JUPITER - THE ARRIVAL OF EXASCALE IN EUROPE

ETP4HPC Webinar

2024-06-14 I ANDREAS HERTEN, BENEDIKT VON ST. VIETH | JÜLICH SUPERCOMPUTING CENTRE

EuroHPC Joint Undertaking

Bundesministerium für Bildung und Forschung

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

GCS Gauss Centre for Supercomputing

JÜLICH Forschungszentrum Shaping Change

Member of the Helmholtz Association

WHAT TO EXPECT FROM TODAYS WEBINAR

- Describe the whole JUPITER *journey*
  - Design decisions
  - Procuring JUPITER
  - Current Status, including JEDI
- Users!
  - During the procurement
  - Recent research on Grace-Hopper
  - Future plans with JUREAP

JÜLICH SUPERCOMPUTING CENTRE

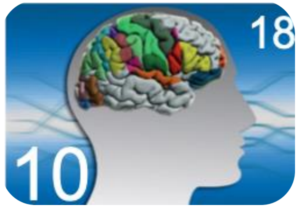

Develop the supercomputing facility **towards Exascale** (2024)

Complement the Exascale by the **quantum computer** facility JUNIQ

Develop **federated data infrastructures** and advanced data science

Provide **most innovative support** structures, tools, algorithms and methods

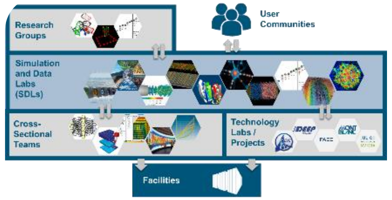

Educate a new generation of simulation and data science specialists

Staff:

- ~300 Total (~280 FTE)
- 210 Scientists
- 22 PhD Students (+ externals)
- 27 Students (Bachelor/Master)

