DEEP Projects

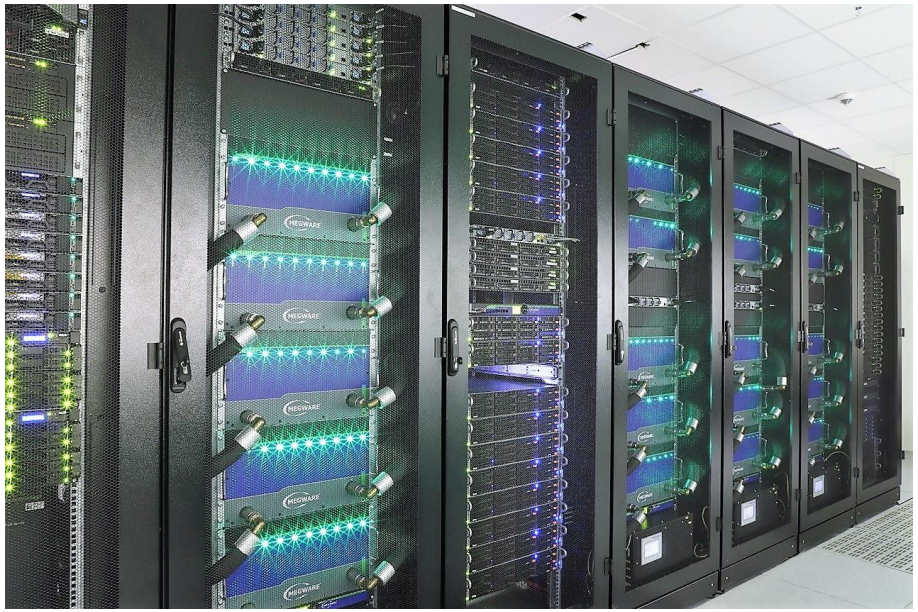

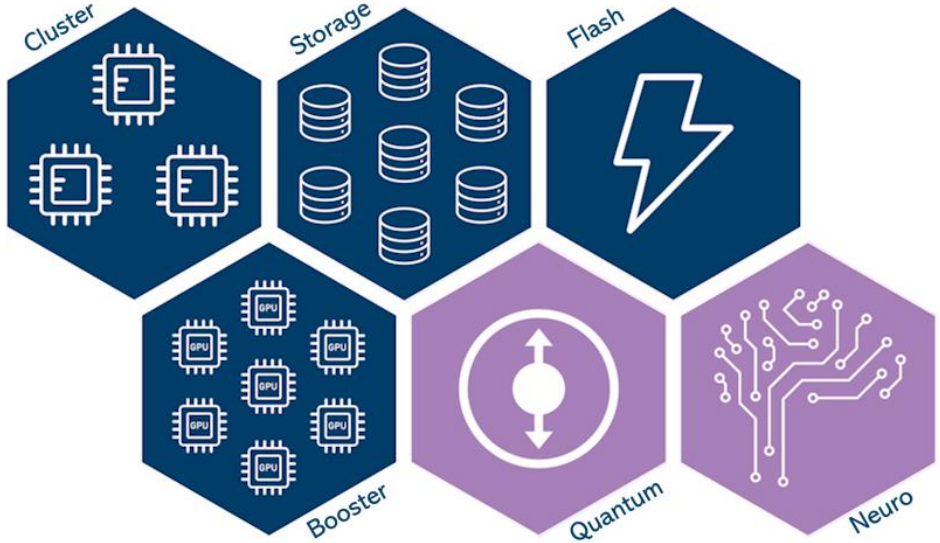

The DEEP-SEA Project

REINVENTING HPC
ISC High Performance
MAY 12 – 16, 2024 | HAMBURG, GERMANY

EVIDEN an atos business
BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación
cea
ECMWF
ETH zürich
FORTH
Fraunhofer ITWM
JÜLICH Forschungszentrum
KTH VETENSKAP OCH KONST
KATHOLIEKE UNIVERSITEIT LEUVEN
ParTec MODULAR SUPERCOMPUTING
lrz Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften
TECHNISCHE UNIVERSITÄT DARMSTADT
Technical University of Munich

Exascale Challenges

Application parallelism
- Applications must support billions of individual threads
- Lower-scaling applications / parts of applications should not run on a full Exascale system

