CoE RAISE Review

9th of September 2022

WP2 "AI- and HPC-Cross Methods at Exascale"

Morris Riedel, UOI

WP2 Objectives

- Build the AI / HPC Exascale backbone by providing the hardware & software infrastructure needed for the implementation of the WP3/WP4 use-cases
- Provide access to available production systems and new prototypes & disruptive technologies for testing, porting, and benchmarking
- Develop and tune cross-sectional HPC/AI methods with WP3/WP4 use cases
- Co-design, implement, and deploy a unique AI framework for Exascale

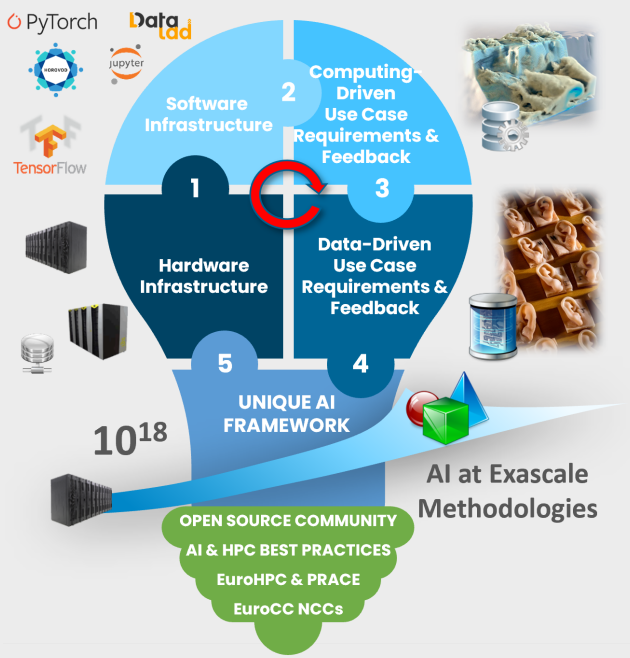

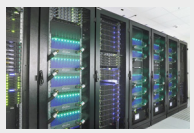

Partners and Tasks of WP2

Partner	FZJ	UOI	RWTH	BSC	CERN	BULL	RTU	FM
PM	43	13	10	8	8	8	22	12

Task	Title	Lead	Duration	Status
2.1	Modular and heterogeneous supercomputing architectures	BSC	M1 – M36	Ongoing
2.2	Hardware prototypes	FZJ	M1 – M18	Done
2.3	Benchmarking on disruptive technologies	FZJ	M19 – M36	Ongoing
2.4	Software design of a unique AI framework	UOI	M4 – M36	Ongoing
2.5	Cross-Sectional AI Methods	UOI	M3 – M36	Ongoing

Deliverables of WP2 (1/2)

ID	Title	Due	Lead	Status
D2.1	Best practice guidelines/tutorials for MSA/heterogeneous systems	M2	BSC	**Submitted**
D2.2	Report on porting and performance analysis	M12	BSC	**Submitted**
D2.3	Report on porting and performance analysis	M24	BSC	**Not started**
D2.4	Report on porting and performance analysis	M36	BSC	**Not started**
D2.5	Best practice guidelines / tutorials prototype	M2	FZJ	**Submitted**
D2.6	Report on support activities	M6	FZJ	**Submitted**
D2.7	Report on support activities	M18	FZJ	**Submitted**
D2.8	Report on benchmarking AI technologies (QA) and on support activities	M24	FZJ	**Not started**
D2.9	Report on benchmarking AI technologies (QA) and on support activities	M36	FZJ	**Not started**
D2.10	Monitoring Report	M18	UOI	**Submitted**

Deliverables of WP2 (2/2) and Milestones

ID	Title	Due	Lead	Status
D2.11	Monitoring Report	M36	UOI	Not started
D2.12	Software Layout Plan for a unique AI Framework	M9	UOI	Submitted
D2.13	Software Layout Plan for a unique AI Framework	M26	UOI	Not started
D2.14	Report on Novel AI Technologies	M12	UOI	Submitted
D2.15	Report on Novel AI Technologies	M24	UOI	Not started
D2.16	Report on Novel AI Technologies	M36	UOI	Not started