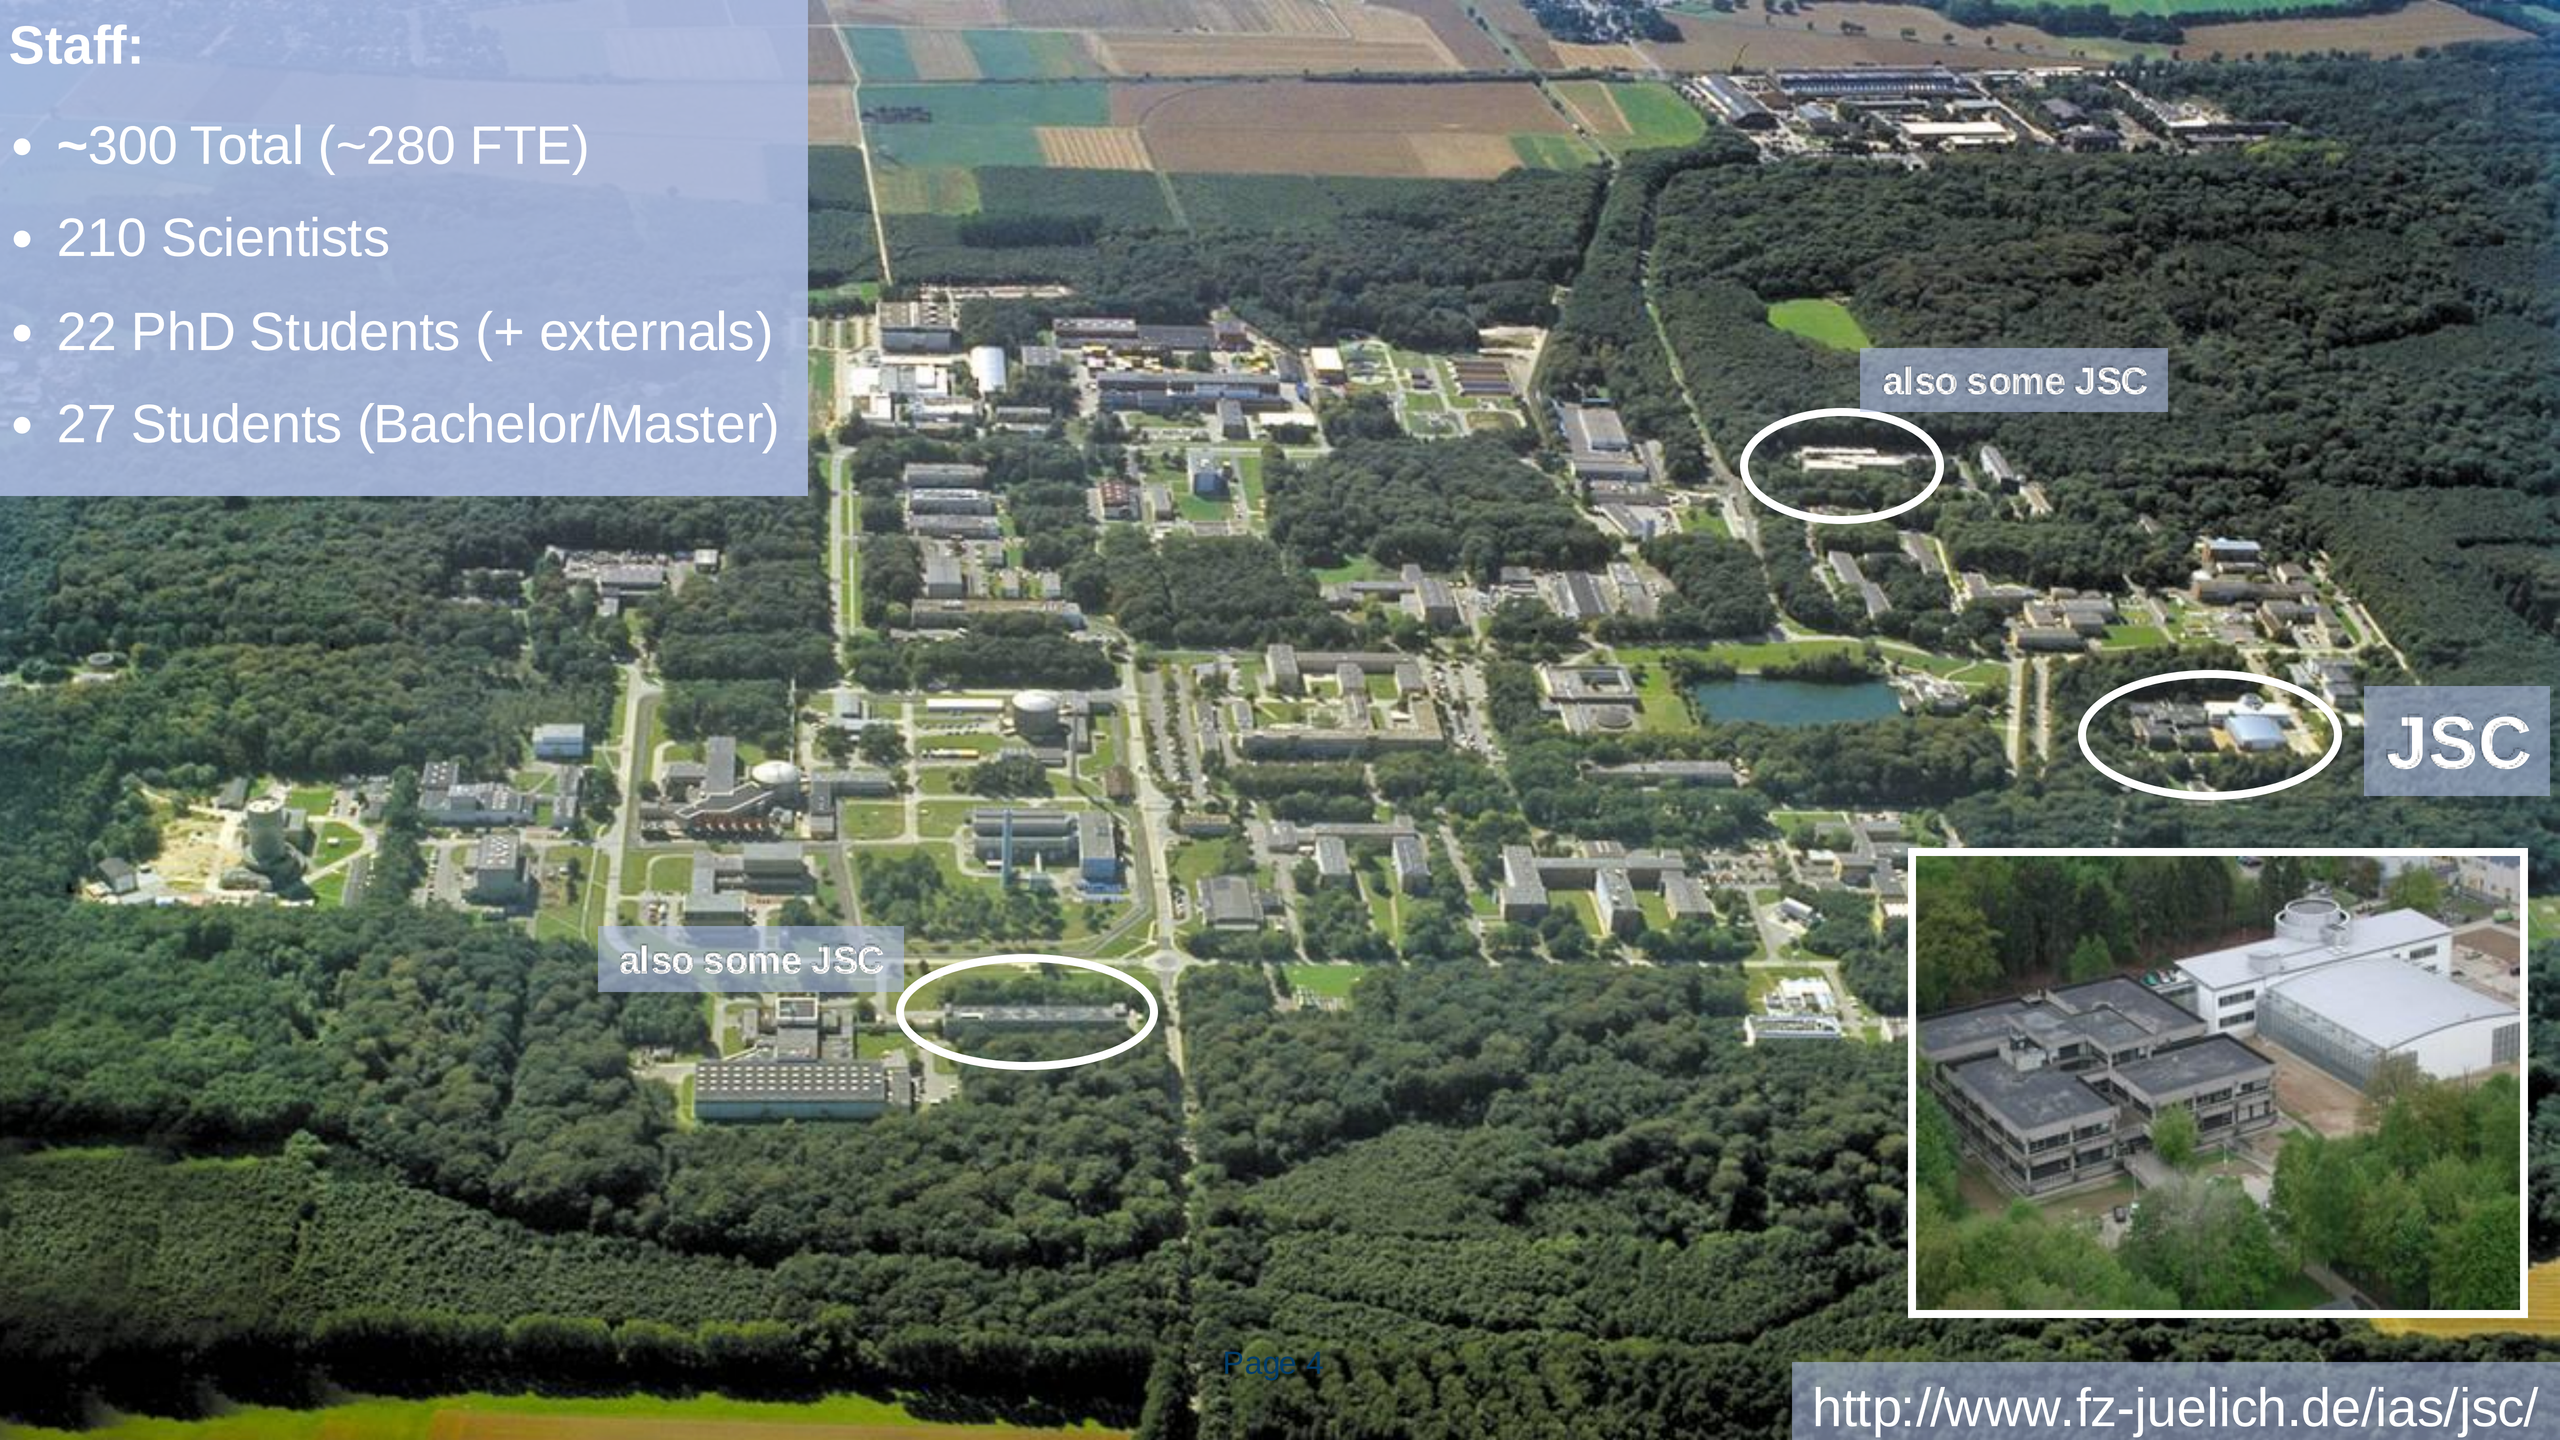

http://www.fz-juelich.de/ias/jsc/

DESIGN DECISIONS

THE PATHFINDER

Implementing the Modular Supercomputing Architecture

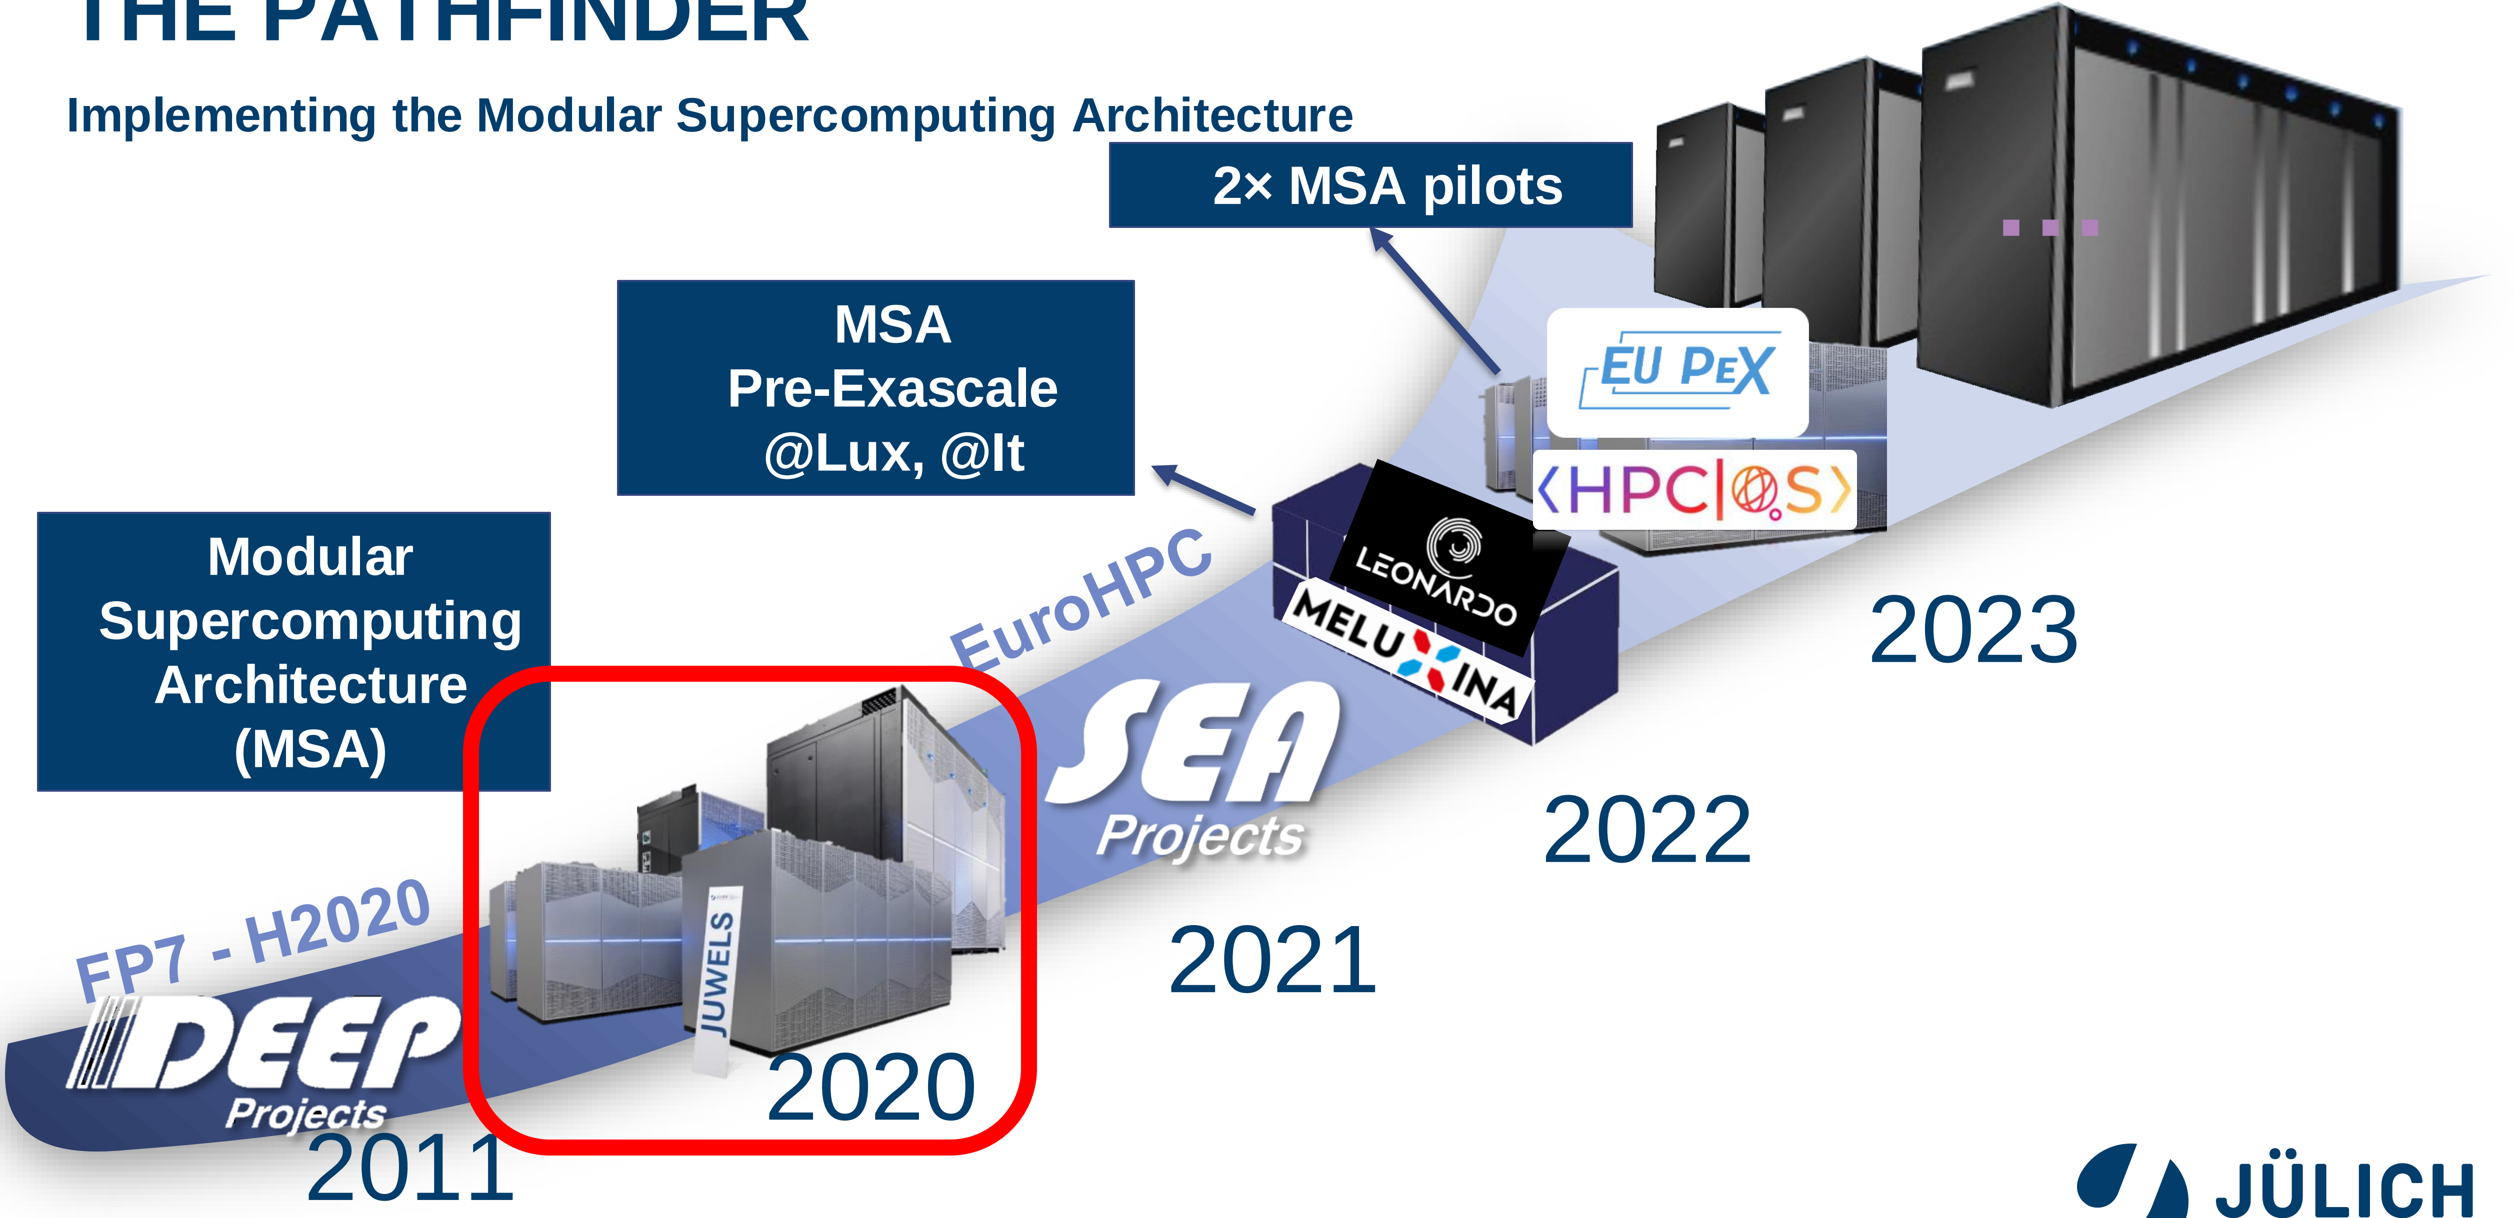

JUWELS Booster

#1 in TOP500 Europe (11/2020), #7 WW

#1 in Green500 among the top 100 in HPL

#5 HPCG500

#4 HPL-AI

PROCURING JUPITER

JUPITER – HOSTING ENTITY DECISION 15.06.2022

tagesschau

Startseite ▸ Wirtschaft ▸ Technologie ▸ Hochleistungs-Rechner: Supercomputer "Jupiter" kommt nach Jülich

Hochleistungs-Rechner

Supercomputer "Jupiter" k

Stand: 15.06.2022 16:43 Uhr

Das Forschungszentrum Jülich wird Standor
Exascale-Computers. "Jupiter" soll die Schal
Rechenoperationen in der Sekunde durchbr

SPIEGEL Netzwelt

»Jupiter«

Jülich bekommt Europas ersten Exascale-Supercomputer

Das Forschungszentrum Jülich bekommt für eine halbe Milliarde Euro einen neuen Vorzeigerechner. Er soll helfen, Fragen zum Klimawandel zu beantworten – mit mehr als einer Trillion Rechenoperationen pro Sekunde.

15.06.2022, 16.52 Uhr

SIGN IN The Register®

HPC

Germany to host Europe's first exascale supercomputer

Jupiter added to HPC solar system

Dan Robinson Thu 16 Jun 2022 // 07:33 UTC

3

Germany will be the host of the first publicly known European exascale supercomputer, along with four other EU sites getting smaller but still powerful systems, the European High Performance Computing Joint Undertaking (EuroHPC JU) announced this week.

Germany will be the home of Jupiter, the "Joint Undertaking Pioneer for Innovative and Transformative Exascale Research." It should be switched on next year in a specially designed building on the campus of the Forschungszentrum Jülich research centre and operated by the Jülich Supercomputing Centre (JSC), alongside the existing Juwels and Jureca supercomputers.

PREPARING FOR LAUNCH

- Mission planning
  - Preparing descriptions, conditions, requirements, evaluation
  - Regular meetings
  - Started already early in year
  - Location: **Earth**

- Target: **JUPITER**
  - Booster
  - Cluster
  - Storage
  - (Machine Hall)

READY FOR TAKE OFF

Competitive Dialogue - Descriptive Document

- Description of procurement procedure
- Overall budget, **273 M€**
- High-level description of targeted system
  - Implementing the MSA
  - Booster to achieve 1 EF
  - Cluster, preferably based on European IP
  - Flash storage module
  - Interconnect expectations
  - Login system sizing
  - System management

EuroHPC Joint Undertaking

European High Performance Computing Joint Undertaking

GENERAL INVITATION TO TENDER

EUROHPC/2023/CD/0001

Descriptive Document

Acquisition, delivery, installation and hardware and software maintenance of JUPITER Exascale Supercomputer for the European High Performance Computing Joint Undertaking (EuroHPC)

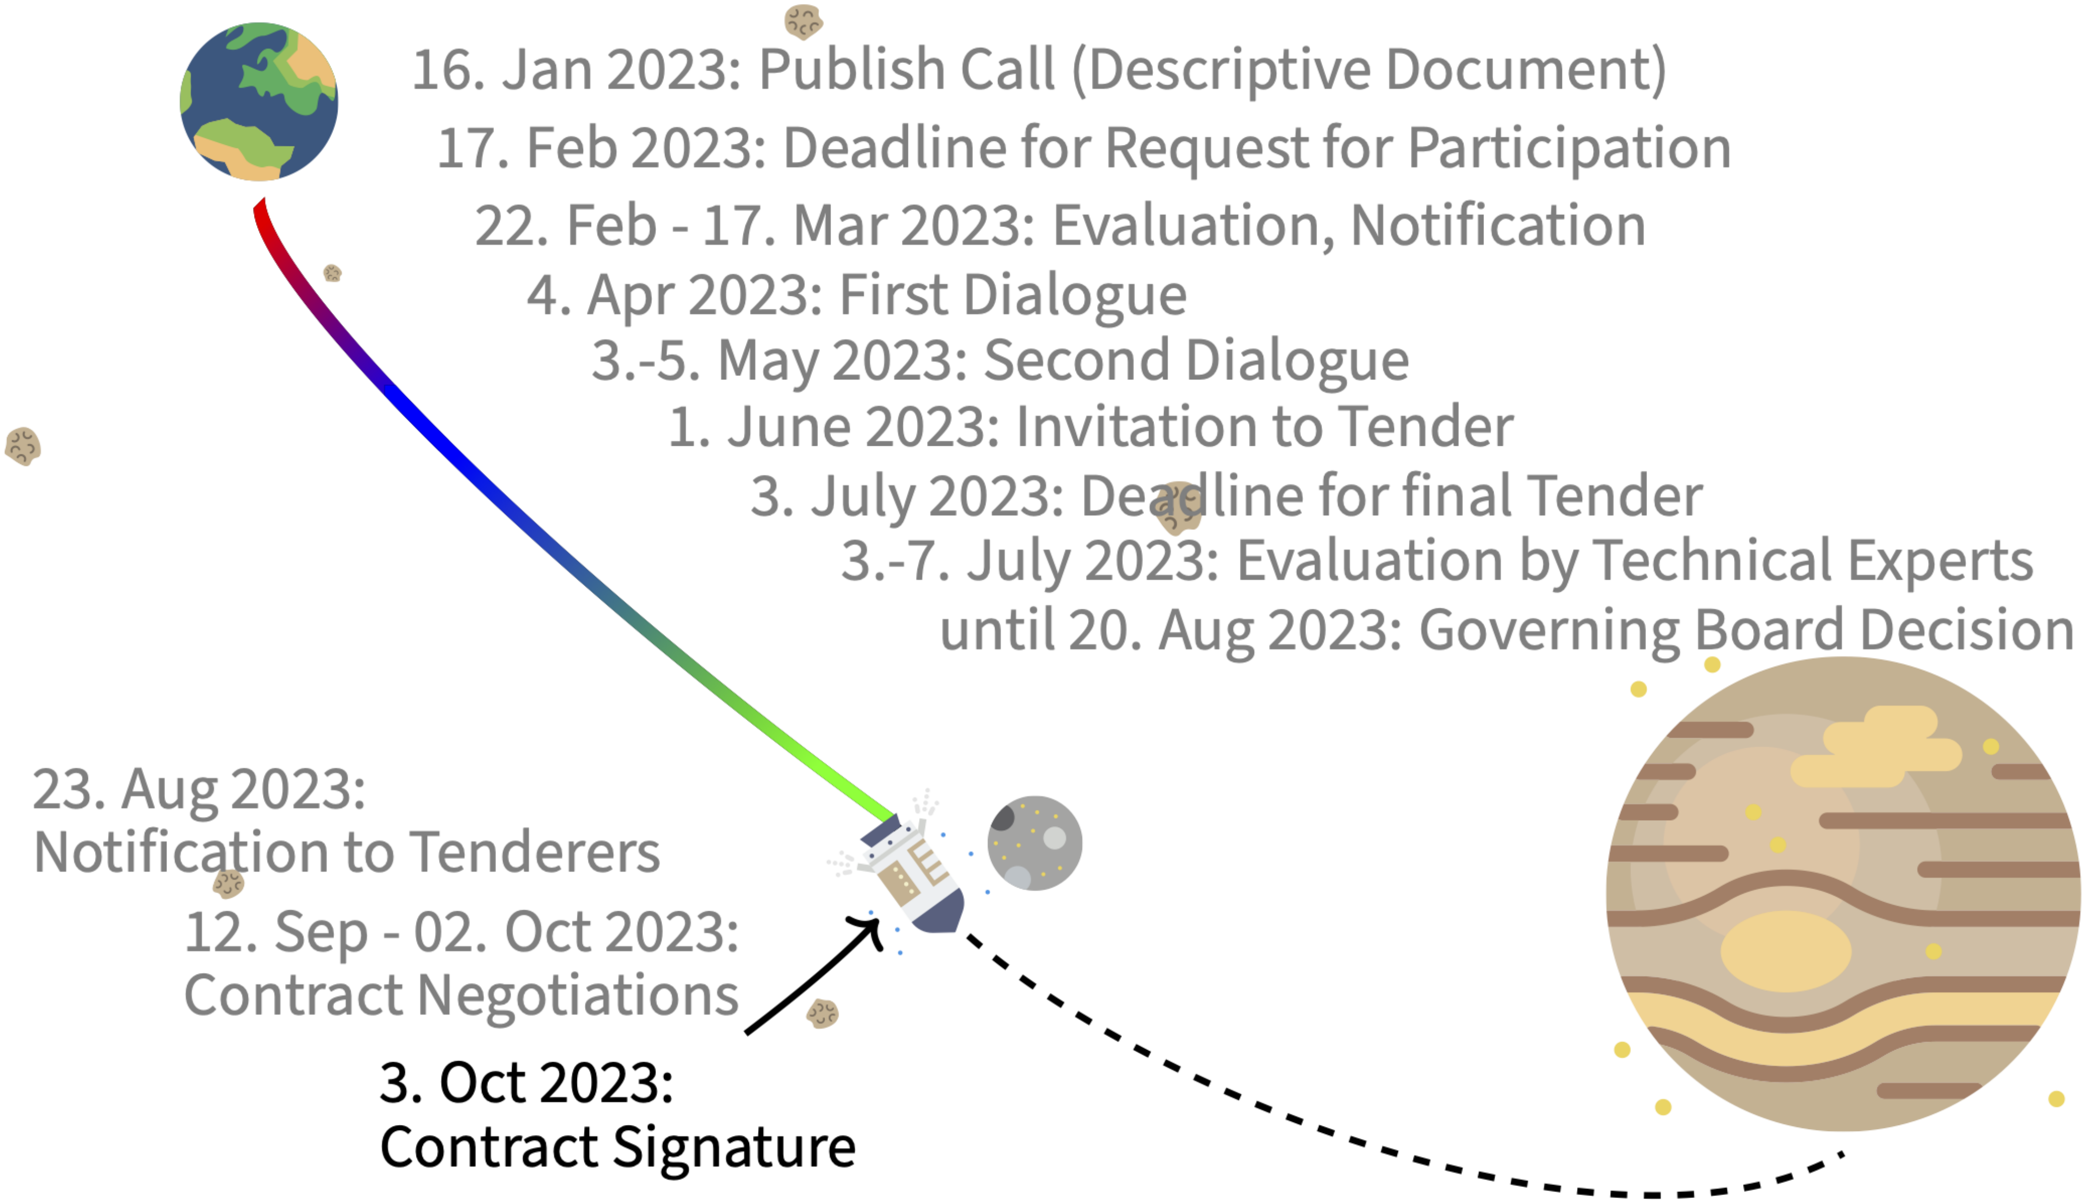
16. Jan 2023: Publish Call (Descriptive Document)
17. Feb 2023: Deadline for Request for Participation
22. Feb - 17. Mar 2023: Evaluation, Notification
4. Apr 2023: First Dialogue
3.-5. May 2023: Second Dialogue
1. June 2023: Invitation to Tender
3. July 2023: Deadline for final Tender
3.-7. July 2023: Evaluation by Technical Experts
until 20. Aug 2023: Governing Board Decision
23. Aug 2023: Notification to Tenderers
12. Sep - 02. Oct 2023: Contract Negotiations
3. Oct 2023: Contract Signature

JUPITER CONTRACT ANNOUNCEMENT 3.10.2023

HPC wire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Topics
- Sectors
- Exascale
- Specials
- Resource Library
- Podcast
- Events
- Solution Channels
- Job Bank
- About
- Subscribe

October 4, 2023

The configuration of Europe's first exascale supercomputer, Jupiter, has been finalized, and it is a win for Nvidia and a disappointment for x86 chip vendors Intel and AMD. The Jupiter supercomputer, which will cost €273 million to build, will pair SiPearl's Rhea processor, which is based on ARM architecture, with accelerator technology from Nvidia.

