



Photo by Herwig Zilken / FZJ

MDC bird's eye view



Photo by Sascha Kieklau / FZJ



110V2




ABB

Inside of IT module

Photo by Sascha Kieklau / FZJ

JUPITER TOP500

- First entry to TOP500 June 2025 at #4 world, #1 EU
- 793 PFLOP/s HPL of 930 PFLOP/s th. peak
- Achieved during build-up phase of system; on 4650 nodes
- Not *really* an exascale system just yet



LLview job report of HPL job

JOINING FORCES




jupiter.fz-juelich.de



Benchmarks Overview

- JUPITER procurement: focus on applications
- 16 application benchmarks (4 de-selected for actual procurement)
Cross-section of domains and methods, 3× AI; including Amber, GROMACS
- 7 synthetic benchmarks
- Extensive [descriptions](#)
- Right: Patterns of 7 Dwarfs



Infrastructure

- Preparation system: **JUWELS Booster** (73 PFLOP/s peak, 44 PFLOP/s HPL), JURECA-DC (15 PFLOP/s, 9 PFLOP/s)
Dependencies: EasyBuild; versioned environment modules
- **JUBE** workflow environment for every benchmark, similar structure; implicit documentation; platform-independence through inheritance
→ Continuous Benchmarking with *exaCB*
- Extensive **description** (incl. guidelines, rules), similar structure
- **Git**, Git submodules for sources (if possible)
- **Management**: Teams with captains and domain scientists, meetings every 2 weeks, hackathons, scale days, *11 step program*, Gitlab issues

```
- name: systemParameter
  init_with: platform.xml
  parameter:
    - name: preprocess
      _: $modules
    - name: executable
      _: myapp
    - name: args_exec
      _: input.json
    - name: queue
      tag: "baseline|scaling_base|scaling"
      _: booster
    - name: queue
      tag: "exa_tiny|exa_small|exa_medium"
      _: largebooster
```


JUBE example

GrandTable Benchmark Details

- Languages, models, libraries
- Licenses
- References nodes Base, High-Scaling
- Memory variants
- Execution targets


	Application Features			Execution Targets		
	Benchmark Name	Progr. Language, [Libraries,]Prog. Models	Licence	Nodes Base	Nodes High-Scale N _{Mem Vars}	Module/ Device B _G B _S C _E M _{GC}
Application	Amber*	Fortran, CUDA	Custom	1		✓
	Arbor	C++, CUDA/HIP	BSD-3-Clause	8	642 ^{T,S,M,L}	✓
	Chroma-QCD	C++, QUDA, CUDA/HIP	JLab	8	512 ^{S,M,L}	✓
	GROMACS	C++, CUDA/SYCL	LGPLv2.1	3/128		✓
	ICON	Fortran/C, OpenACC/CUDA/HIP	BSD-3-Clause	120/300		✓
	JUQCS	Fortran, CUDA/OpenMP	None	8	512 ^{S,L}	✓
	nekRS	C++/C, OCCA, CUDA/HIP/SYCL	BSD-3-Clause	8	642 ^{S,M,L}	✓
	ParFlow*	C, Hypre, CUDA/HIP	LGPL	4		✓
	PICongPU	C++, Alpaka, CUDA/HIP	GPLv3+	8	640 ^{S,M,L}	✓
	Quantum Espresso	Fortran, ELPA, OpenACC/CUF	GPL	8		✓
	SOMA*	C, OpenACC	LGPL	8		✓
	MMoCLIP	Python, PyTorch, CUDA/ROCm	MIT	8		✓
	Megatron-LM	Python, PyTorch/Apex, CUDA/ROCm	BSD-3-Clause	96		✓
	ResNet*	Python, TensorFlow, CUDA/ROCm	Apache-2.0	10		✓
	DynQCD	C, OpenMP	None	8		✓
Synthetic	NASTJA	C++, MPI	MPL-2.0	8		✓
	Graph500	C, MPI	MIT	4/16/all		✓
	HPCG	C++, OpenMP, CUDA/HIP	BSD-3-Clause	1/4/all		✓
	HPL	C, BLAS, OpenMP, CUDA/HIP	BSD-4-Clause	1/16/all		✓
	IOR	C, MPI	GPLv2	-/ > 64		✓
	LinkTest	C++, MPI/SIONlib	BSD-4-Clause+	all		✓
	OSU	C, MPI, CUDA	BSD	1/2		✓
	STREAM	C, CUDA/ROCm/OpenACC	Custom	1		✓
						✓
						✓

TCO Base Result Grid



Execution on reference number of nodes ($x = 1$) resulting in reference timing ($y = 1$); absolute numbers super-imposed; strong scaling to $0.5\times - 2\times$