DEEP-SEA

Truly scalable systems
- Huge numbers of devices need to exchange data with each other
- Collective communication operations are “slowing down” due to larger system sizes
- Network contention and reliability become worries

Energy efficiency
- Accelerators clearly beat CPUs for many (most?) codes
- System heterogeneity is a must
- Yet – portable accelerator programming is hard

DEEP-SEA

Memory and storage
- Ever growing gap between compute throughput and memory bandwidth
- New technologies like HBM suffer from capacity limitations & high energy consumption

DEEP-SEA

Workload diversity
- Exascale centers must run a wide variety of HPC, AI and data analytics workloads with highest energy efficiency
- One size does not fit all

DEEP-SEA

Modular Supercomputing Architecture

The MSA achieves composability of heterogeneous resources

- Cost-effective scaling
- Effective resource-sharing
- Match workload diversity

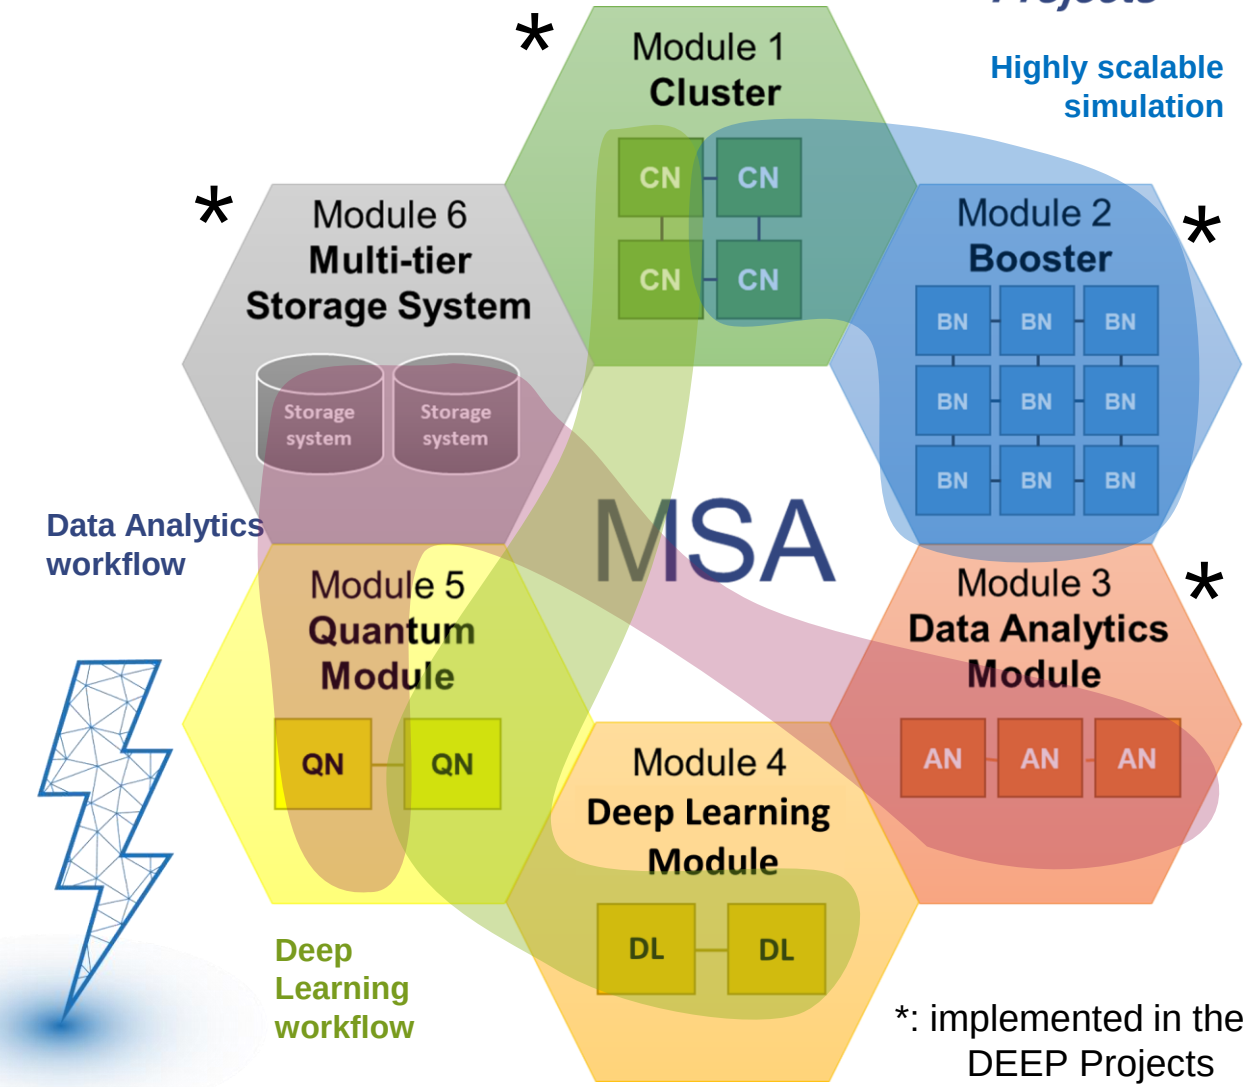

E. Suarez, N. Eicker, T. Moschny, S. Pickartz, C. Clauss, V. Plugaru, A. Herten, Kristel Michielsen, T. Lippert, *"Modular Supercomputing Architecture – A Success Story of European R&D",* ETP4HPC White Paper. (2022) Available at https://www.etp4hpc.eu/white-papers.html#msa.

E. Suarez, N. Eicker, Th. Lippert, "*Modular Supercomputing Architecture: from idea to production*", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, p 223-251, CRC Press. (2019)