The supercomputer is being built by the European High-Performance Computing Joint Undertaking (EuroHPC JU) and a consortium including Eviden and ParTec. Eviden is an Atos business focusing on advanced computing initiatives tha

The Jülich Supercompu
(*Correction: Jülich Supe*
about 600 km or 375 mi

Specifically, the superco
CPUs, and the initial co
supercomputers in the w
on ARM.

That is a big disappoint
invest €33 billion to buil
development initiatives
leaders in a bid to get m

Jülich's fastest system,
November 2021 and is
Europe are the third-ran
performance of 309 peta

Off The Wire Industry Headlines

October 13, 2023

- Coherent File Format Accelerates Time-to-Solution with OpenFOAM
- HealthyCloud Project Unveils Roadmap to Maximize Impact of Health Data and Research Across Europe
- NCSA Welcomes 2023-24 Fellows
- Berkeley Lab CS Area to Share Computing Expertise at SC23

October 12, 2023

- Samsung Electronics to Host AI Forum 2023 Highlighting AI and Computer Engineering Innovation
- PacBio Announces Complete Computational Workflow for Human Whole Genome Sequencing Data Analysis
- SiFive Announces Differentiated Solutions for Generative AI and ML Applications
- EQTC 2023: Europe's Quantum Sector to Showcase Successes and Its Roadmap for Global Leadership
- EuroHPC JU Announces Procurement Call for Upgrading Discoverer Supercomputer
- Los Alamos Partners with AirMettle for Efficient In-Storage Data Analysis
- Caltech Researchers Demonstrate Quantum Eraser to Combat Erasure Errors in Quantum Systems

THENEXTPLATFORM

HOME | COMPUTE | STORE | CONNECT | CONTROL | CODE | AI | HPC | ENTERPRISE | HYPERSCALE | CLOUD

LATEST › Intel To Set Its FPGA Unit Free To Pursue Its Own Path ‣ COMPUTE — Search ...

HOME › HPC › Details Emerge On Europe's First Exascale Supercomputer

DETAILS EMERGE ON EUROPE'S FIRST EXASCALE SUPERCOMPUTER

October 5, 2023 Timothy Prickett Morgan

me details are emerging on Europe's first exascale system, codenamed "Jupiter" and to be installed at the ich Supercomputing Center in Germany in 2024. There has been a lot of speculation about what Jupiter ll include for its compute engines and networking and who will build and maintain the system. We now ow some of this and can infer some more from the statements that were made by the organizations rticipating in the Jupiter effort.

June 2022, the Forschungszentrum Jülich in Germany, which has played host to many supercomputers ice it was founded in 1987, was chosen to host the first of three European exascale-class percomputers to be funded through the EuroHPC Joint Undertaking and through the European national d state governments countries who are essentially paying to make sure these HPC and AI clusters are here they want them. With Germany having the largest economy in Europe and being a heavy user of C thanks to its manufacturing focus, Jülich was the obvious place to park the first machine in Europe to eak the exaflops barrier.

at barrier is as much an economic one as it is a technical one. The six-year budget for Jupiter weighs in at 00 million, which is around $526.1 million at current exchange rates between the US dollar and the ropean euro. That is in the same ballpark price as what the "Frontier" exascale machine at Oak Ridge tional Laboratory and the "El Capitan" machine that is being installed right now at Lawrence Livermore tional Laboratory – both of which are based on a combination of AMD CPUs and GPUs and Hewlett ckard Enterprise's Slingshot variant of Ethernet with HPE as the prime contractor.

erybody knows that Jupiter was going to use SiPearl's first generation Arm processor based on the overse "Zeus" V1 core from Arm Ltd, which is codenamed "Rhea" by SiPearl and which is appropriate

SIGN IN / UP — The Register®

HPC

Atos subsidiary Eviden scores contract win in Europe's first exascale system

$526M Jupiter set to rule EU's tech orbit by 2024

Dan Robinson — Wed 4 Oct 2023 // 16:45 UTC

The EU's supercomputing initiative, the European High Performance Computing Joint Undertaking (EuroHPC JU), has awarded a procurement contract for Europe's first exascale system, with installation due to start in early 2024.

Known as Jupiter (Joint Undertaking Pioneer for Innovative and Transformative Exascale Research), the system was announced last year followed by a call for tender in January of this year.

EuroHPC JU said the procurement contract for Jupiter was awarded to a consortium comprising of Eviden, the professional services side of French IT giant Atos, and ParTec, a German supercomputing hardware company.

The project is expected to have a total cost of €273 million ($287 million) covering the build, delivery, installation, and maintenance of Jupiter, according to the EuroHPC.

However, Eviden put the overall project cost at €500 million ($526 million), saying that this is the figure for the entire project, including the system manufacturing and its

JUPITER – OVERALL TIMELINE

Bundesministerium für Bildung und Forschung

EuroHPC Joint Undertaking

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

- 17.12.2021: Call for Expression of Interest (EoI) for Hosting Entity
- 14.02.2022: Deadline EoI Submission
- Q2 2022: Hearings & Hosting site decision and announcement
- Q1-Q3 2023: Procurement (Competitive Dialogue)
- **03.10.2023: Contract Signature**
- **Q3/Q4 2023: Installation Planning**
- Q1/2024: Start installation of JUPITER
- Beg. of 2025: Put in operation JUPITER

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Digital Europe programme, as well as by Germany through the BMBF and the MKW.

JUPITER ARCHITECTURE

DISCOVERING JUPITER

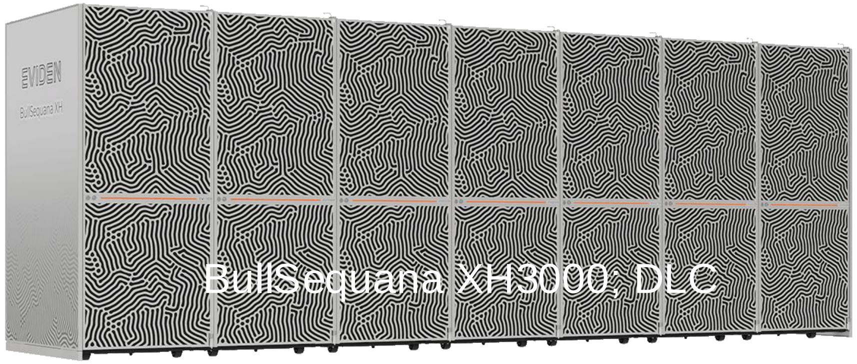

- First Exascale system in Europe (HPL); modular system
- Procured/funded by: EuroHPC JU, BMBF/NRW-MKW
  - Contract signed end of 2023
  - Installation starting soon
- JUPITER Booster: High scalablibty; 1 EFLOP/s HPL, >70 EFLOP/s FP8
  JUPITER Cluster: High versatility; 0.5 B/FLOP balance
- Network: InfiniBand NDR; Storage: 20 PB NVMe, 200 PB HDD
- Deployed in Modular Datacenter
- Building on: MSA (JUWELS); DEEP, EPI; ThunderX2, Ampere; …
- About 1.936.000 Arm cores

EVIDEN

ParTec MODULAR SUPERCOMPUTING

NVIDIA

SIPEARL

JUPITER MODULES

JUPITER Booster

- Node design
  - ~6000 nodes
  - 4× NVIDIA CG1 per node
- CG1: Grace-Hopper
  - 72 Arm Neoverse V2 cores (4×128b SVE2); 120 GB LPDDR5
  - H100 (132 SMs); 96 GB HBM3
  - NVLink C2C (900 GB/s)

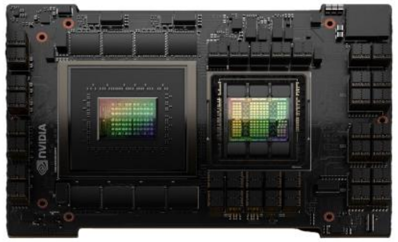

JUPITER Cluster

- Node design
  - ~1300 nodes
  - 2× SiPearl Rhea1 per node
- Rhea1
  - 80 Arm Neoverse V1 cores (2×256b SVE)
  - 256 GB DDR5, 64 GB HBM2e

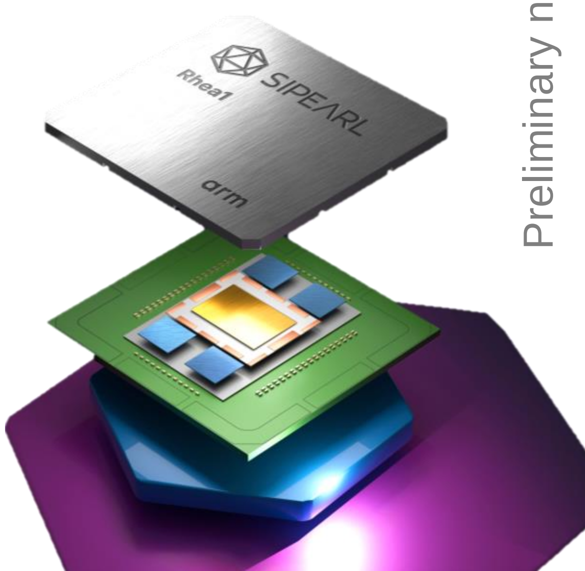

Preliminary numbers, might change during installation

BOOSTER NODE DESIGN

Node Specs

- 4× NVIDIA **Grace-Hopper** in SXM5 Board (4× 680W)
- 4× NVIDIA InfiniBand NDR200
- 480 GB LPDDR5X / 360 GB HBM3 (usable)
- NVLink 4
  - GPU-GPU 150 GB/s per dir
- Links: CPU-GPU 450 GB/s per dir, CPU-CPU 100 GB/s per dir

CPU Specs

- ARM Neoverse V2
  - SVE2/NEON (4x 128 bit vector op)
- 72 cores @ ~2.4GHz (~3.2 GHz turbo)
- 120 GB LPDDR5X (8 channels)
  - *≥450 GB/s*
  - *~150 ns latency*

GPU Specs

- H100
  - ~50 TFLOP/s (HPL single GPU)
- 96 GB HBM3
  - ~4 TB/s

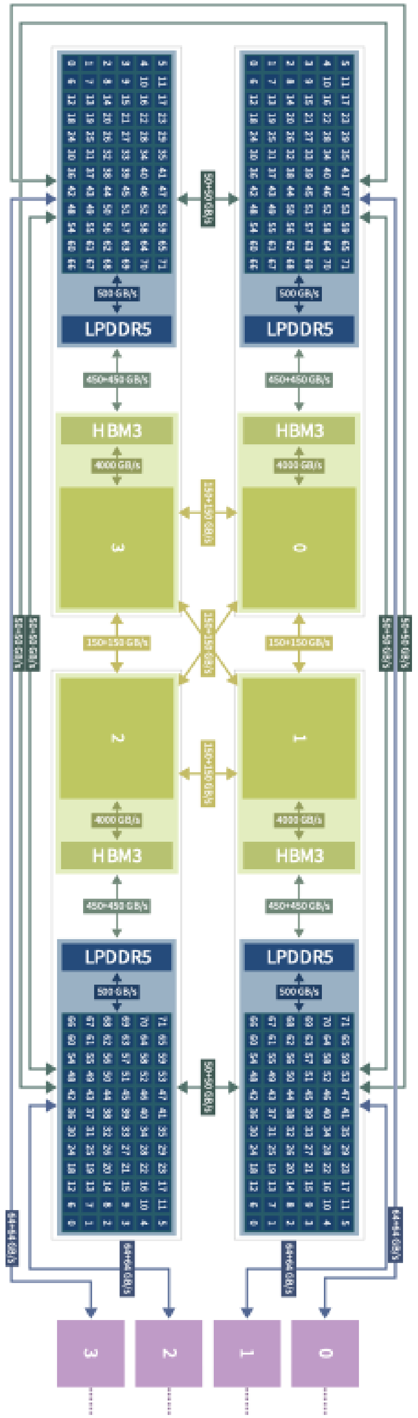

JUPITER – LOGIN/VISUALIZATION

Login Partition and Visualization Capabilities

- Login Nodes
  - Booster: 12 nodes, 1× CG1
  - Cluster: 5 nodes, 2× Rhea1
- Visualization Nodes
  - Booster: 3 nodes, 1× CG1
  - Cluster: 3 nodes, 2× Rhea1 and 2× NVIDIA A40
- 2× 100 Gbit Ethernet for external connectivity

JUPITER – INTERCONNECT

One Network to Rule Them All

- NVIDIA Mellanox InfiniBand NDR/NDR200
  - NVIDIA Quantum-2 switches
  - NVIDIA Connect-X7 HCAs
- Dragonfly+ topology
  - 27 Dragonfly groups
  - Within each group: full fat tree
- Approximately 51000 links, 102000 logical ports, 25400 endpoints, 867 switches
- Adaptive Routing
- In-network processing on switch level (SHARPv3), tentatively

JUPITER – INTERCONNECT

One Network to Rule Them All

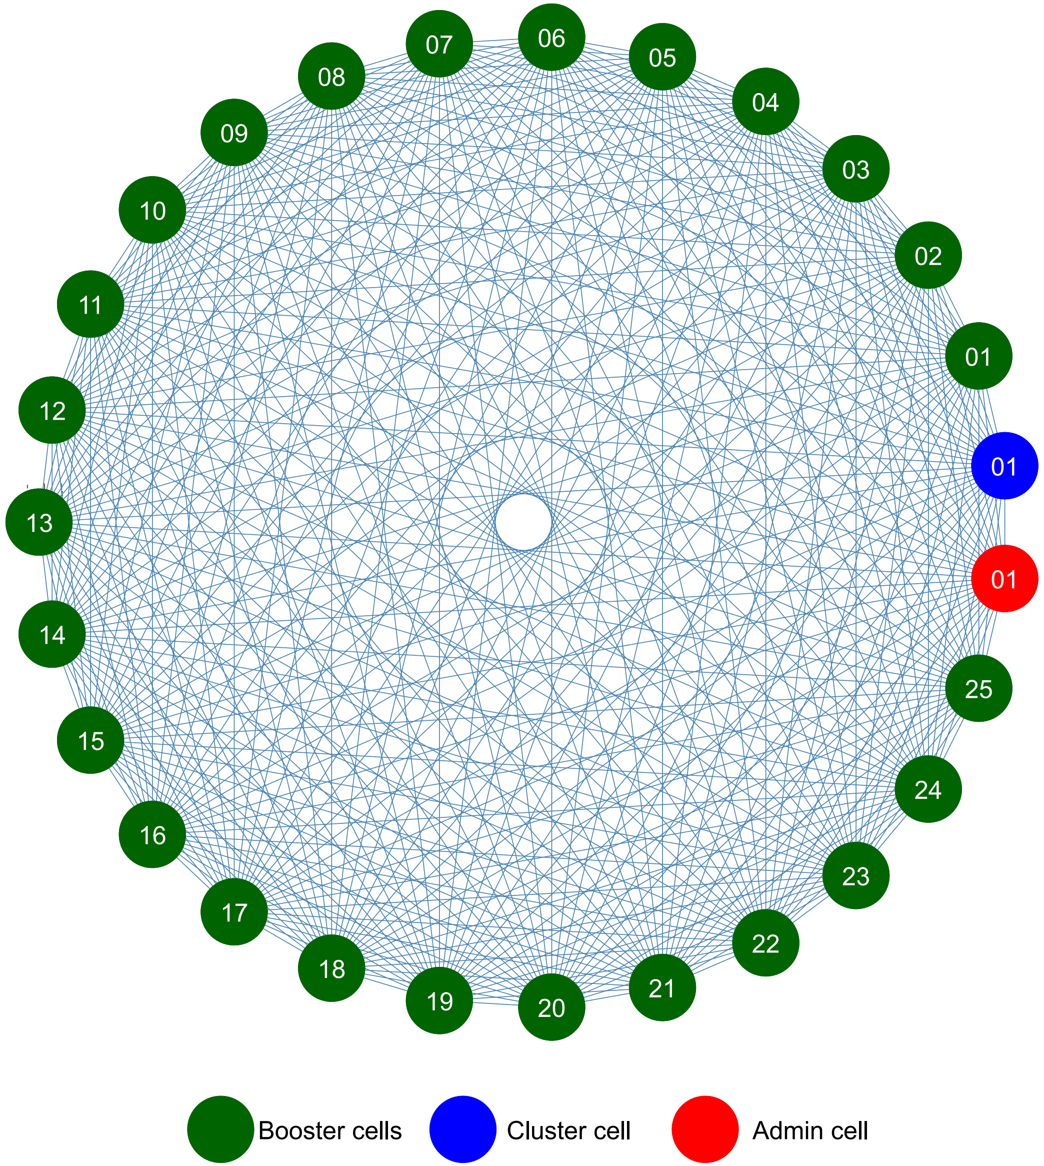

JUPITER – STORAGE (SCRATCH)