ID	Title	Due	Lead	Status
MS2	AI/HPC Methods	M7	UOI	Achieved
MS4	Technical implementations functional	M24	FZJ	Not yet achieved
MS6	All final reports	M36	FZJ	Not yet achieved

Task 2.1

Modular and heterogeneous supercomputing architectures

Task 2.1 – Modular and heterogeneous supercomputing …

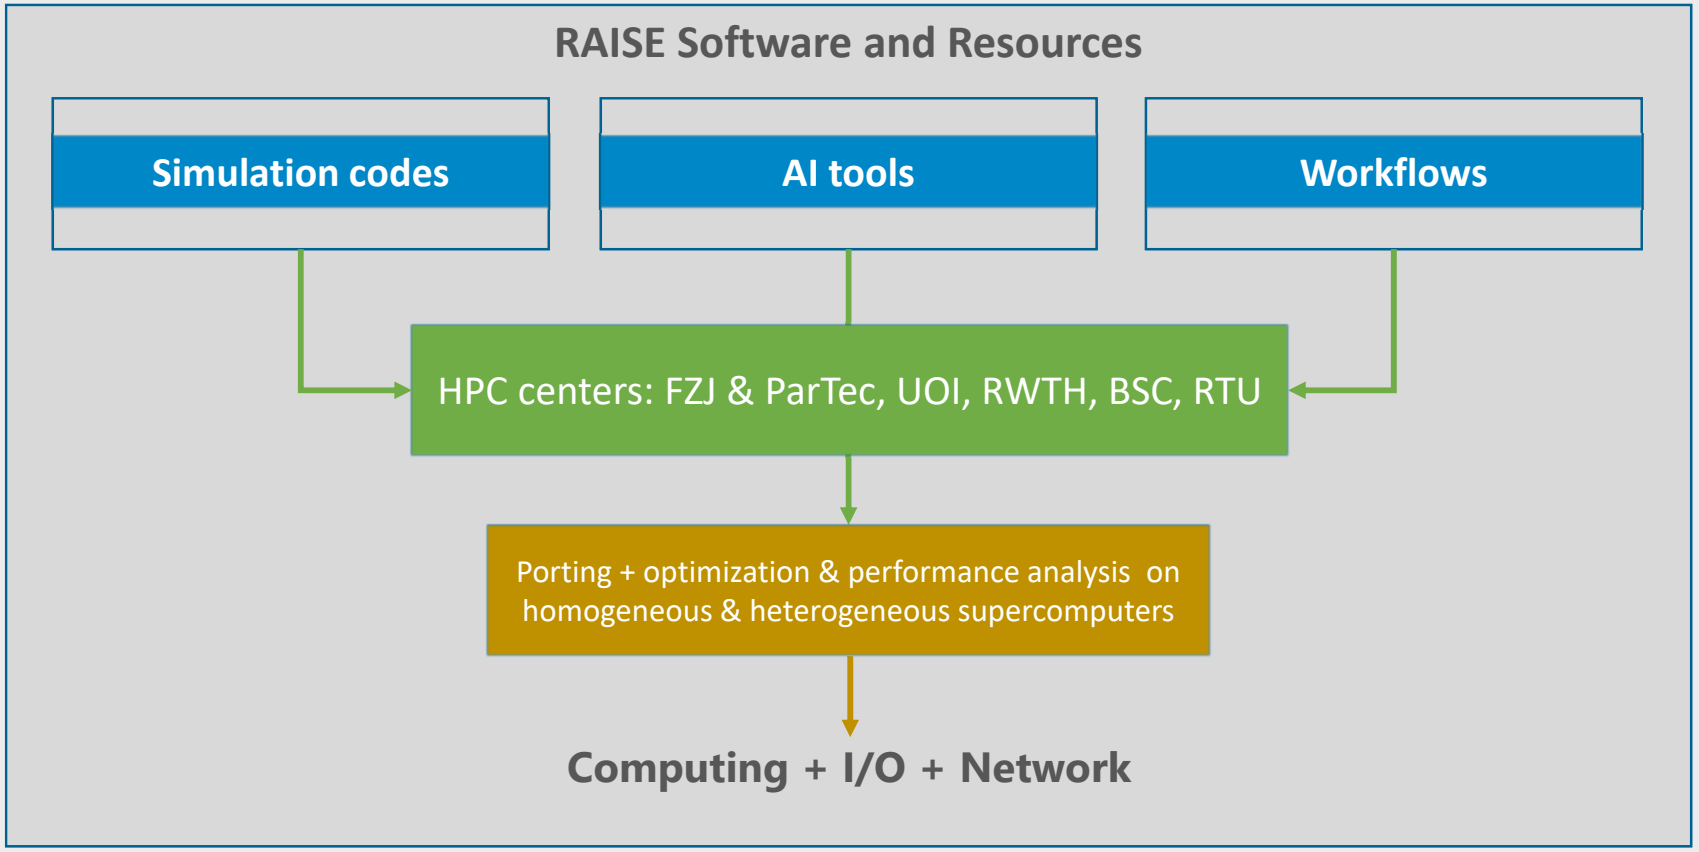

Task 2.1 – Modular and heterogeneous supercomputing …

Where do the supercomputing resources come from?

- Local supercomputers
  - TIER-3, TIER-2: UOI, RWTH, RTU
  - TIER-1, TIER-0: FZJ, BSC
- PRACE specific access calls to COEs:
  - 2021-1: 1.6M
  - 2021-2: 0.6M
  - 2022-1: 0.5M)
- Competitive access calls

Partner	HPC System	Hours
CERFACS	GENCI Jean Zay	20k GPU
CERFACS	GENCI Jean Zay	20k GPU
CYI	Cyclone, acNational	1M CPU
CYI	VEGA, EuroHPC	12M CPU
FM	Cloud, regional center	35k CPU
FZJ	JURECA-DC GPU, JARA	8.6M GPU
FZJ / CERN / UOI	JURECA-DC GPU, JARA	8.3M GPU
FZJ / CERN / UOI	JUPSI, JARA	5 hours
RTU	In-kind for RAISE	550k CPU
RWTH	JUWELS, GCS	6M CPU

Task 2.1 – Modular and heterogeneous supercomputing …

Deliverable D2.1: Best practice guidelines and tutorials for the various HPC systems