Integrated Exascale-Ready SW Stack

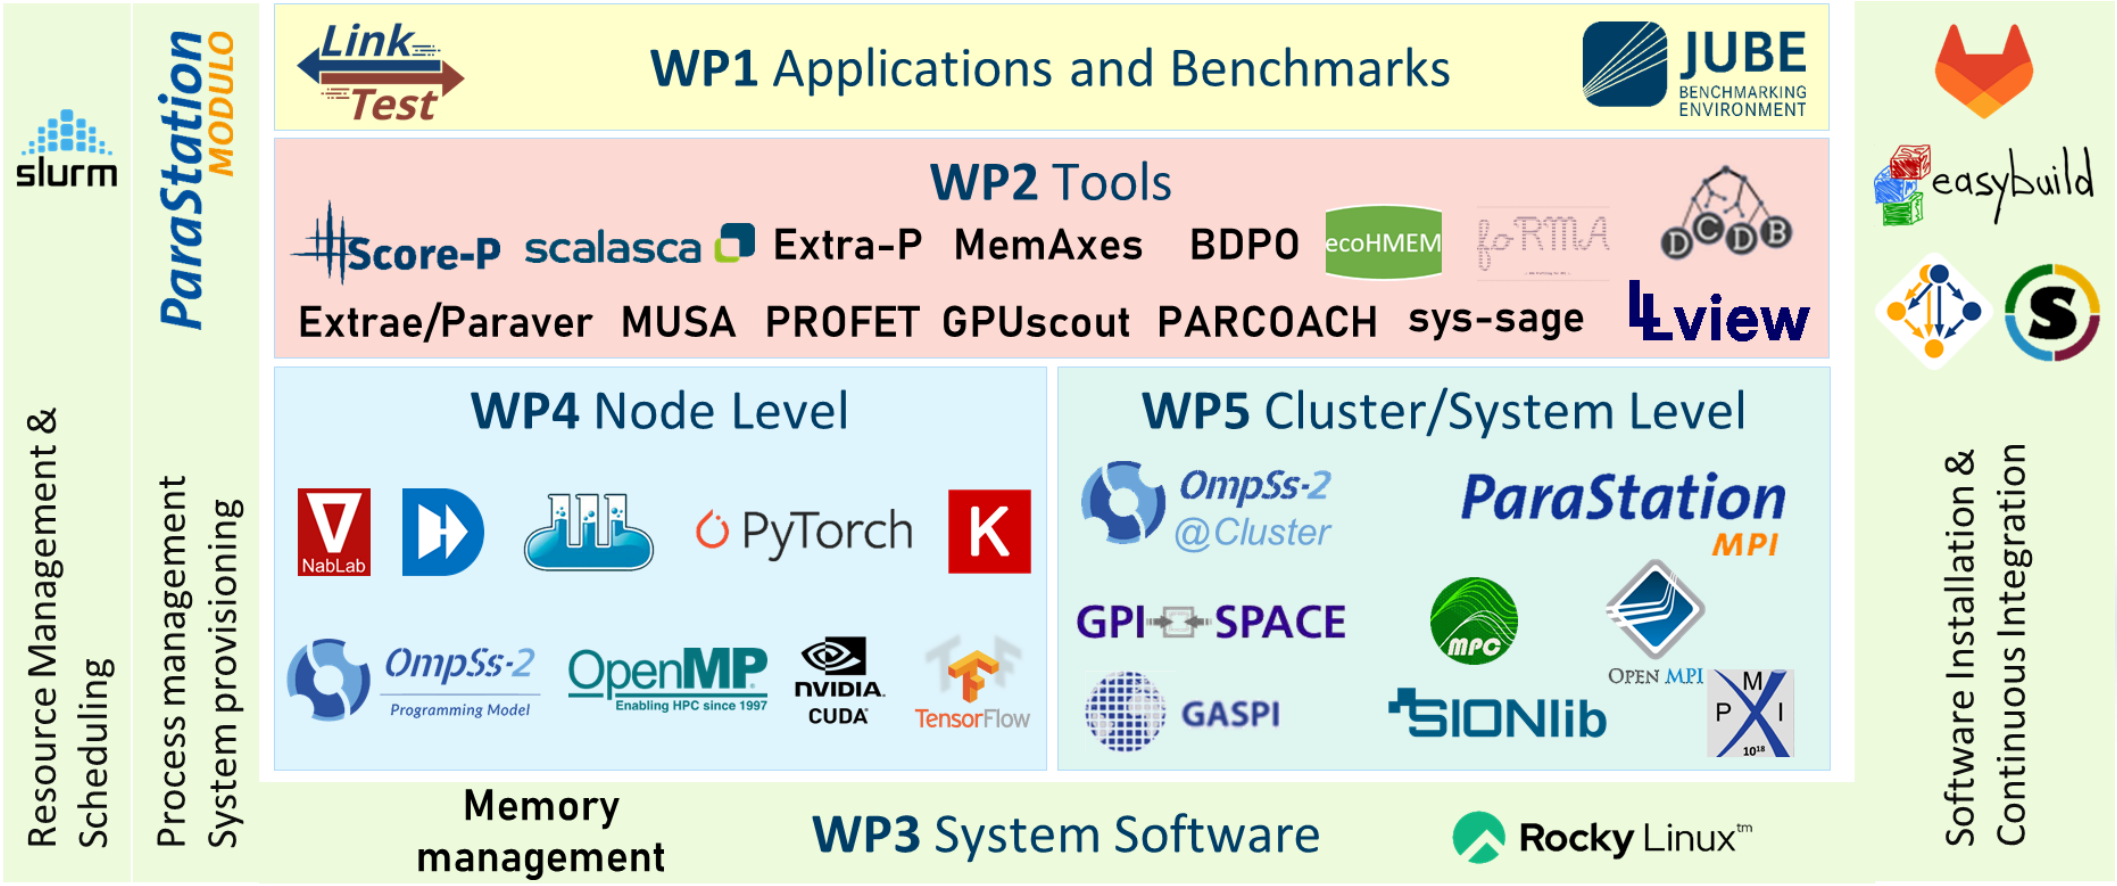

At the heart of the JUPITER system

Public release at https://gitlab.jsc.fz-juelich.de/deep-sea/easybuild-repository-public-release

Optimisation Cycles

Large variety of SW tools available to HPC SW developers for analysis and optimisation – in DEEP-SEA alone, these:

Score-P scalasca Extra-P MemAxes BDPO ecoHMEM

Extrae/Paraver MUSA PROFET GPUscout PARCOACH sys-sage Lview

Optimisation cycles encapsulate (complex) tool workflows for *specific purposes*

- Like assessing load balance or optimising energy use

They guide SW developers and make it easier to achieve specific goals

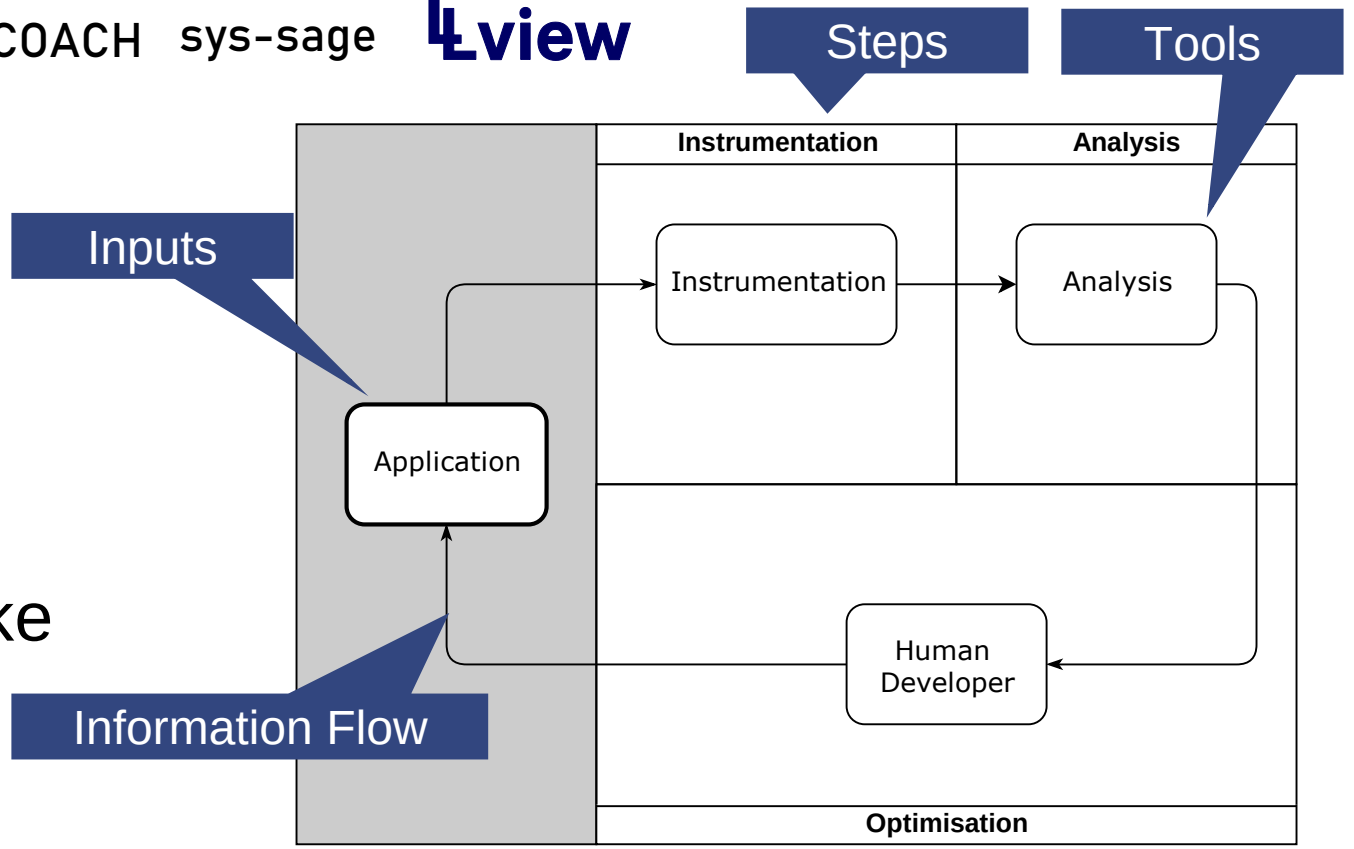

Based on material by A. Geiß, TU Darmstadt

Heterogeneous/Hierarchical Memory Tools

Combine fast (like HBM) with conventional memory (like DDR or persistent)

DEEP-SEA tools support

- Restructuring & tool-guided optimization of applications / automatic use by unmodified code

Modifying the application

OpenMP

Standard interfaces
OpenMP 5.0, SICM

Transparent to the application

Static
(compile time)

ecoHMEM

Object mapping by source-to-source compiler based on profile run

Dynamic
(run time)