ParTec MODULAR SUPERCOMPUTING IBM

- Gross Capacity: 29 PB; Net Capacity: 21 PB
- Bandwidth: 2.1 TB/s Write, 3.1 TB/s Read
- 20× IBM SSS6000 Building Blocks (40 servers)
  - 2× NDR400 per server
  - 48× 30 TB NVMe drives per block
  - IBM Storage Scale (aka Spectrum Scale/GPFS)
- Manager and Datamover Nodes
- Exclusive for JUPITER
  - Integrated into InfiniBand fabric

Preliminary numbers, might change during installation

JUPITER – STORAGE (EXASTORE)

In kind contribution from JSC, not part of the JUPITER procurement

- Gross Capacity: 308 PB; Net Capacity: 210 PB
- Bandwidth: 1.1 TB/s Write, 1.4 TB/s Read
- 22× IBM SSS6000 Building Blocks (44 servers)
  - 2× NDR200 per server
  - 7× JBOD enclosures, each with 91x 22 TB Spinning Disks per block
  - IBM Storage Scale (aka Spectrum Scale/GPFS)
- Manager and Datamover Nodes
- Exclusive for JUPITER
  - Integrated into InfiniBand fabric

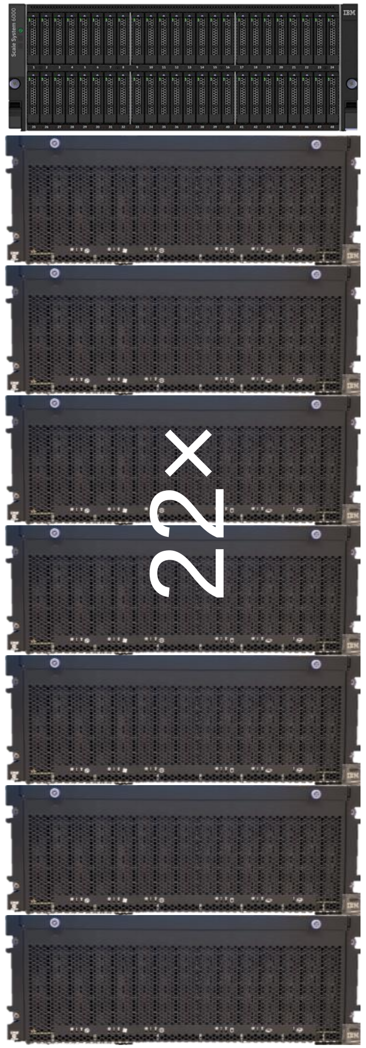

Preliminary numbers, might change during installation

JUPITER MANAGEMENT STACK

”Power is nothing without control"

- Eviden SMC xScale
- ParaStation Modulo
    - Resource management
    - ParaStation MPI
- Ansible as provisioning system
- SLURM as scheduler
- EasyBuild as scientific software package management
- RedHat Enterprise Linux 9


OPTIMUS MAXIMUS
(AKA MODULAR HPC DATACENTRE OR MDC)

DATA CENTER FOR JUPITER

JÜLICH
Forschungszentrum

DATA CENTER FOR JUPITER

JÜLICH
Forschungszentrum

CAMPUS INFRASTRUCTURE PREPARATIONS

- Upgrade of campus power supply 110 kV / 35 kV from 2× 40 MVA to 2× 60-80 MVA
- Upgrade of 110 kV power line to next electric power transformation substation
- New powerline 2× 35 kV from campus power transformation substation to Exascale site
- Construction of new weather tower further away from Exascale site
- Construction of concrete slab 85 m × 42 m × 0.5 m
- Connection to campus cooling loop, river water, potable water, process water, wastewater, network
- Connection to heating network for heat reusage

Preliminary numbers, might change during installation

CAMPUS ELECTRIC POWER TRANSFORMER SUBSTATION

- Upgrade of transformers 110 kV / 35 kV from 2 × 40 MVA to 2 × 60-80 MVA

CAMPUS ELECTRIC POWER TRANSFORMER SUBSTATION

- Upgrade of 110 kV power line to next
- electric power transformation substation

EXASCALE SITE

Construction of concrete slab 85 m × 42 m × 0.5 m

MODULAR DATA CENTER FOR JUPITER

EVIDEN
an atos business

MODULAR DATA CENTER FOR JUPITER

MDC provided by Eviden

- Base area: 2300 m2
- 7 IT modules (2 containers each)
  - Each container can be equipped with 10 Sequana XH3000 racks
- Data hall (4 containers) for standard racks with RDHx in 4 rows, electrical power max. 1 MW
- 15 energy supply modules with transformer 35 kV / 400 V, 2.5 MVA
  - 14 IT containers + data hall
- Logistic part with entrance area, control room, workshop, warehouse, restroom
- 1 glycol free hybrid cooling unit on the roofs for each IT container ~ 36°C / 48°C
- Redundancy by modules

