



JUPITER: FIRST INSIGHTS FROM AN EXASCALE SYSTEM

31 PARS WORKSHOP

17 September 2025 | Andreas Herten | Forschungszentrum Jülich, Jülich Supercomputing Centre

Content

Exascale Era

JUPITER

System Overview

Comparison

Pictures

Build-Up

Inauguration

Procurement

Framework

Benchmark Suite

Benchmarks Overview

Infrastructure

Benchmark Details

TCO Results

High-Scaling Results

JUREAP

Overview

exaCB

First JUPITER Results

Synthetic Benchmarks

Hopper Power Analysis

JUREAP & Early Users

Scaled Block Vecchia

CFD

LQCD

ICON

ECMWF

Soon: JUPITER AI Factory

Conclusion

The Exascale Era

- Exascale systems with GPUs:  Frontier^A,  Aurora^I;  El Capitan^A,  JUPITER^N
- Homogeneous heterogeneity: 1 system in TOP500 top 20 *not* with GPUs


The Exascale Era

- Exascale systems with GPUs:  Frontier^A,  Aurora^I;  El Capitan^A,  JUPITER^N
- Homogeneous heterogeneity: 1 system in TOP500 top 20 *not* with GPUs

The Exascale Era

- Exascale systems with GPUs: 🇺🇸 Frontier^A, 🇺🇸 Aurora^I; 🇺🇸 El Capitan^A, 🇪🇺 JUPITER^N
- Homogeneous heterogeneity: 1 system in TOP500 top 20 *not* with GPUs


GPU Cores in Top500 (SMs etc.)



The Exascale Era

- Exascale systems with GPUs: 🇺🇸 Frontier^A, 🇺🇸 Aurora^I; 🇺🇸 El Capitan^A, 🇪🇺 JUPITER^N
- Homogeneous heterogeneity: 1 system in TOP500 top 20 *not* with GPUs

Energy Efficiency of GPUs




⚡ 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

- CPU-based based design (no accelerators)
- Nodes with $2 \times$ SiPearl Rhea1 CPU, 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 *intra*

- 5 racks
- 240 nodes
- 15 L1 (lower), 16 L2 (upper) switches
- Per rack: 48 nodes, 192 GH200
- Link bandwidth from HCA to L1: 200 Gbit/s; from L1 to L2: 400 Gbit/s


■ Beyond group *inter*

- 16 links to each other group
- Link bandwidth 400 Gbit/s

Other Exascale Systems


Frontier

ORNL, USA: 1.2 EFLOP/s




Aurora

ANL, USA: 1 EFLOP/s



El Capitan

LLNL, USA: 1.7 EFLOP/s



Overview Exascale Systems

Unique features

Frontier AMD GPUs in GCD design (two chip dies), few cores

Aurora 6 Intel GPUs, HBM-equipped Intel CPUs

El Capitan AMD APU design, no *host* memory

JUPITER Superchip design, tight combination between CPU-GPU

Overview

System	Nodes	per Node					per GPU		
		TFLOP/s [§]	GPUs	CPUs	NICs	Cores	TFLOP/s	Cores	NICs
Frontier	9604	141	4	1	4	64	35	16	1
Aurora	10 624	95	6	2	8	112	16	19	1.3
El Capitan	11 136	156	4	4	4	96	39	24	1
JUPITER	~6000	170*	4	4	4	288	43	72	1

§: HPL value divided by number of nodes; for theoretical values, see [appendix](#)

*: For TOP500-reported HPL run with 4650 nodes

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



Thomas Lippert and the press



Photo by Jenö Gellinek / FZJ

Friedrich Merz in speech




Photo by Kurt Steinhausen / FZJ

Buttons! So many Buttons!



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



Preparations

Benchmarks: Applications

- Intensive benchmarking endeavor
- Many community applications

JEDI: Preparation system

- 48 nodes
($\frac{1}{5}$ DragonFly group)
- May 2024: #1 Green500
(72.7 GFLOP/(s W))


JUREAP: Research & Early Access Program

- > 130 participants
- 33 Lighthouse Applications selected (EU+DE)



JSC Workload


- JSC: HPC resources for Forschungszentrum campus, state (NRW), Germany, Europe
- Compute time through peer-review
- Heterogeneous workload
- Physics, climate, biology, chemistry, AI, ...
- **Goal:** Respect current & anticipated workload in procurements of new systems; incl. domains, methods, programming languages, profiles



*Programs / GPUh (2020)
JUWELS Booster*

JSC Workload

- JSC: HPC resources for Forschungszentrum campus, state (NRW), Germany, Europe
- Compute time through peer-review
- Heterogeneous workload
- Physics, climate, biology, chemistry, AI, ...
- **Goal:** Respect current & anticipated workload in procurements of new systems; incl. domains, methods, programming languages, profiles




Framework


- Procuring entity: EuroHPC JU
 - Hosting entity: JSC
 - Procurement: JSC, with support from EuroHPC, national ministries
 - Compute time allocations: GCS (Germany), JU (Europe)
 - 500 M€ project budget
- Strong investment from all sides → replicability, reproducibility, reusability R^3
 - Benchmarks
 - **Applications** (Base TCO, High-Scaling, MSA)
 - Synthetic
 - JUPITER Cluster  and Booster 

Benchmarks Overview

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




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 $0.5 \times - 2 \times$

High-Scaling Results

Weak-scaling
relative to
reference up to
642 nodes
(50 PFLOP/s th.)