SHAMBLES

Run time page-level migration by OS

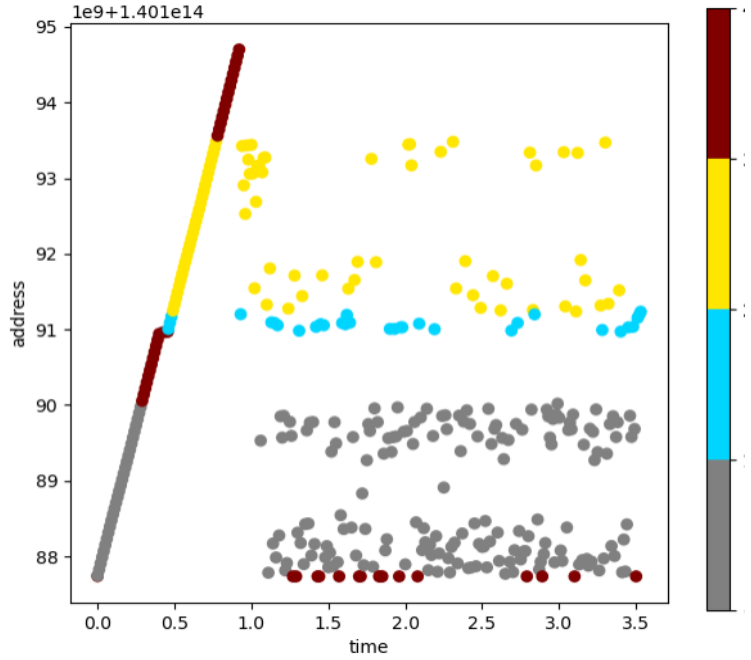

SHAMBLES scatter plot example for sparse kernel

Malleability

- Use of resources varies over time (yellow curve)
- Constant allocation of resources (blue)
- Improved dynamic allocation of resources (green)