Preliminary numbers, might change during installation

MODULAR DATA CENTER FOR JUPITER

EVIDEN
an atos business

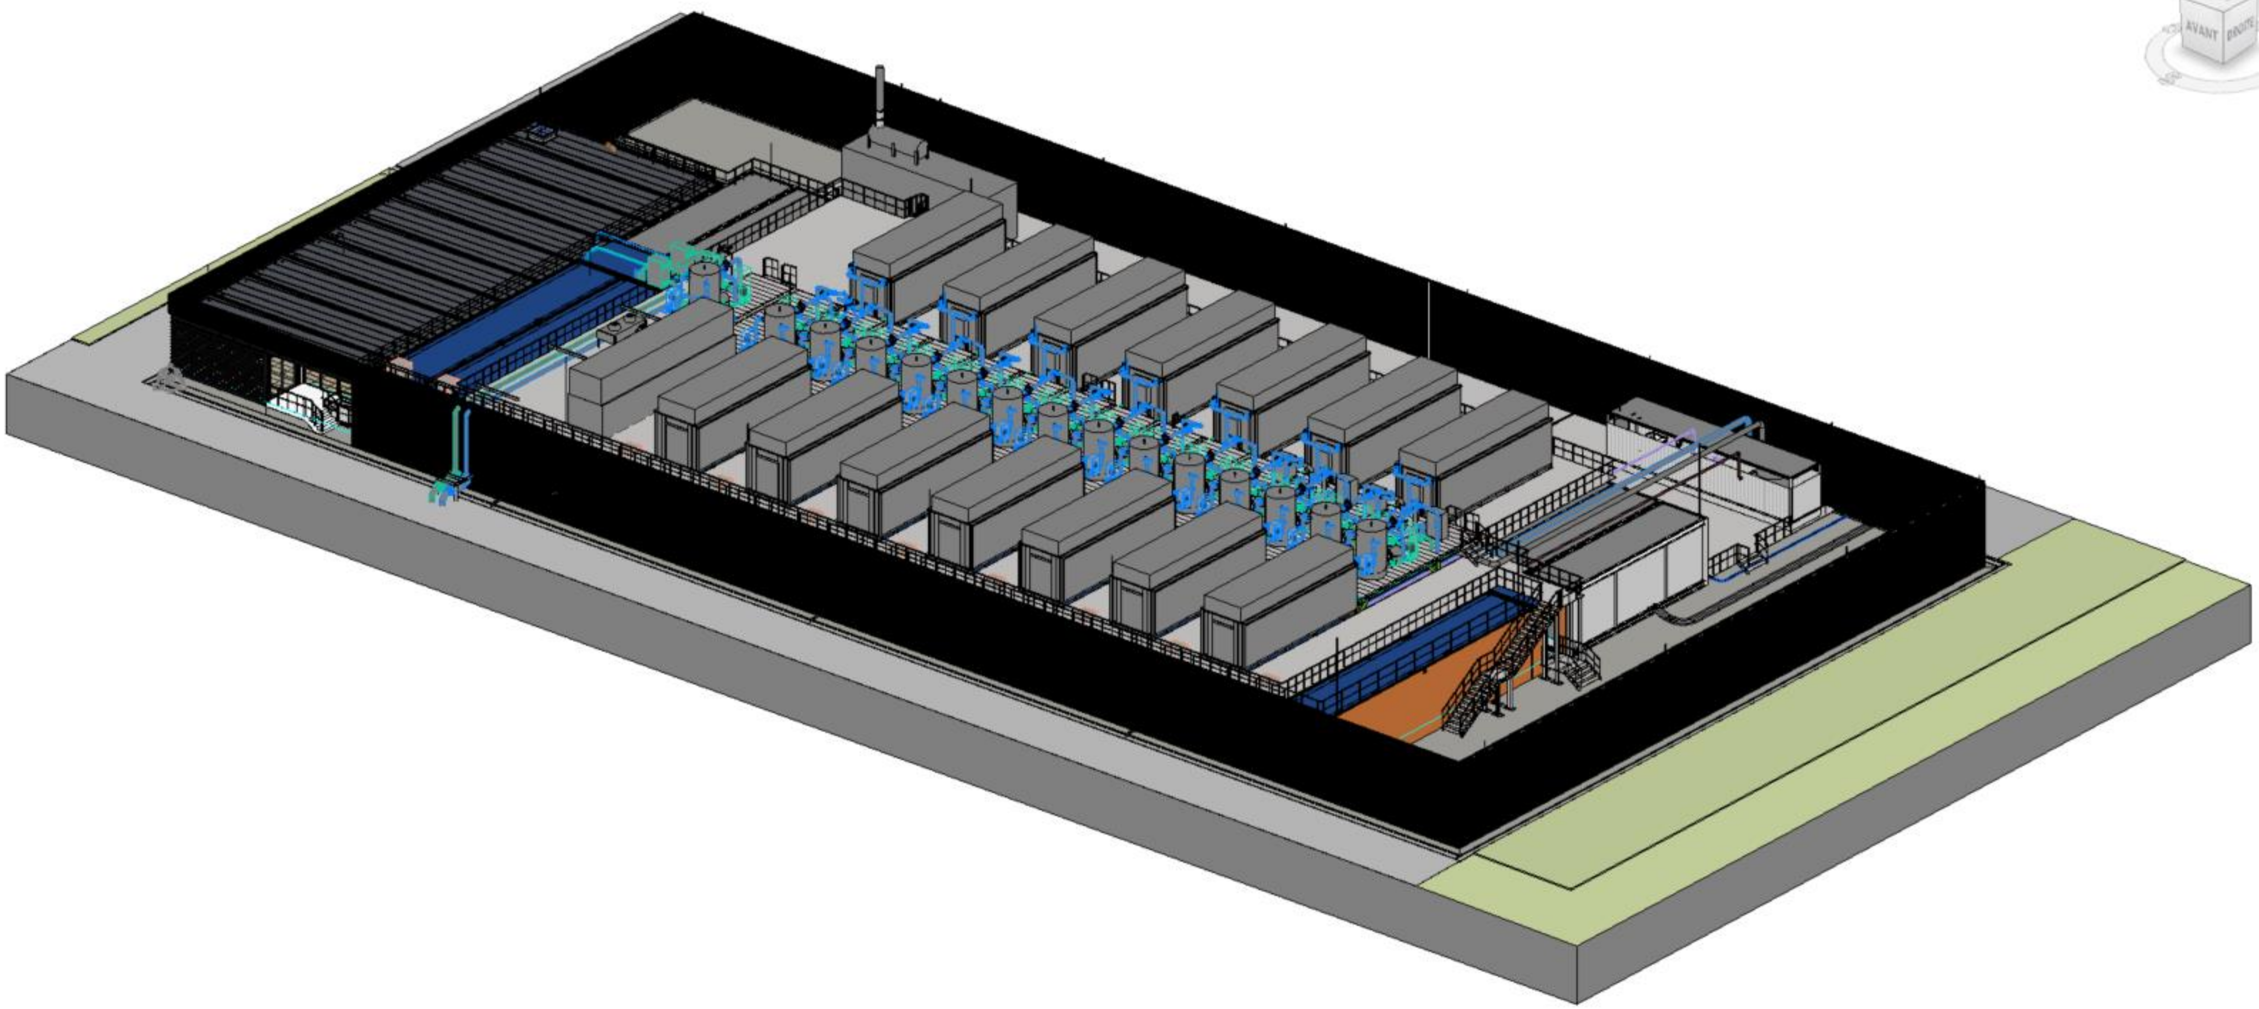

MODULAR DATA CENTER FOR JUPITER

EVIDEN
an atos business

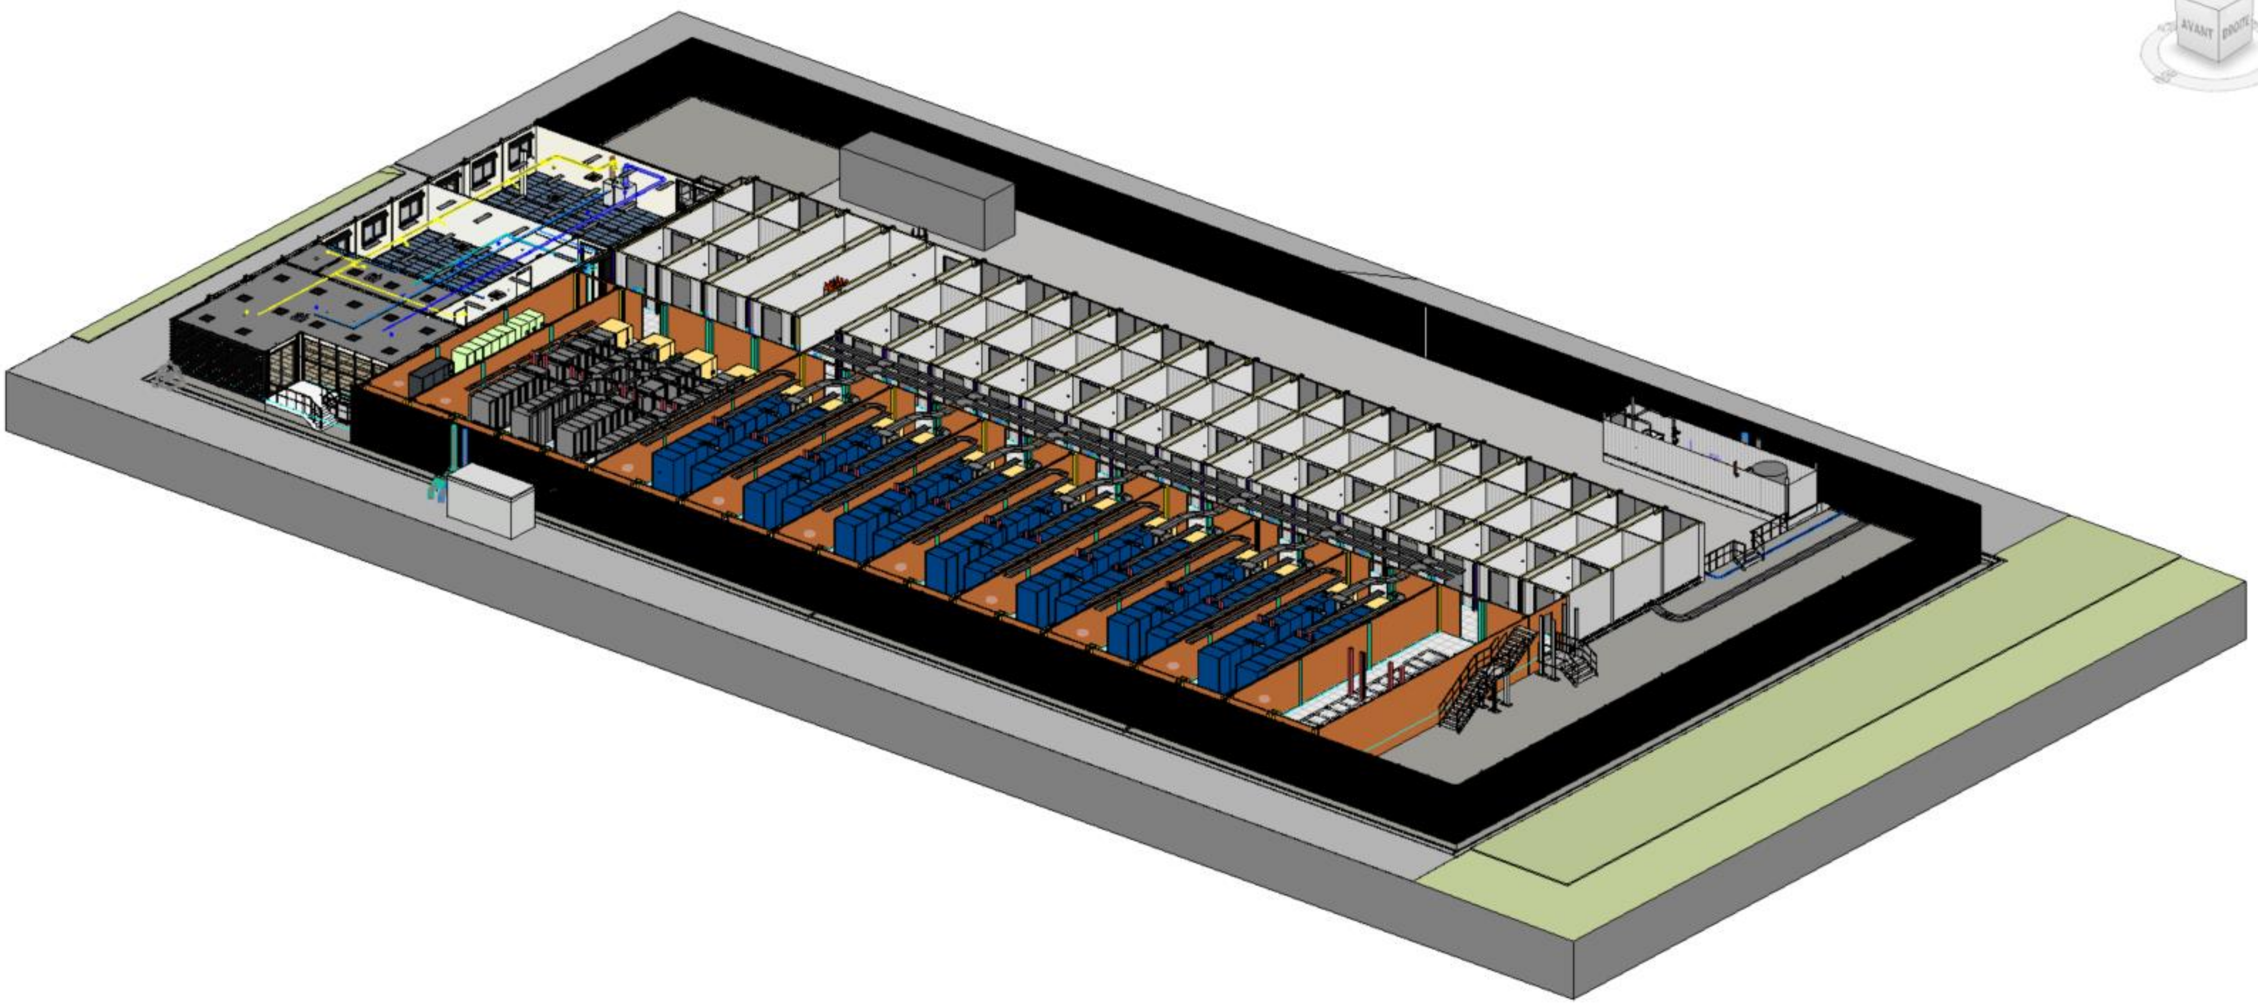

Member of the Helmholtz Association

JÜLICH
Forschungszentrum

JEDI

JUPITER EXASCALE DEVELOPMENT INSTRUMENT

JEDI

#1 in Green500 (05/2024)
#189 in TOP500

JUPITER EXASCALE DEVELOPMENT INSTRUMENT

EuroHPC / Forschungszentrum Jülich

- Eviden BullSequana XH3000
  - 48× compute nodes (24× blades)
    - NVIDIA quad-GH200 96 GB Grace-Hopper Superchip
    - Memory: 480 GB on CPUs + 384 GB on GPUs
    - NVIDIA quad-rail InfiniBand NDR200
- Usage
  - System management preparations
  - JUREAP
  - Application porting
- Also: GH200 COTS test nodes

POWER MEASUREMENT

Choice of Source

- XH3000 rack sensors vs Power meter
  - Similar measurements
  - In all tests, variation between 0.5 and 1.3%
  - 73.17 GF/W vs 72.73 GF/W
- Choice to use data from power meter:
  - More conservative measurement results

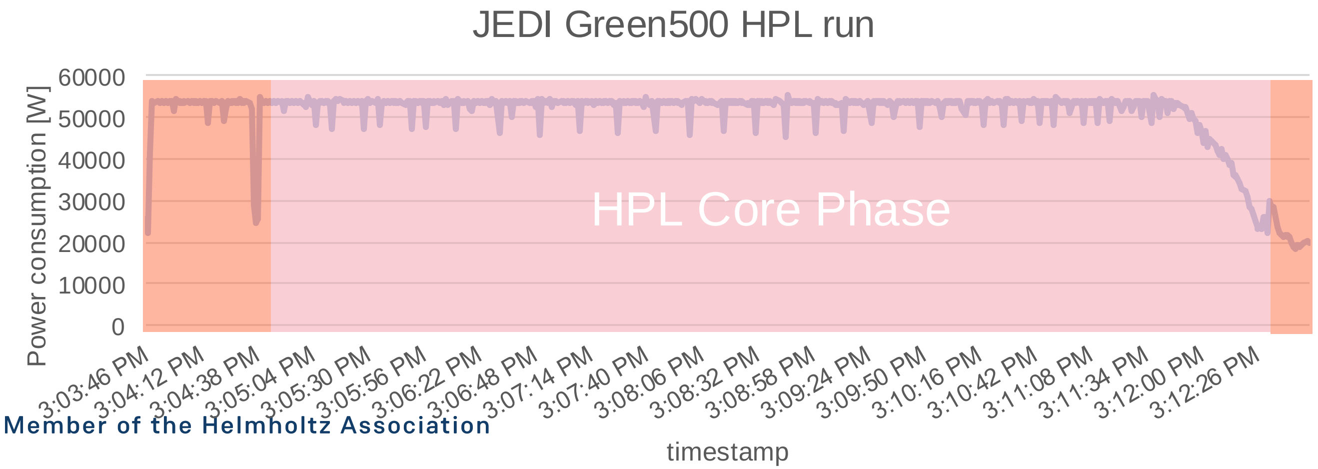

Sensor	Measurement	Difference
Rack sensors	51838.4 W	0.0%
Power meter	52154.3 W	+0.6%

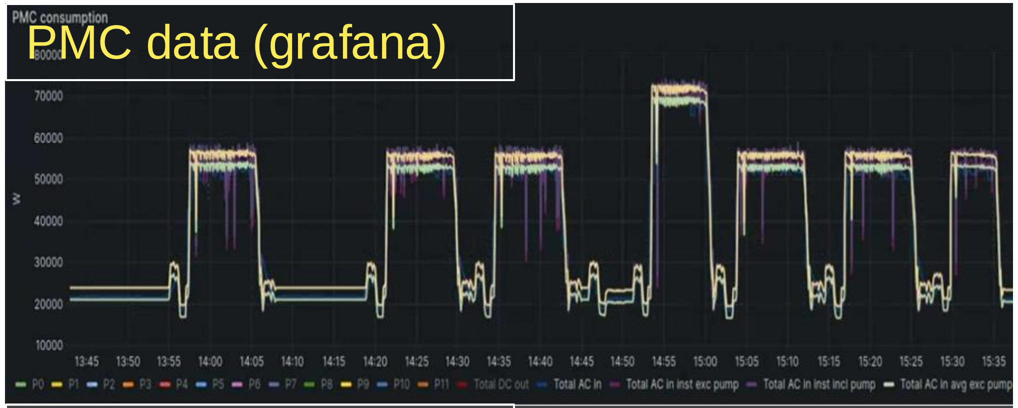

JÜLICH Forschungszentrum

GREEN500 – JUNE 2024

Green500 Data

Rank	TOP500 Rank	System	Cores	Rmax (PFlop/s)	Power (kW)	Energy Efficiency (GFlops/watts)
1	189	**JEDI** - BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200, ParTec/EVIDEN EuroHPC/FZJ Germany	19,584	4.50	67	72.733

APPLICATIONS

ASSESSING WITH APPLICATIONS

- Theoretical FLOP/s and GB/s are nice; but building machines for <u>users</u>
- → **Applications** core of procurement assessment
- Define representative benchmarks, *ExaBench*
  1. Analyze workload (JSC, DE, EU, 🌍)
  2. Select fitting applications
  3. Benchmark*ize* them
  4. Submit as part of specification

  → Get best machine

EVALUATION

- Criteria
  - Requirements to project planning, etc.
  - Technical requirements to overarching design and details
  - Performance of applications, benchmarks
    - Total cost of ownership (TCO): How much science for money
    - Further categories (Synthetic Benchmarks, High-Scaling Applications)
- Quantified evaluation

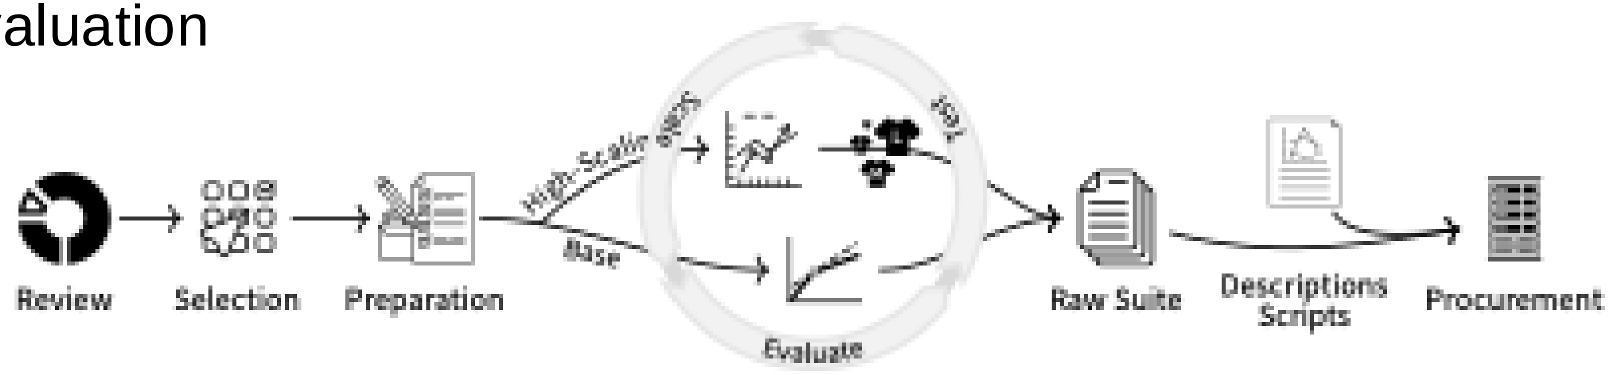

JÜLICH Forschungszentrum

JSC USERS

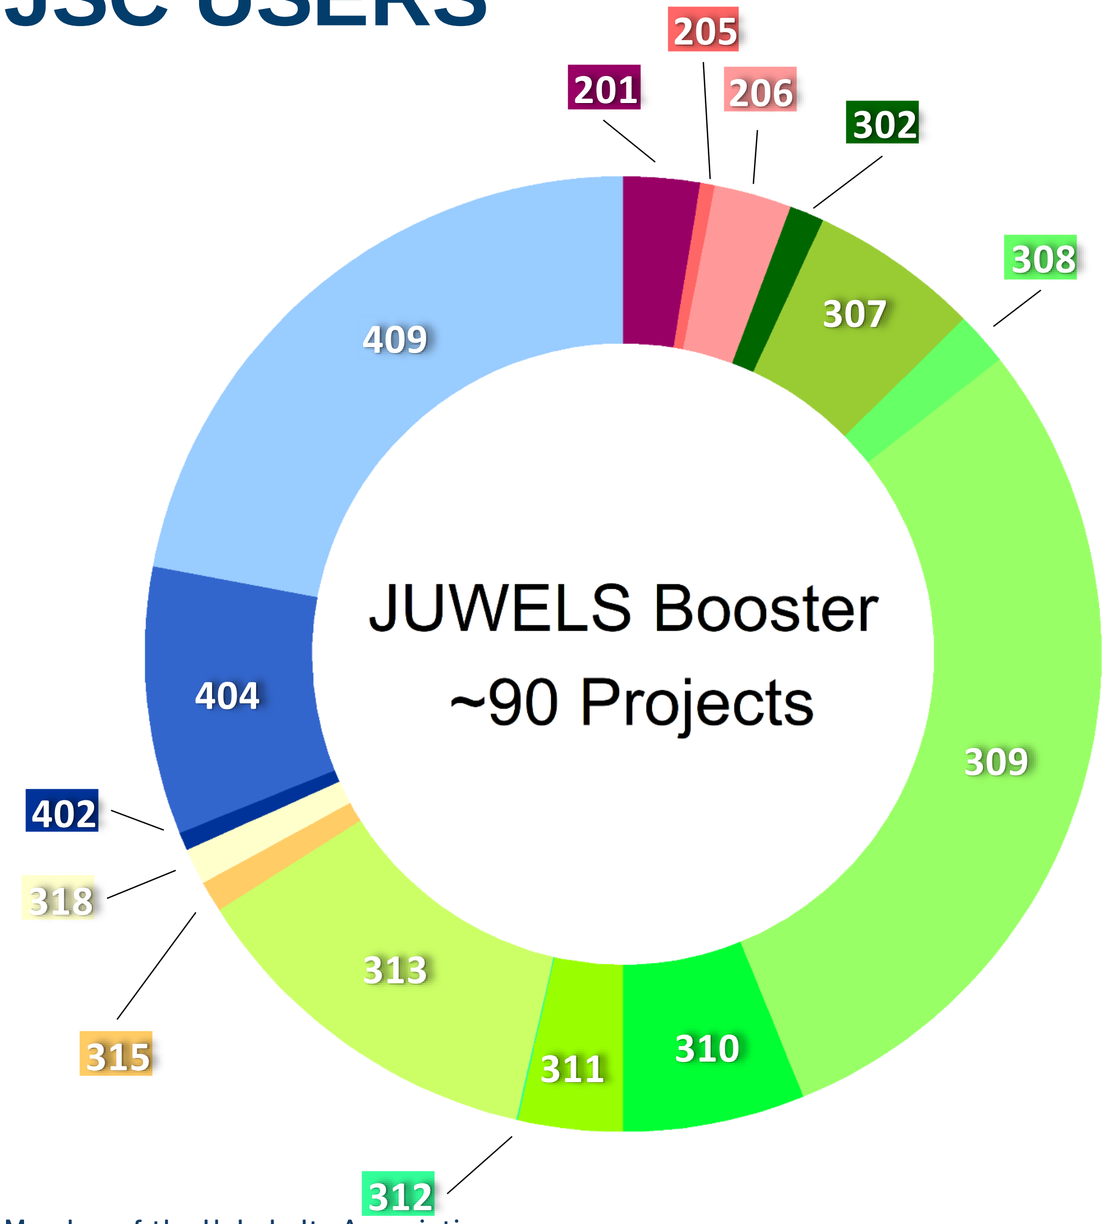

Reseach Fields

- **201** Basic Biological and Medical Research
- **205** Medicine
- **206** Neurosciences
- **302** Chemical Solid State and Surface Research
- **307** Condensed Matter Physics
- **308** Optics, Quantum Optics and Physics of Atoms, Molecules and Plasmas
- **309** Particles, Nuclei and Fields
- **310** Statistical Physics, Soft Matter, Biological Physics, Nonlinear Dynamics
- **311** Astrophysics and Astronomy
- **312** Mathematics
- **313** Atmospheric Science, Oceanography and Climate Research
- **315** Geophysics and Geodesy
- **318** Water Research
- **402** Mechanics and Constructive Mechanical Engineering
- **404** Heat Energy Technology, Thermal Machines, Fluid Mechanics
- **409** Computer Science

→ Define Benchmarks

SUB-BENCHMARKS, VARIANTS

- Type of benchmarks
  - Applications benchmarks
  - Synthetic benchmarks
- Execution targets
  - JUPITER Booster (GPU, CPU)
  - JUPITER Cluster (CPU)
  - MSA
- Application benchmark categories
  - TCO
  - High-Scaling

		Booster			Cluster	MSA
Before Dialogue	After Dialogue	GPU	GPU High-Scale	CPU	CPU	
Amber	~~Amber~~	~~✓~~				
Arbor	Arbor	✓	✓			
Chroma	Chroma	✓	✓			
Gromacs	Gromacs (2)	✓				
ICON	ICON (2)	✓				
JUQCS	JUQCS	✓	✓			✓
nekRS	nekRS	✓	✓			
ParFlow	ParFlow	✓				
PIConGPU	PIConGPU	✓	✓			
Quantum ESPRESSO	Quantum ESPRESSO	✓				
SOMA	~~SOMA~~	~~✓~~				
AI-MMoCLIP	AI-MMoCLIP	✓				
AI-NLP	AI-NLP	✓				
AI-ResNet	~~AI-Resnet~~	~~✓~~				
dynQCD	dynQCD				✓	
NAStJA	NAStJA				✓	
Graph500	Graph500			✓		
HPCG	HPCG	✓			✓	
HPL	HPL	✓			✓	
IOR	IOR			✓	✓	
LinkTest	LinkTest			✓	✓	✓
Multi-Flow IP	~~Multi-Flow IP~~			~~✓~~		
OSU	OSU (2)	✓		✓	✓	
STREAM	STREAM	✓			✓	