High-Scaling Results



JUREAP Overview

- Competitive early access program since 2024
- Open phase: 130 proposals
- Various domains, methods (simulation, AI)
- Preparations on JEDI, JUWELS, other EU systems
- Reproducibility, continuous integration (*exaCB*)
- Lighthouse application: peer-reviewed selection, 15  + 18  apps
- Currently preparing intensively (*pre-production*) for imminent production runs
- Also: GCS AI competition





exacb.pages.jsc.fz-juelich.de

Features

- Continuous Benchmarking for HPC
 - Template-based declarative syntax
 - Reduced barrier to entry for CI/CB
 - Reproducible benchmarks (JUBE, ...)
 - Ease of sharing workflows
 - Pre-defined runs, experiments
- Currently extending, publishing

JUREAP default example:

```
include:  
- component: jureap/jube@v3.2  
  inputs:  
    prefix: "jedi.strong.tiny"  
    variant: "strong.tiny"  
    machine: "jedi"  
    queue: "all"  
    project: "cjsc"  
    budget: "zam"  
    jube_file: "simple.yaml"
```



exacb.pages.jsc.fz-juelich.de

Features

- Continuous Benchmarking for HPC
 - Template-based declarative syntax
 - Reduced barrier to entry for CI/CB
 - Reproducible benchmarks (JUBE, ...)
 - Ease of sharing workflows
 - Pre-defined runs, experiments
- Currently extending, publishing

Energy measurement example:


```
include:  
- component: jureap/energy@v3.2  
inputs:  
  prefix: "jedi.strong.tiny"  
  variant: "strong.tiny"  
  machine: "jedi"  
  queue: "all"  
  project: "cjsc"  
  budget: "zam"  
  jube_file: "simple.yaml"
```

First JUPITER Results

GH200/A100 Application Performance/Efficiency


- Test application: **MPTRAC**
- Lagrangian transport of particles
- Analysis of Energy-to-Solution for Grace, Hopper, GH200; A100
- Horizontal: Different GPU clock frequencies

➔ Energy-saving potential




Scaled Block Vecchia Approximation

- Investigation by KAUST, supported by JSC
- New algorithm for approximate modelling of simulations with GPUs
- Extensive tests on JUPITER
- Focus point: GPU energy efficiency
- Preprint on [arXiv \(2504.12004\)](https://arxiv.org/abs/2504.12004)



A100: 2M points


Preliminary data, in publication



GH200: 5M points


Heat Convection Simulation

- US-DE collaboration in CFD; software: nekRS
- Simulate dynamics heat plumes in Rayleigh–Bénard convection
- Large scale of simulation, needs JUPITER
- Interesting results how convection forms out small plumes
- Published in [PNAS 10.1073/pnas.2502972122](https://doi.org/10.1073/pnas.2502972122)




QCD Lattice Calculation of Tin


- FZJ work simulating Sn isotopes
- Fine lattice, many nodes of JUPITER
- First ab-initio calculation; insight: binding energy below extrapolated values
- Preprint on [arXiv \(2509.08579\)](https://arxiv.org/abs/2509.08579)



- European weather/climate model by Max Planck Institute for Meteorology and others
- Components for simulation of atmosphere, ocean, land, ...
- Nominated for Gordon Bell Prize for Climate Modelling
- Part of JUREAP
- Use full GH200 superchip (CPU, GPU) with different components
- Simulate full earth model on 1 km resolution

Strong/Weak Scaling - ICON Earth System on Alps/JUPITER






*ICON Simulation: Phytoplankton,
surface wind, air-sea CO₂ flux*



- European Centre for Medium-Range Weather Forecasts
- Intergovernmental organization for weather, climate research/forecasts
- Destination Earth: Digital twin for earth simulation, coordinated by ECMWF
- Nominated for Gordon Bell Prize for Climate Modelling
- Part of JUREAP
- Applications: IFS, AIFS, ...
- Utilizing significant portions of JUPITER for applications



Recently presented by Nils Wedi

Soon: JUPITER AI Factory

JUPITER AI FACTORY (JAIF)



Fraunhofer



hessian.AI

Associated
partners:




WEST AI
KI-Servicezentrum



KI BUNDESVERBAND



JAIF OVERVIEW

- AI Factory around JUPITER infrastructure
- One-Stop Shop for AI services
- Focus on key sector
- JUPITER extension: JARVIS
Cloud-based inference system (experimental)