- 4 countries
- 8 systems

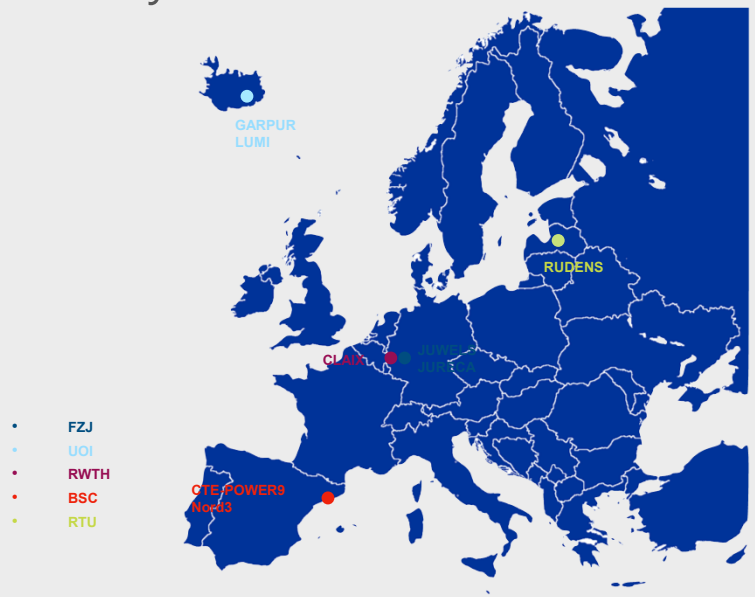

CTE-Power, Nord3

JUWELS

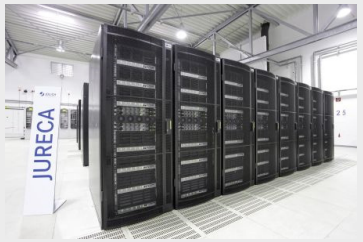

JURECA

RUDENS

CLAIX

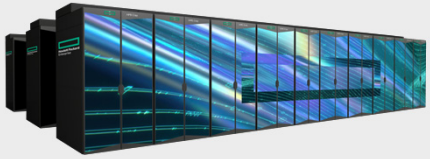

GARPUR & LUMI

Task 2.1 – Modular and heterogeneous supercomputing …

Deliverable D2.1: Report on porting and performance engineering activities

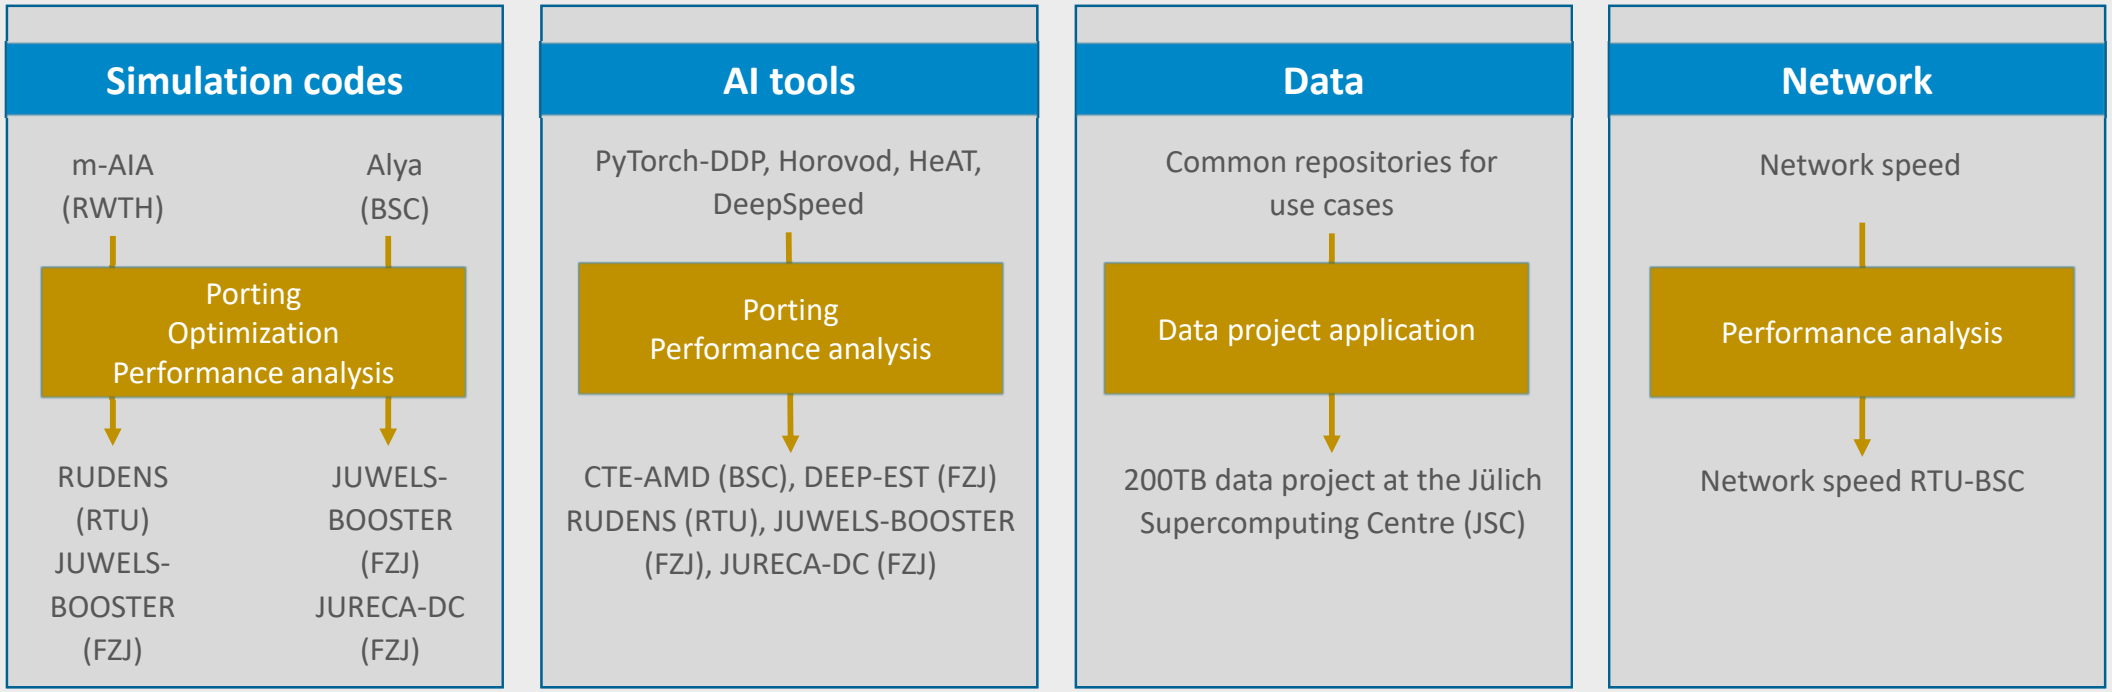