TCO APPLICATIONS

Total Cost of Ownership

- Traditional benchmark category
- ***How much of benchmark suite can be run in lifetime of system?*** Also: energy
- Key: same metric for each benchmark
  - Unit: time / s
  - Needed to convert rate → time
- One reference run for formula (e.g. 8 nodes); additional strong-scaled runs (e.g. 4, 16)
- Weights per individual benchmark

HIGH-SCALING APPLICATIONS

- Novel category for us
- Give benchmarks a focus on large-scaleness of system
- Compare execution on full* JUWELS Booster to full* JUPITER Booster
  - *: Use 50 PFLOP/s$^{\text{th. peak}}$ part of JUWELS Booster

  → compare to 1000 PFLOP/s$^{\text{th. peak}}$ part of JUPITER Booster
- Challenging design, challenging commitments
  - Design for unknown system, unknown device, unknown memory size
    Introduce 3 memory variants: small ($^2/_4$), medium ($^3/_4$), high ($^4/_4$ JWB A100 memory)
  - Many tests on scale at JUWELS Booster
  - Some internal extrapolations to JUPITER scale

- Arbor
  tiny ($^1/_4$), small, medium, large
- Chroma
  small, medium, large
- JUQCS
  small, large
- nekRS
  small, medium, large
- PIConGPU
  small, medium, large

JUREAP

Seeding Exascale in Europe!

jureap@fz-juelich.de • https://www.fz-juelich.de/en/ias/jsc/jupiter/jureap

JUPITER Research and Early Access Program

Mitglied der Helmholtz-Gemeinschaft

Foto von Robert Wiedemann auf Unsplash

JÜLICH Forschungszentrum

JUREAP OVERVIEW

One necessary goal – and three sufficient goals for success of JUREAP with ~20 applications

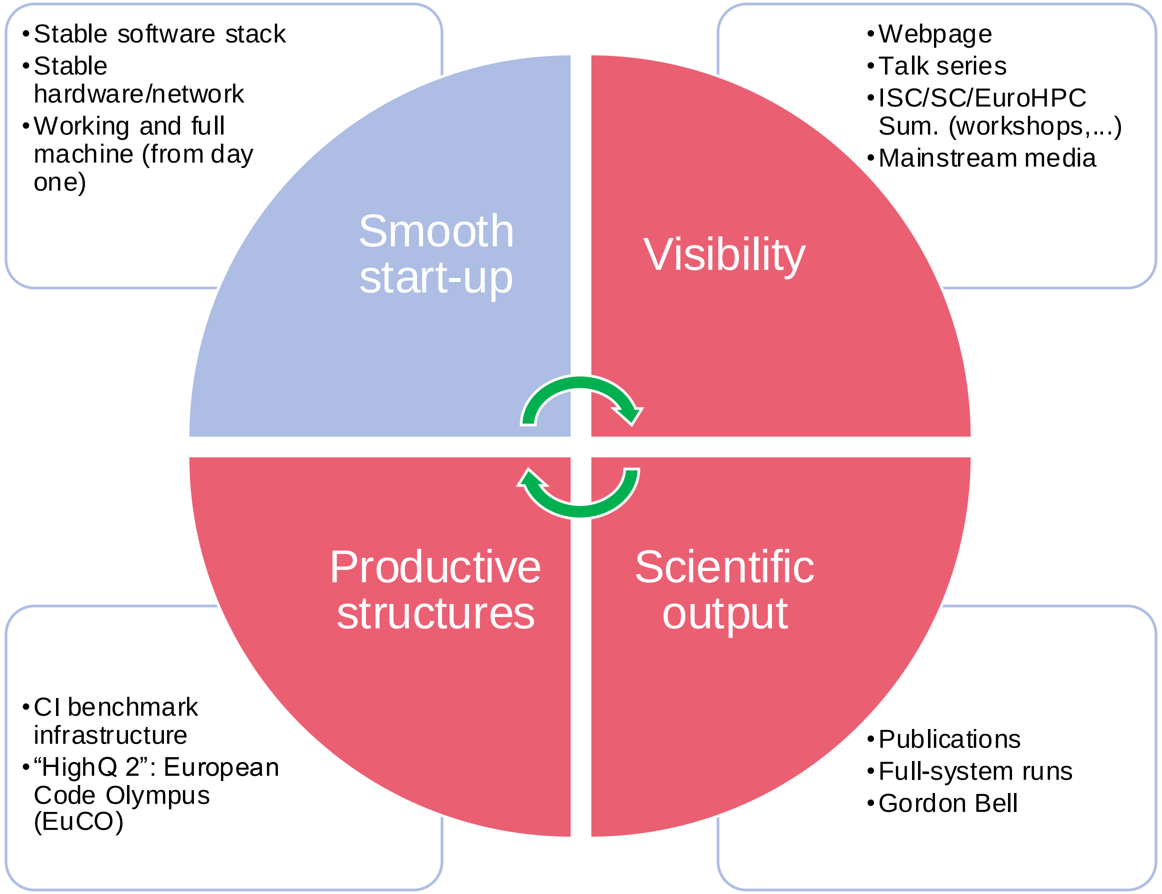

~20 participating teams with codes/applications

- EXA: Exascale Applications
- LAX: Large Applications with X
- Balanced in terms of simulation, AI, and scientific domains
- Computing time via GCS and EuroHPC

JUREAP

JUPITER Research and Early Access Program

IMPORTANT: Timeline not final yet!

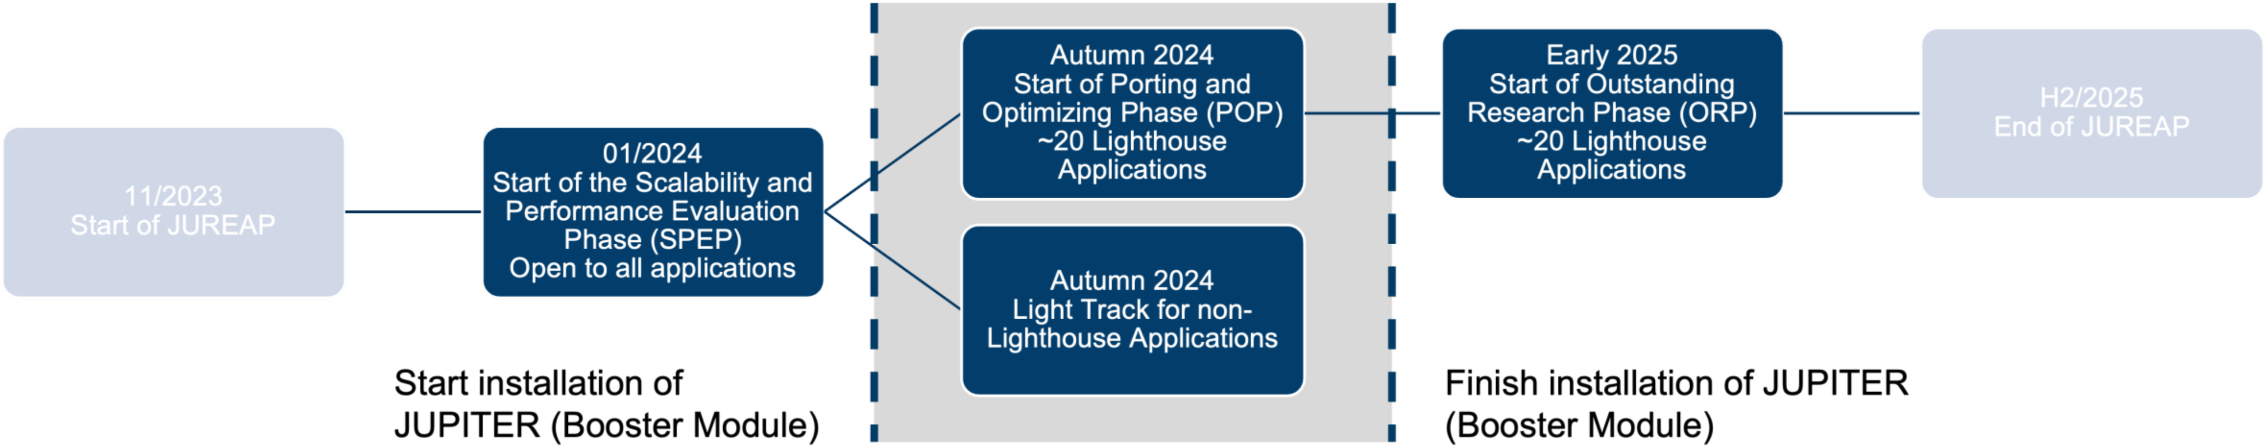

Phase 1: Scalability and Performance Evaluation Phase (SPEP)
Phase 2: Porting and Optimizing Phase (POP)
Phase 3: Outstanding Research Phase (ORP)

JUREAP call open since Jan‘2024
https://events.hifis.net/event/1239/

GH RESULTS

GH200 TEST NODES

- GH200 Prototype
- 2× Grace-Hopper superchips
  - 1 Grace CPU (72 cores), 480 GB LPDDR5X RAM
  - 1 H100 GPU
  - TDP 700-1000 W
- Slightly different variant compared to JUPITER node design

GPU STREAM

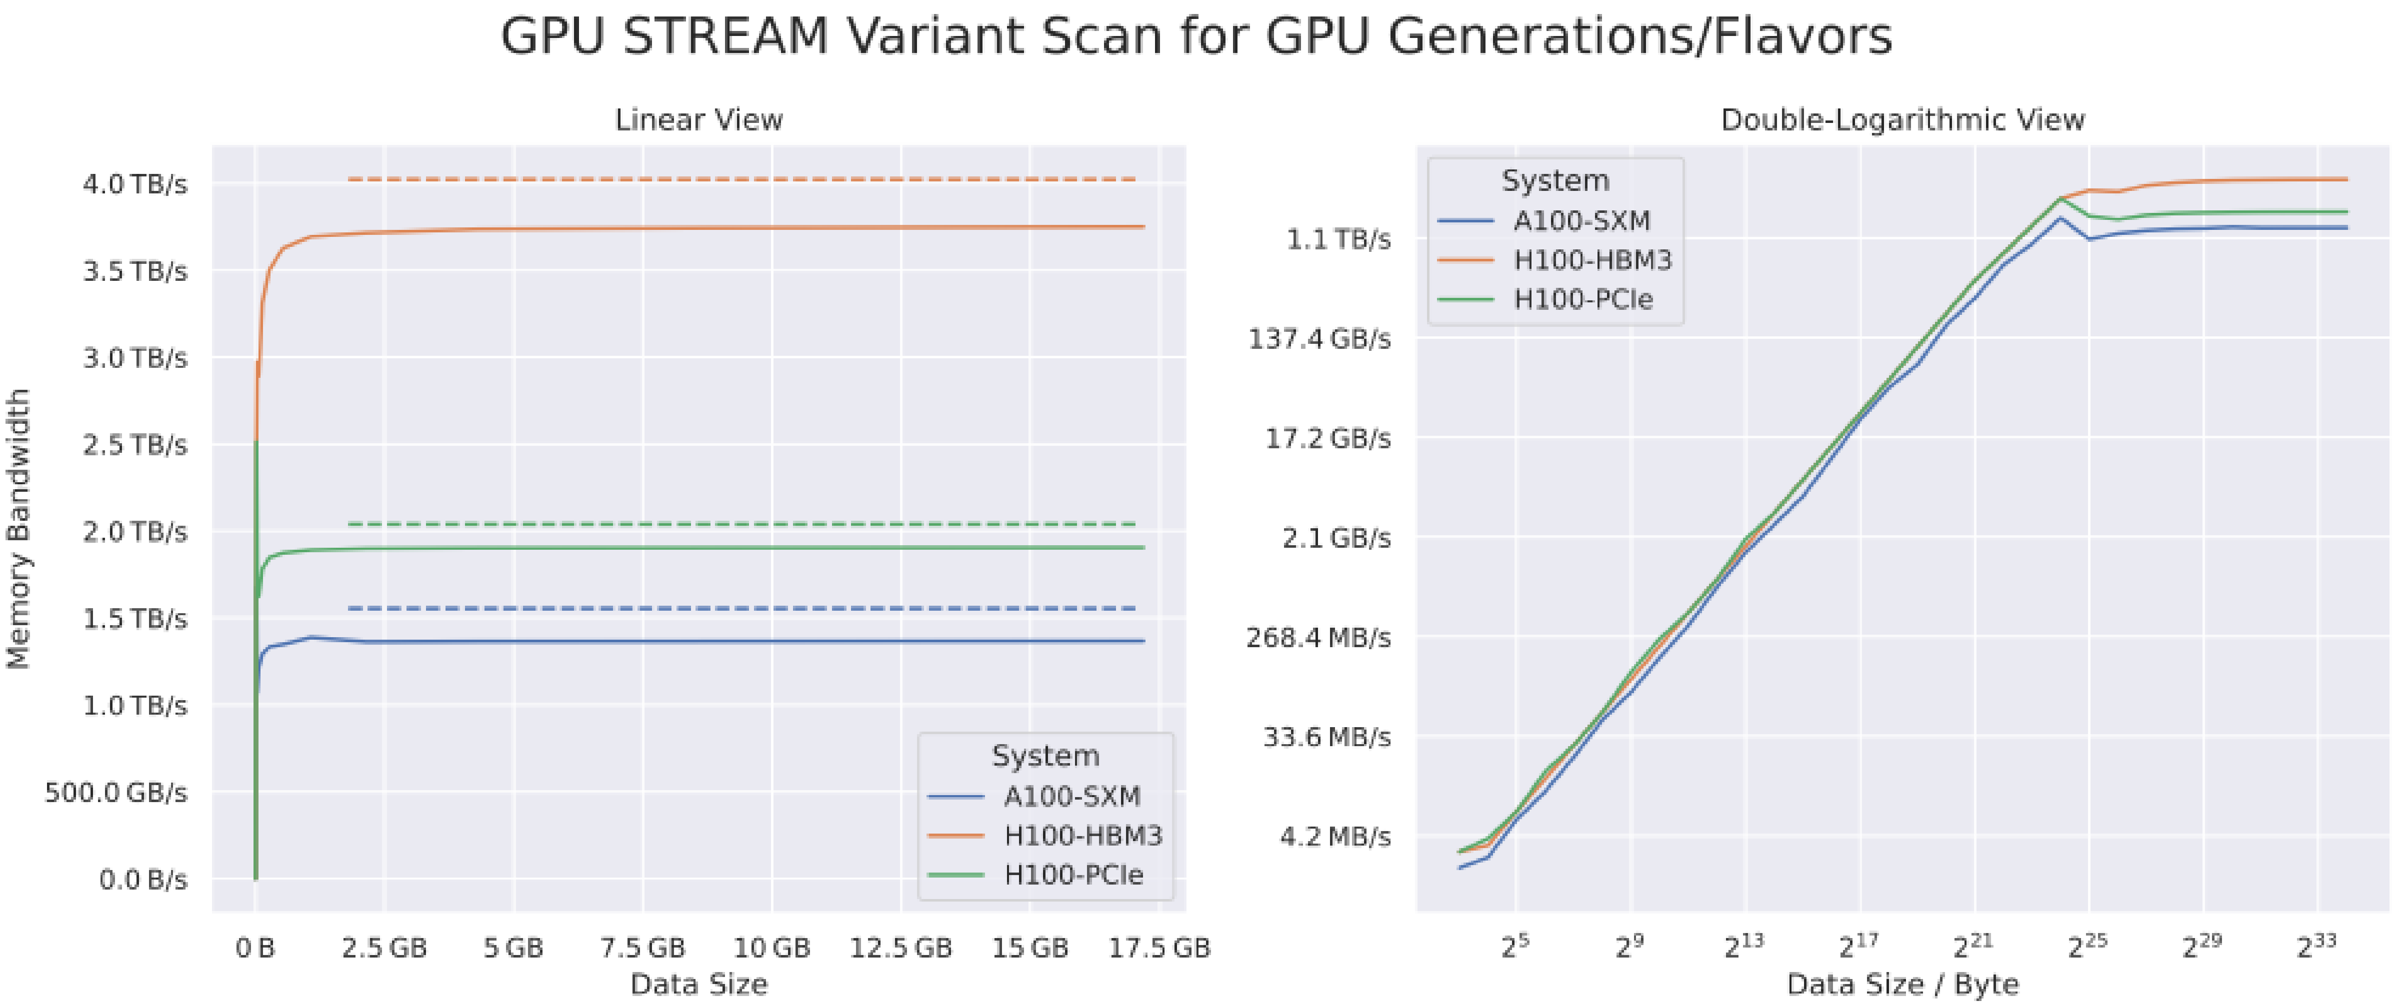

https://github.com/AndiH/CUDA-Cpp-STREAM

KOKKOS LQCD BENCHMARK

By Simon Schlepphorst / JSC

- Simplified Staggered Fermion Dirac Operator
- MI250
- A100
- H100 (PCIe)
- H100 (GH200)