Core Configuration
Future Call Modules

JUPITER
AI FACTORY



*Depending on procurements and available functionality

**This is a subset of the available software

Conclusion

Conclusions

- JUPITER: European exascale system, just inaugurated
 - 24 000 GH200 with plenty of performance
 - TOP500: #1 Europe, #4 world
- JUPITER Benchmark Suite with community involvement; sustainability
- JUREAP: Competitive program for early access of JUPITER
- First results from users coming in

Conclusions

- JUPITER: European exascale system, just inaugurated
 - 24 000 GH200 with plenty of performance
 - TOP500: #1 Europe #1
 - JUPITER Benchmark Suite
 - JUREAP: Competitive performance of JUPITER
 - First results from users coming in
- Thank you
for your attention!*
a.herten@fz-juelich.de
- ... involvement; sustainability

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



ParTec
MODULAR SUPERCOMPUTING

EVIDEN



SIPEARL

IBM

fz-juelich.de/jupiter

Appendix

GPU Performance Details

Vend.	System	Device	TDP	Compute Performance							Memory	
				64	64 ^M	32	32 ^M	16 ^M	8 ^M	8 ^S	Size	BW
			W	TFLOP/s							GB	TB/s
AMD	Frontier	MI250X	560	48	96	48	96	383			128	3.2
	El Capitan	MI300A	760	61	123	123	123	981	1961	3922	128	5.3
Intel	Aurora	PVC	600	52		52	419	839			128	3.2
NVIDIA	JUPITER	GH200	680	34	67	67	494	989	1979	3958	96	4

GPU Performance Details

Vend.	System	Device	TDP	Matrix Compute			Sparse+Matrix Compute				Memory	
				Compute Performance			Float Precision					
				64	64 ^M	32	32 ^M	16 ^M	8 ^M	8 ^S	Size	BW
				W	TFLOP/s					GB	TB/s	
AMD	Frontier	MI250X	560	48	96	48	96	383			128	3.2
	El Capitan	MI300A	760	61	123	123	123	981	1961	3922	128	5.3
Intel	Aurora	PVC	600	52		52	419	839			128	3.2
NVIDIA	JUPITER	GH200	680	34	67	67	494	989	1979	3958	96	4
		Superchip TDP					TF32					

Superchip TDP

TF32

Evaluation Target

Mainly: Total Cost of Ownership (TCO)

- Proposals ranked by *workload intensity (how much workload over system lifespan)*
- Also energy consumption respected (simplified)
- Master formula: calculate value for ranking
- Normalized metric/FOM: runtime
- Applications-based

Evaluation Target

Mainly: Total Cost of Ownership (TCO)

- Proposals ranked by *workload intensity (how much workload over system lifespan)*
- Also energy consumption respected (simplified)
- Master formula: calculate value for ranking
- Normalized metric/FOM: runtime
- Applications-based

New: High-Scaling Benchmarks

- For Exascale procurement → respect large-scaleness of system
- Run dedicated workload on entire system
- Applications-based

Evaluation Target

Mainly: Total Cost of Ownership (TCO)

- Proposals ranked by *workload intensity (how much workload over system lifespan)*
- Also energy consumption respected (simplified)
- Master formula: calculate value for ranking
- Normalized metric/FOM: runtime
- Applications-based

New: High-Scaling Benchmarks

- For Exascale procurement → respect large-scaleness of system
- Run dedicated workload on entire system
- Applications-based
- Method:

Full JUWELS Booster Workload



20× workload on full JUPITER Booster

- *Full*: 50 PFLOP/s vs. 1000 PFLOP/s th. peak
- Instructions/rules to determine workload
- Memory variants: (tiny,) small, medium, large



Evaluation Target

Mainly: Total Cost of Ownership (TCO)

- Proposals ranked by *workload intensity (how much workload over system lifespan)*
- Also energy consumption respected (simplified)
- Master formula: calculate value for ranking
- Normalized metric/FOM: runtime
- Applications-based

Also: Synthetic benchmarks

New: High-Scaling Benchmarks

- For Exascale procurement → respect large-scaleness of system
- Run dedicated workload on entire system
- Applications-based
- Method:

Full JUWELS Booster Workload



20× workload on full JUPITER Booster

- *Full*: 50 PFLOP/s vs. 1000 PFLOP/s th. peak
- Instructions/rules to determine workload
- Memory variants: (tiny,) small, medium, large



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 <i>N_{Mem Vars}</i>	Module/ Device 
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		✓
Synthetic	DynQCD	C, OpenMP	None	8		✓
	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		✓ ✓

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 <i>N_{Mem Vars}</i>	Module/ Device 
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		✓
Synthetic	DynQCD	C, OpenMP	None	8		✓
	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		✓ ✓

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 PIconGPU 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