Malleable applications can request or shed resources

MPI & Slurm prototype for enabling application-driven malleability

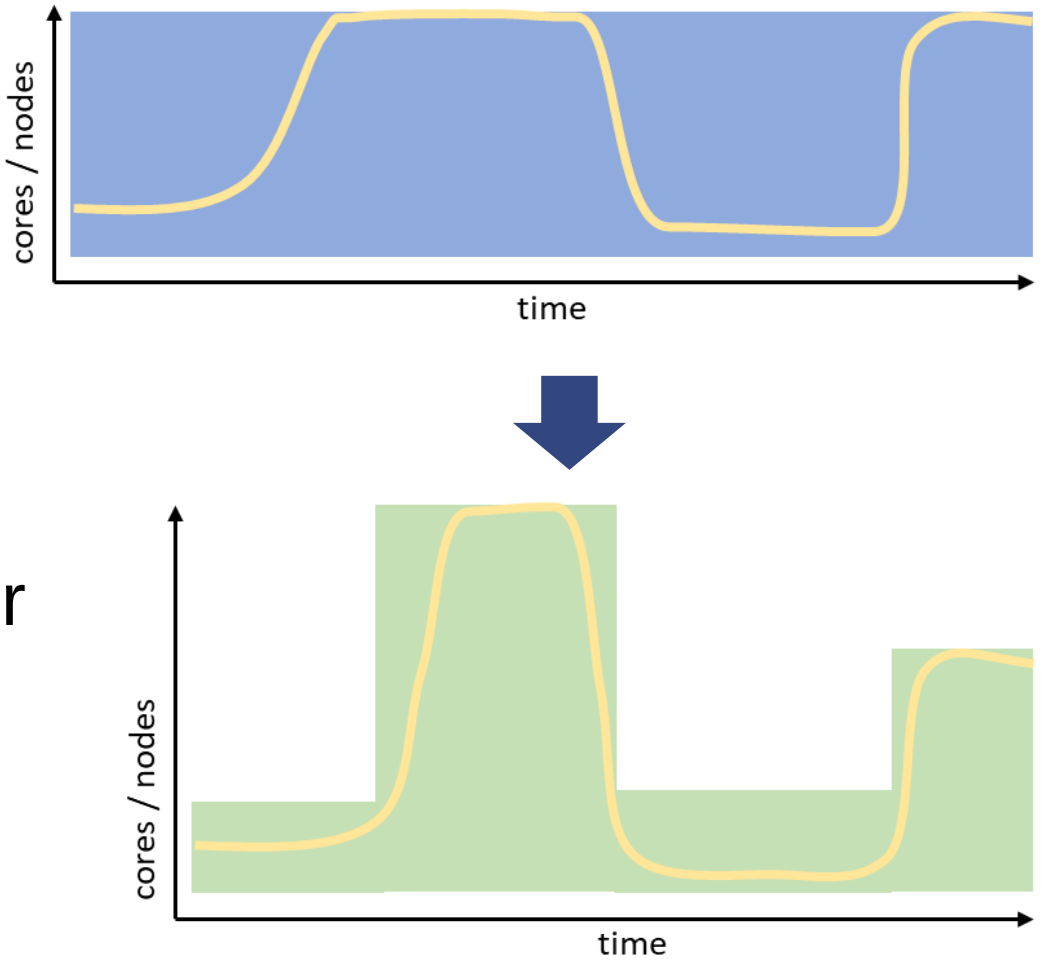

Ten Co-Design Applications

Weather and Climate
ML, Fortran,
OpenMP, MPI

Molecular Dynamics
C++, OpenMP, MPI,
CUDA

Neutron Monte-Carlo
Transport
C++, OpenMP, MPI

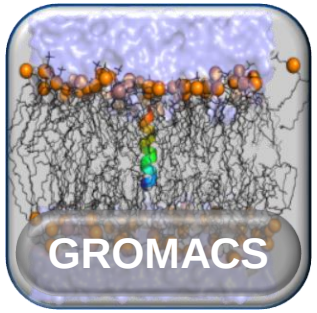