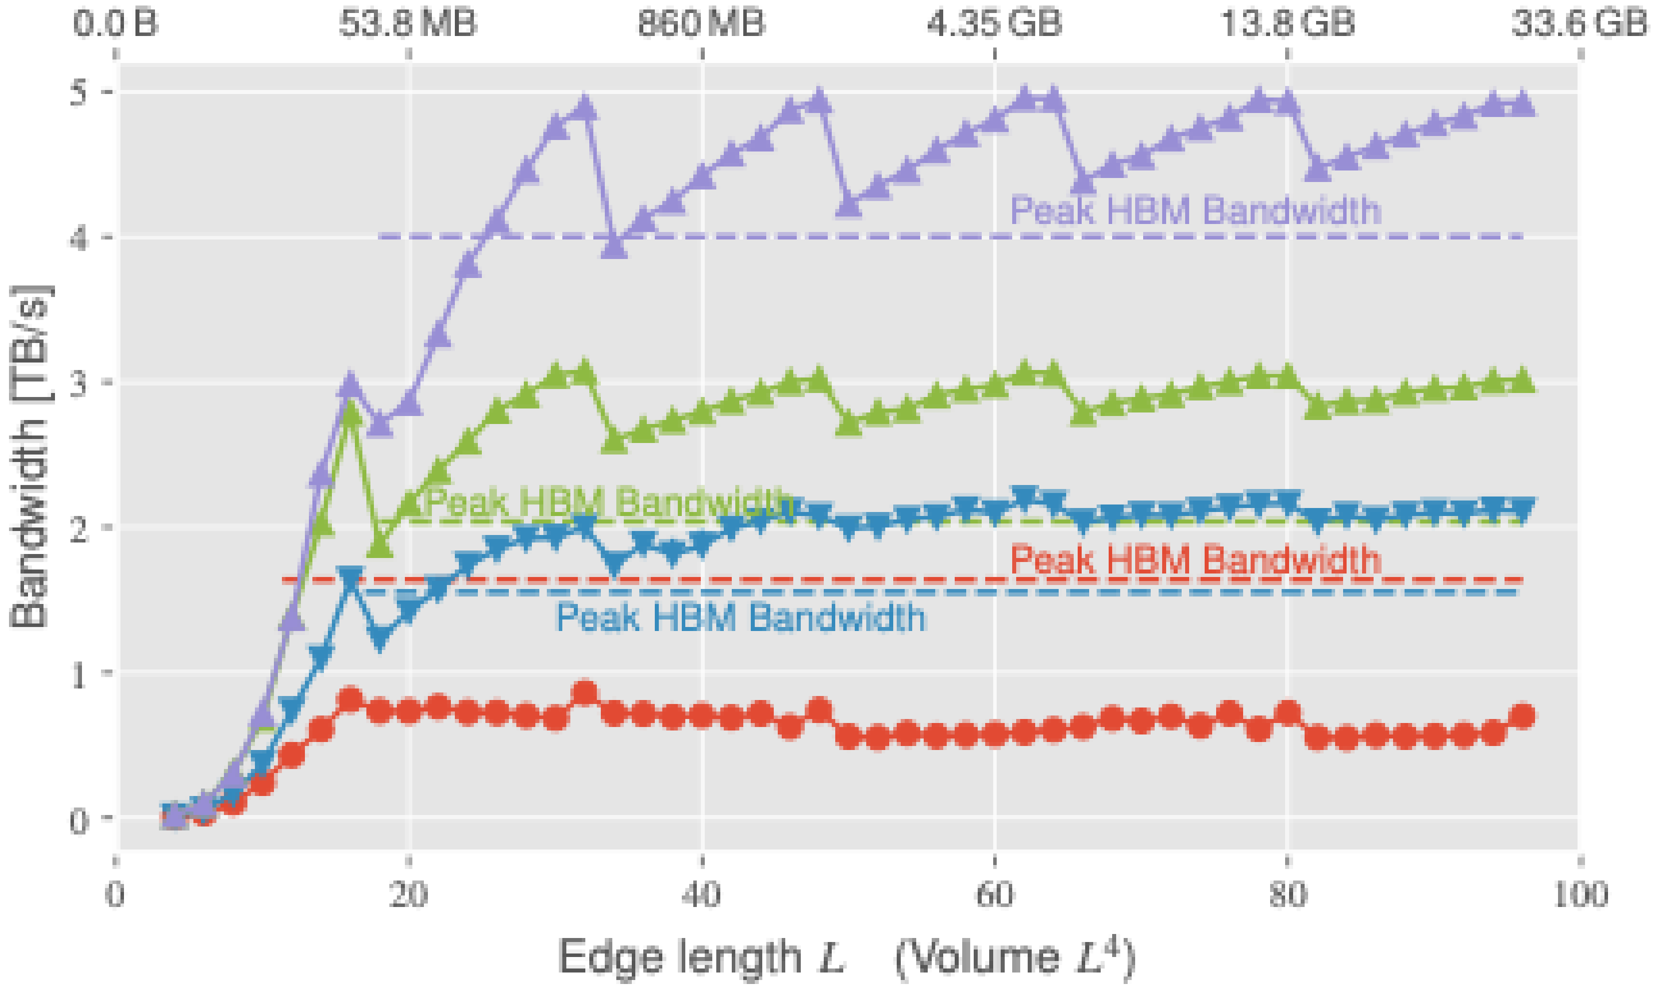

https://gitlab.jsc.fz-juelich.de/sdl-nqft/kokkos-lqcd-benchmarks

MEGATRON-LM

By Chelsea John / JSC

- OpenGPT-X: BMWK project for LLM with EU languages, Open
- 1 node benchmark:
  800M GPT Model
- Increasing Batchsizes
- Excellent performance on H100, especially GH200 variant (HBM3, TDP)

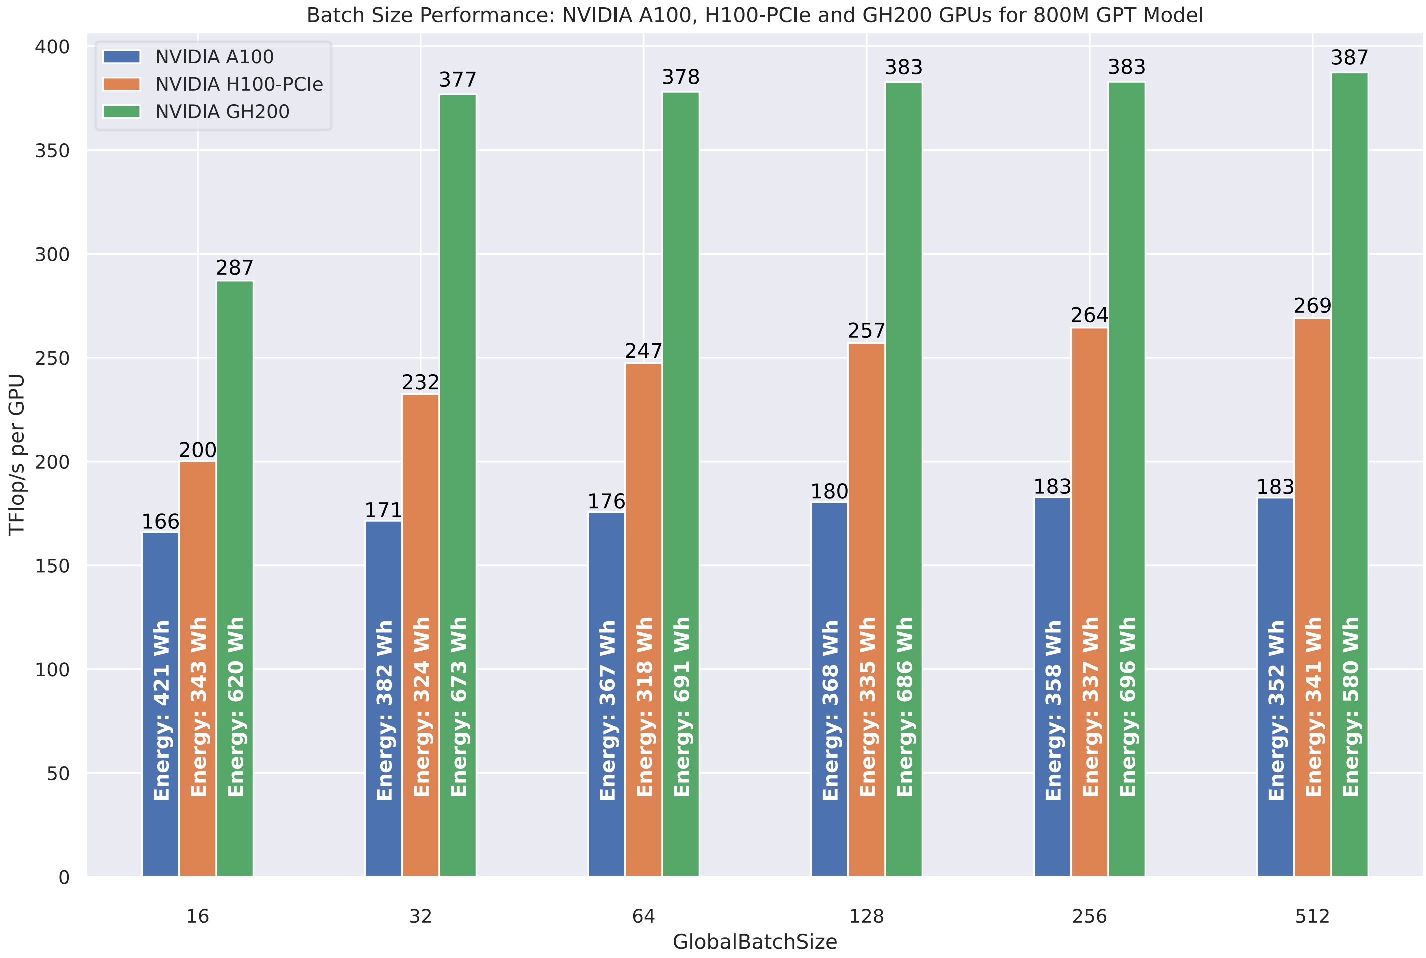

https://github.com/OpenGPTX/Megatron-LM/

MPTRAC

By Lars Hoffmann / JSC

- Lagrangian particle dispersion model: atmospheric transport processes (troposphere/stratosphere)
  → volcanic emissions
- Continuously optimized for GPUs
  Recently: Significant speedup on A100
- First test on GH200

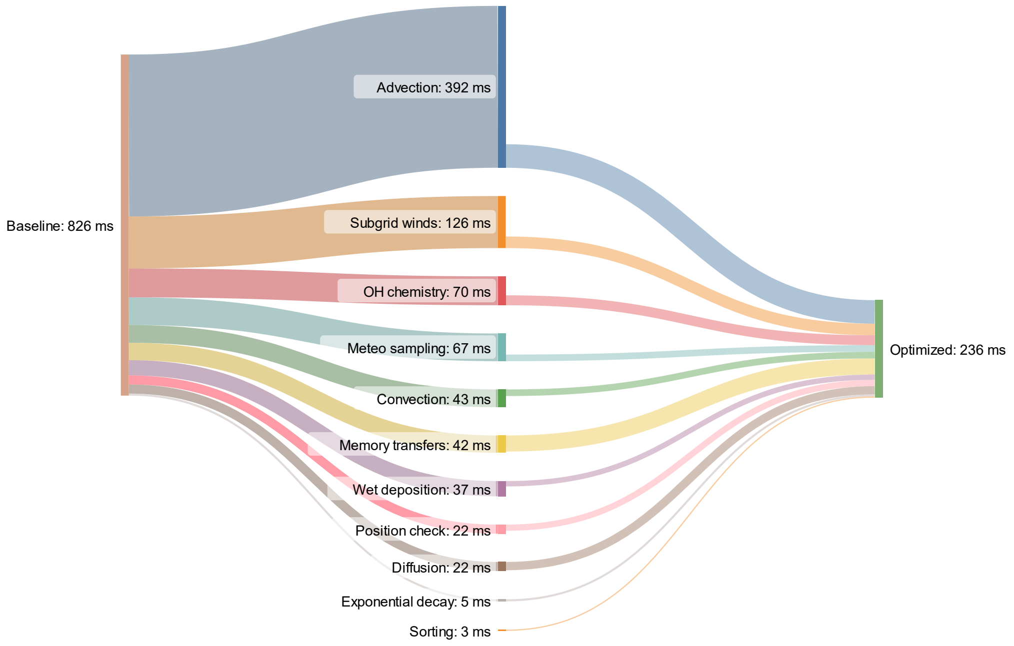

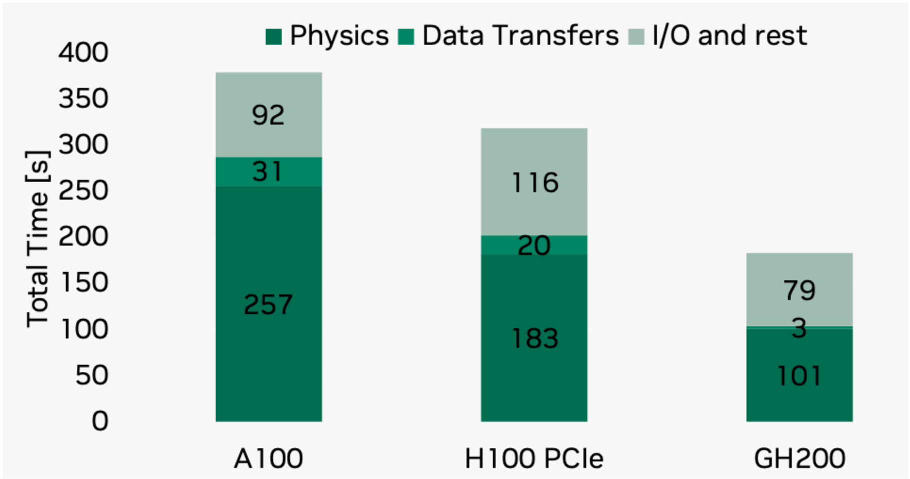

See also GTC talk by Mathias Wagner

OTHER GH200 RESULTS

JUWELS Booster A100 vs. GH200

- ICON:
  - R2B4 Benchmark: 563 s vs. 343 s → 1.62×
  - By Manoel Römmer / JSC
- nekRS
  - Rayleigh–Bénard Convection: 2.16×
  - By Mathis Bode / JSC
- Arbor
  - Busyring Benchmark: 330 s vs 167 s → 1.97×
  - By Thorsten Hater / JSC

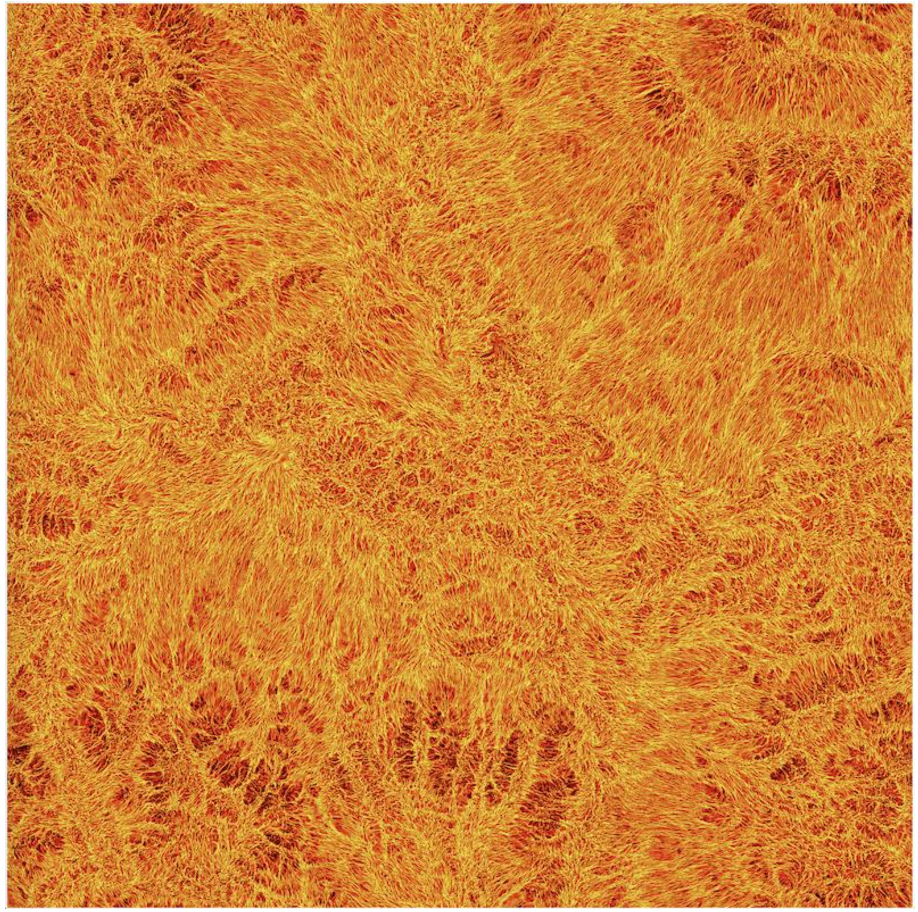

RBC Visualization (M. Bode)

MAELSTROM TESTS

- MEALSTROM: EuroHPC-JU-funded project, ended in March 2024
- Enablement/optimization of 6 weather & climate applications using ML
- D3.7 (Feb 2024): Benchmarking of final applications on diverse hardware