TCO Base Result Grid




Execution on reference number of nodes ($x = 1$) resulting in reference timing ($y = 1$); absolute numbers super-imposed; strong scaling to $0.5\times - 2\times$

JUREAP Overview

- Competitive early access program since 2024
- Open phase: 130 proposals
3× GROMACS
- Various domains, methods (simulation, AI)
- Preparations on JEDI, JUWELS, other EU systems
- Reproducibility, continuous integration (*exaCB*)
- Lighthouse application: peer-reviewed selection,
15  + 18  apps



Amber on GH200



- Comparing STMV production NPT 4FS benchmark (Exacct)
- JUPITER Benchmark Suite: STMV NVE
- JSC: different TDPs

Conclusion

Conclusions

- JUPITER: First European exascale system
 - 24 000 GH200 with plenty of performance
 - TOP500 at ISC25 last week
 - No exascale just yet
- Procurement with focus on applications (also Amber, GROMACS)
- First users using JUPITER since 6 days!

Conclusions

- JUPITER: First European exascale system
 - 24 000 GH200 with plenty of performance
 - TOP500 at ISC25 last year
 - No exascale just yet
- Procurement with focus on open source (e.g. LAMMPS, Amber, GROMACS)
- First users using JUPITER since 6 days!

**Thank you
for your attention!**

a.herten@fz-juelich.de

JUPITER

The Arrival of
Exascale in Europe

fz-juelich.de/jupiter | [#exa_jupiter](https://twitter.com/#!/exa_jupiter)



Funding Agencies:



Ministry of Culture and Science
of the State of
North Rhine-Westphalia



JOINING FORCES



EuroHPC
Joint Undertaking



Bundesministerium
für Forschung, Technologie
und Raumfahrt

Ministerium für
Kultur und Wissenschaft
des Landes Nordrhein-Westfalen



GCS
Gauss Centre for Supercomputing

JÜLICH
Forschungszentrum



EVIDEN



IBM

fz-juelich.de/jupiter

Appendix

Application Descriptions

TCO

- Amber** Molecular dynamics; STMV use-case; single node
- ParFlow** Hydrology; ClayL use-case
- SOMA** Polymer simulation; Monte-Carlo
- ResNet** AI: Computer vision; TensorFlow
- DynQCD** Particle physics; CPU-only
- GROMACS** Molecular dynamics: GluCL, STMV use-cases; multi-node
- ICON** Climate simulation; atmosphere with R02B09, R02B10; many nodes
- Megatron-LM** AI: LLM; PyTorch
- MMoCLIP** AI: Mixed-Modal (text, image); PyTorch
- Quantum ESPRESSO** Electronic structure; Car-Parrinello test-case
- NASStJA** Biology; CPU-only, MPI-only


TCO & High-Scaling

- Arbor** Neuroscience; busyring benchmark
- Chroma** Particle physics; hybrid-Monte-Carlo test; QUDA with JIT; max 512 nodes
- JUQCS** Quantum computer simulator; gate-based simulation; communication-heavy; max 512 nodes; Cluster-Booster version (MSA)
- nekRS** Fluid dynamics; Rayleigh-Bénard convection use-case
- PICongPU** Plasma physics; Kelvin-Helmholtz instability use-case

JUPITER Application Benchmarks

- JUPITER: Largest procurement to date
- >18 months of work
- >30 people involved
- 1(-3) associated people (*captains*) per benchmark
- Meetings every two weeks
- Gitlab issue tracker, status tracker (11 points)

1. Source code available
2. Input data available
4. JUBE integration
11. Description, documentation



Lessons Learned

Performance models useful, even if simple

Domain decomposition scripts/rules for unknown system makeup

Intensive feedback for app devs

Verification is hard



Applications

Suite: resource-intensive
→ aim for short turn-around times during dev

Artificially limit benchmarks on prep system to mimic future system

Extensive, balanced execution rules/guidelines



Benchmarks

Multi-system procurement
→ benchmark balance 🤖

Collaboration, tools

Bias towards incremental update

Limiter: time, on *all* sides
→ reuse!



Procurement

Benchmark Suite Availability

- All benchmark workflows, descriptions, data released online
- **Code** GitHub jubench meta-repository, Zenodo meta-record
Individual repos: Arbor Amber Chroma-LQCD GROMACS ICON JUQCS Megatron-LM
MMoCLIP nekRS ParFlow PConGPU Quantum ESPRESSO ResNet SOMA DynQCD
(CPU) NASTJA (GPU) Graph500 HPCG HPL IOR LinkTest OSU Micro-Benchmarks
STREAM STREAM (GPU)
- **Paper** SC24 Proceedings, arXiv:2408.17211
Including extensive SC reproducibility appendix



GitHub



Proceedings