Space Weather
ML, Python, C++,
OpenMP, MPI, CUDA

Seismic Imaging
C++, OpenMP, MPI,
CUDA, GASPI

Computational Fluid
Dynamics
Fortran, MPI

Earth Systems
Modelling
C, C++, Fortran, MPI

More than 40 HPC SW Components

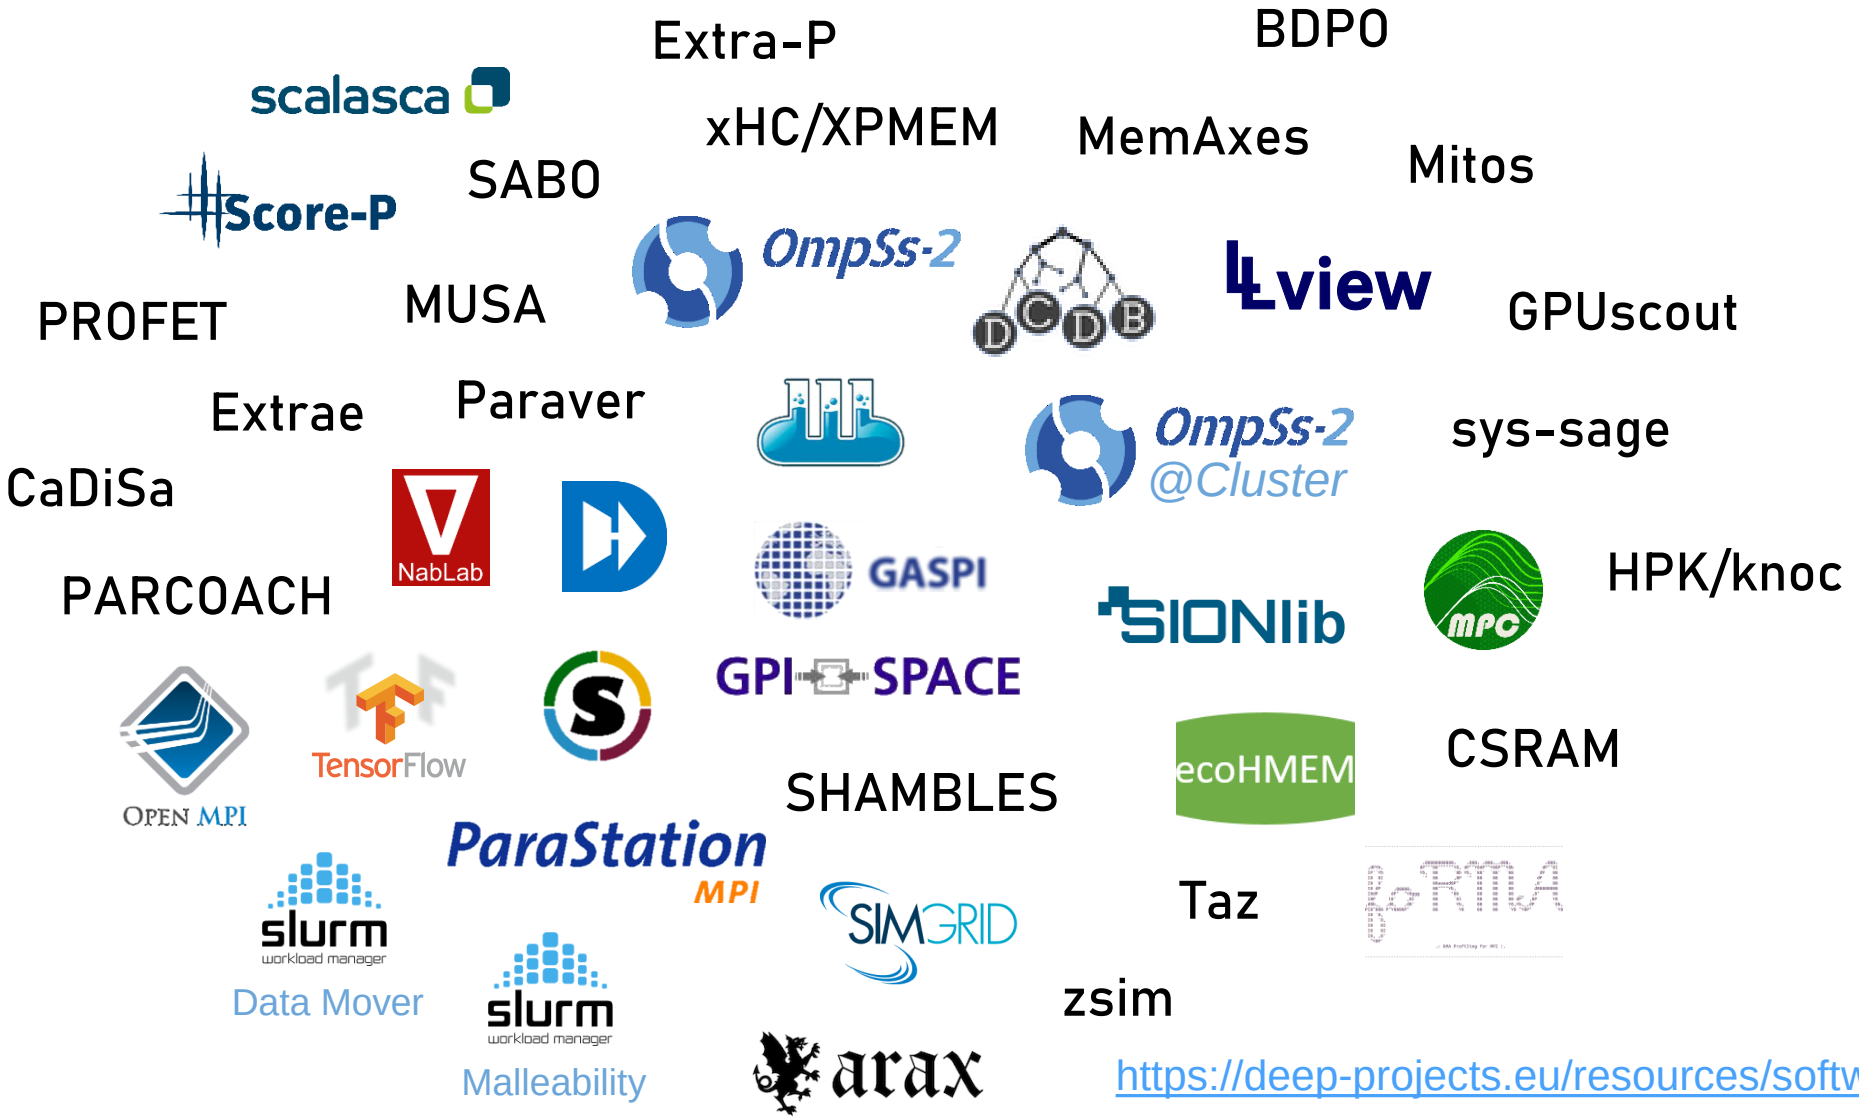

Funding Acknowledgement

EuroHPC Joint Undertaking

SPONSORED BY THE

Federal Ministry of Education and Research

Swedish Research Council

The DEEP Projects have received funding from the European Commission's FP7, H2020, and EuroHPC JU Programmes, under Grant Agreements n° 287530, 610476, 754304, and 955606. The DEEP-SEA project receives also support from Belgium, France, Germany, Greece, Spain, Sweden, and Switzerland

www.deep-projects.eu

@DEEPprojects