→ www.maelstrom-eurohpc.eu

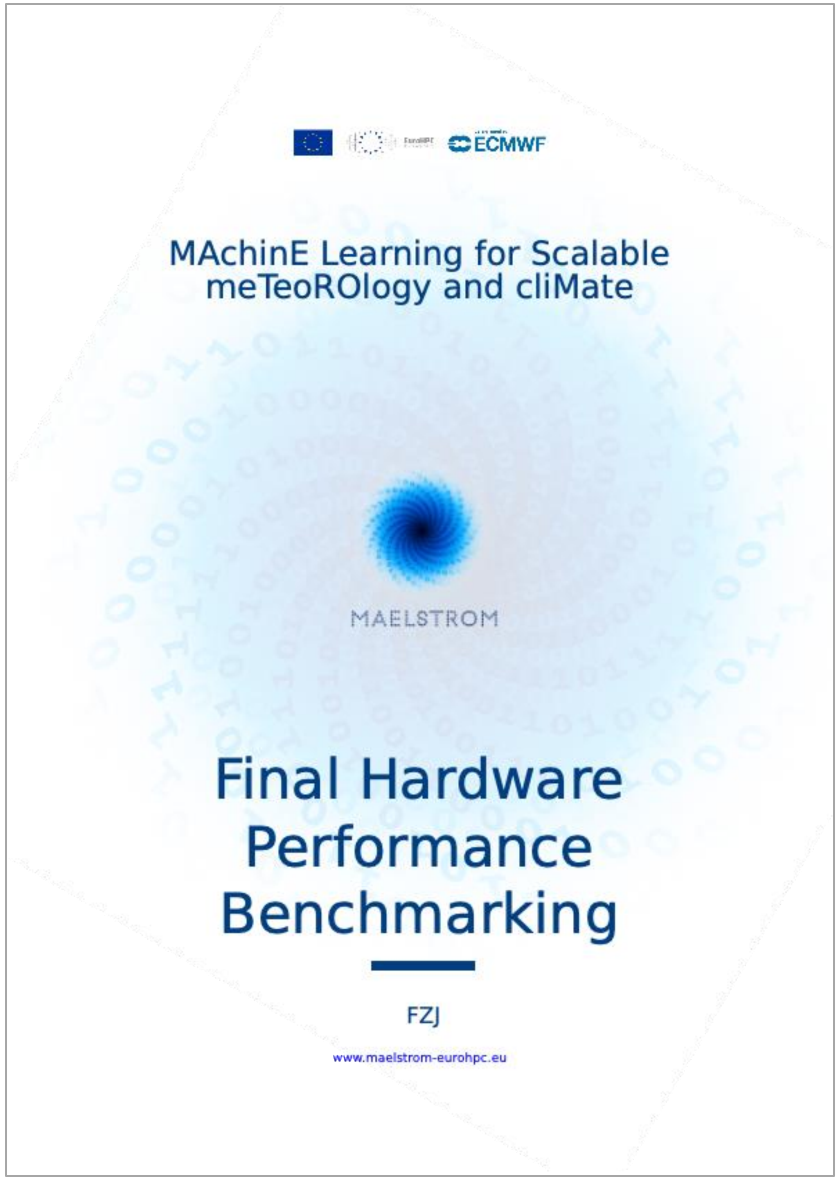

MAELSTROM TESTS

Energy-to-Solution Application Overview

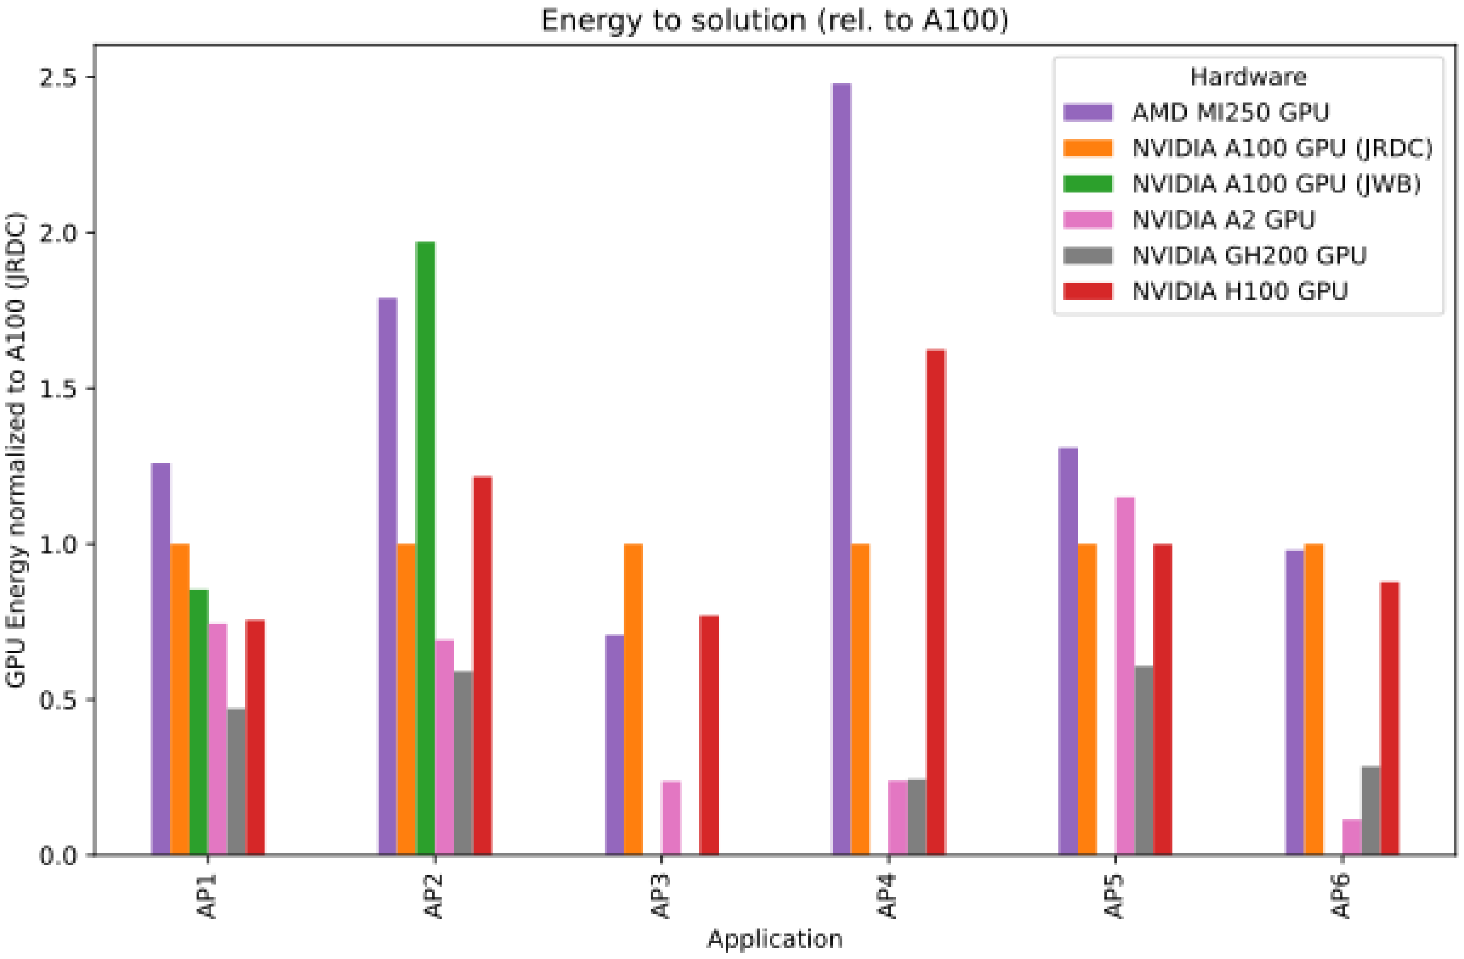

JÜLICH
Forschungszentrum

MAELSTROM TESTS

AP1: MetNor

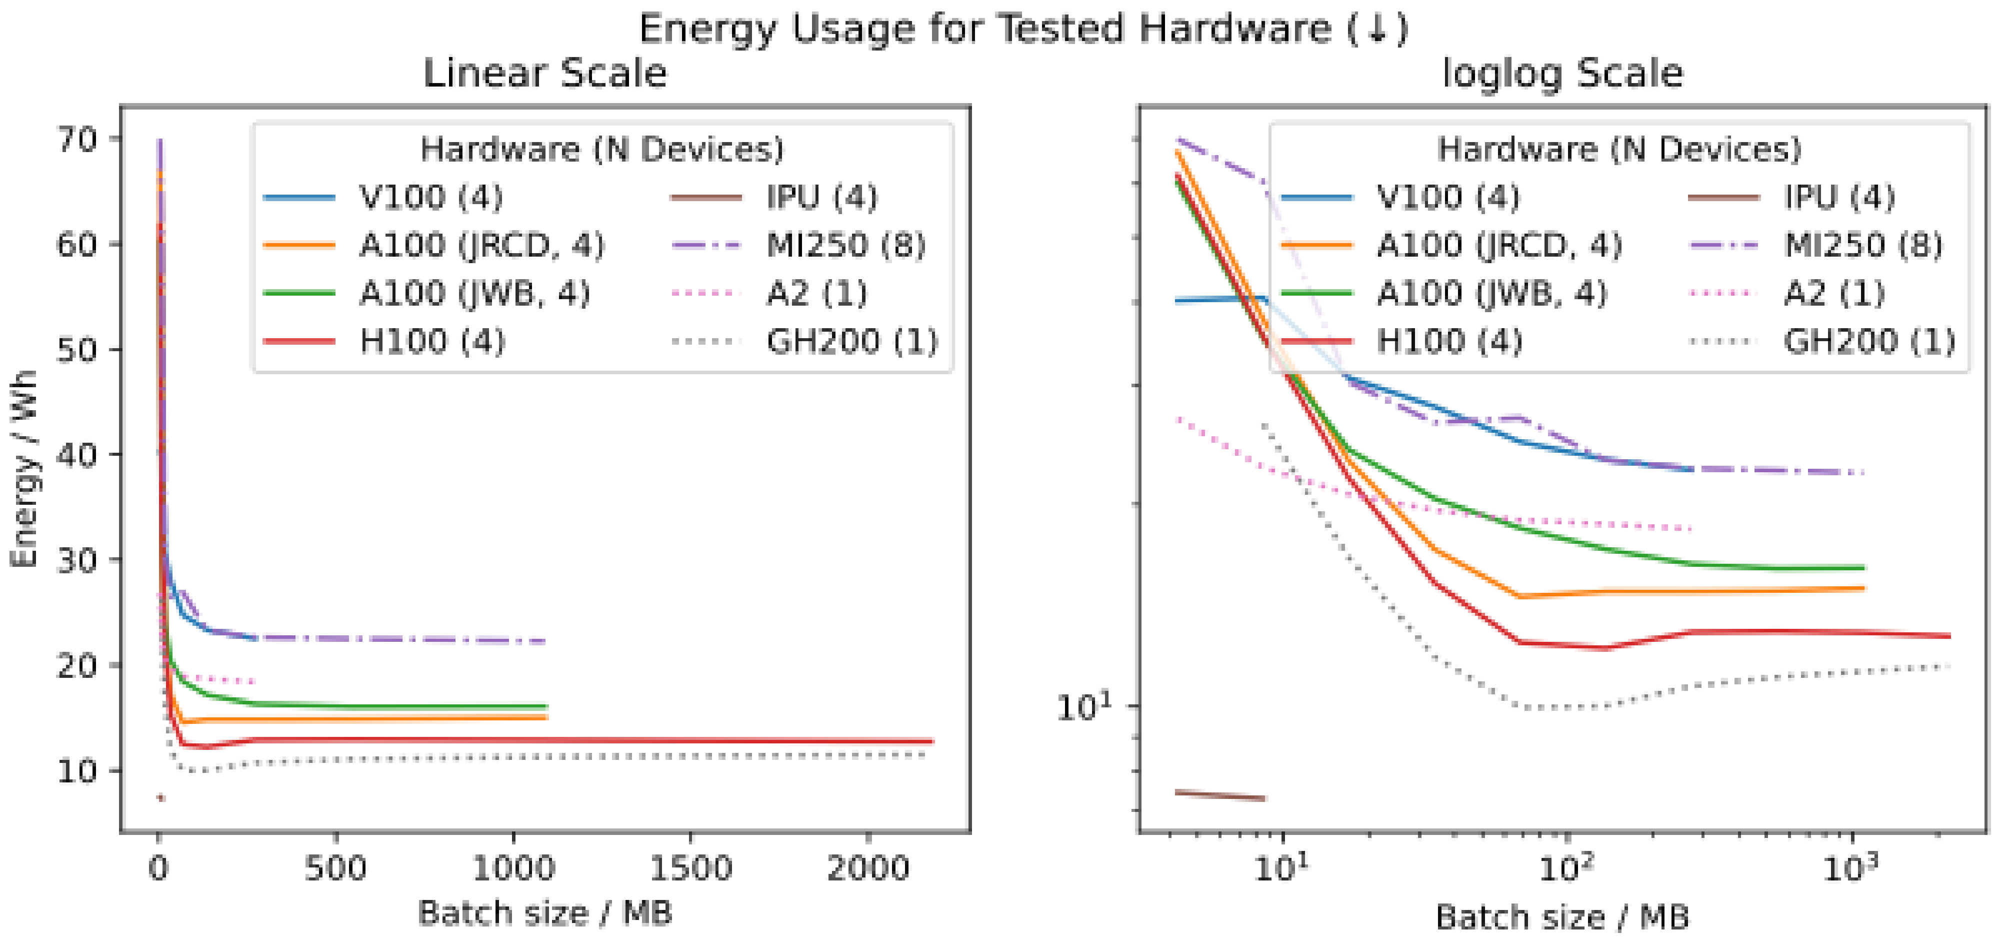

FIRST CPU INVESTIGATIONS *(GRACE)*

- Focus mostly on GPU currently
- Some first results on Grace hardware

→ Very competitive performance, especially wrt TDP (but still early)

DynQCD: 1.5× vs. EPYC Rome 7742 (2×64 cores)

- Best: Grace-Clang, ACfL
- Slightly worse: GCC
- Investigating FMLA instructions
- *(Auto-Vectorization works well!)*

NAStJA:

- 2.3× vs. EPYC Rome 7402 (2×24 cores)
- 5.6× vs Intel Skylake 8168 (2×24 cores)

JUQCS: 1.35× vs. EPYC Rome 7402 for 31 Qubits (2×24 cores)

FLEUR:

- 1.2× vs. Intel Skylake 8168 (2×24 cores, 400 W TDP)
- 0.8× vs. EPYC Rome 7742 (2×64 cores, 450 W TDP)
- 1.5× vs. Intel SPR-HBM (2×32 cores, 700 W TDP)

MAX DRIVING THE EXASCALE TRANSITION

MISSION BRIEFING OVERVIEW

- En route to JUPITER: EuroHPC JU system hosted at JSC
- Launched with focus on applications
- ~6000 nodes,
  24 000 H100 GPUs, 1 728 000 Arm cores, 24 000 NDR200 endpoints
- Landing in Modular Data Center
- Preparing for descent:
  - JUREAP
  - GH200 test systems

→ jupiter.fz-juelich.de

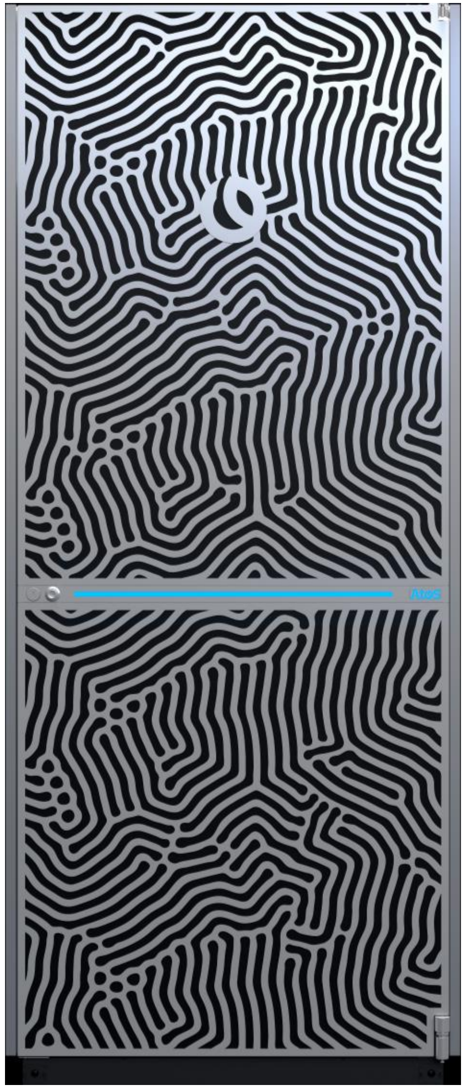

JUPITER

The Arrival of
Exascale in Europe

fz-juelich.de/jupiter | #exa_jupiter

EuroHPC
Joint Undertaking

Federal Ministry
of Education
and Research

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

GCS
Gauss Centre for Supercomputing