Task 2.1 – AI Benchmarking: PyTorch-DDP

Parallel performance using PyTorch-DDP on JUWELS Booster (4 x NVIDIA A100 / node):

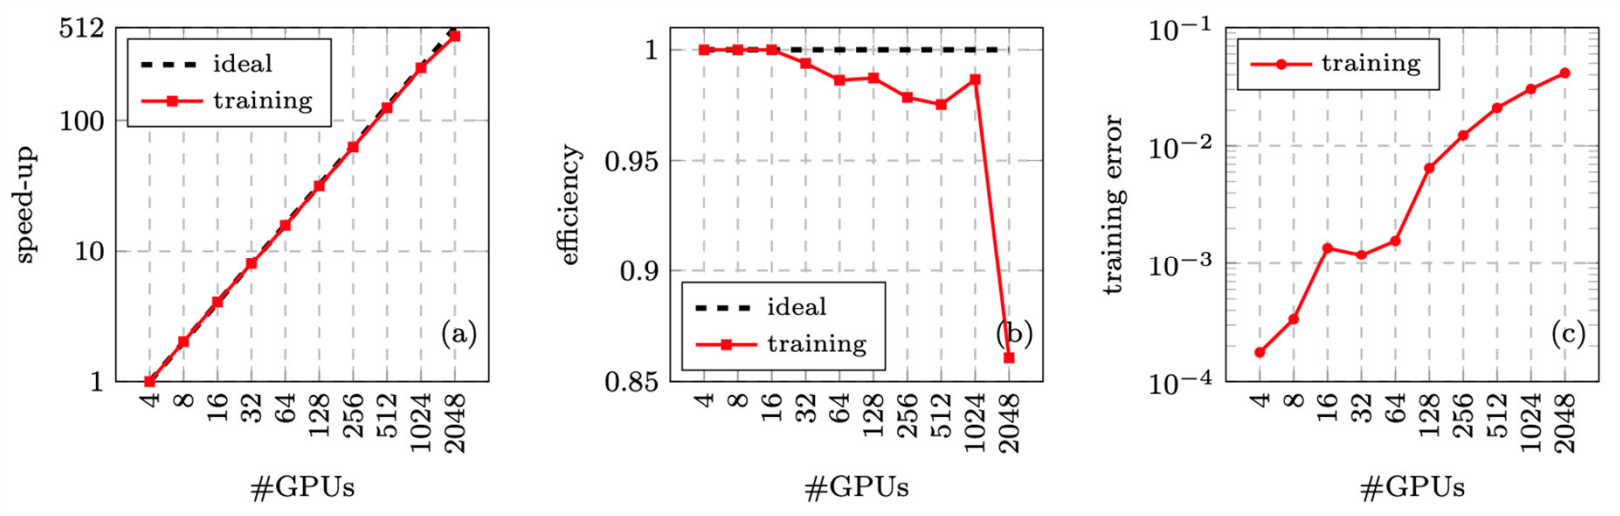

Application use case Task 3.1: Autoencoders for Turbulent Boundary Layer Flows

Task 2.1 – AI Benchmarking: ImageNet

Performance of Horovod and PyTorch-DDP (with DALI dataloader) on up to 1,024 GPUs:

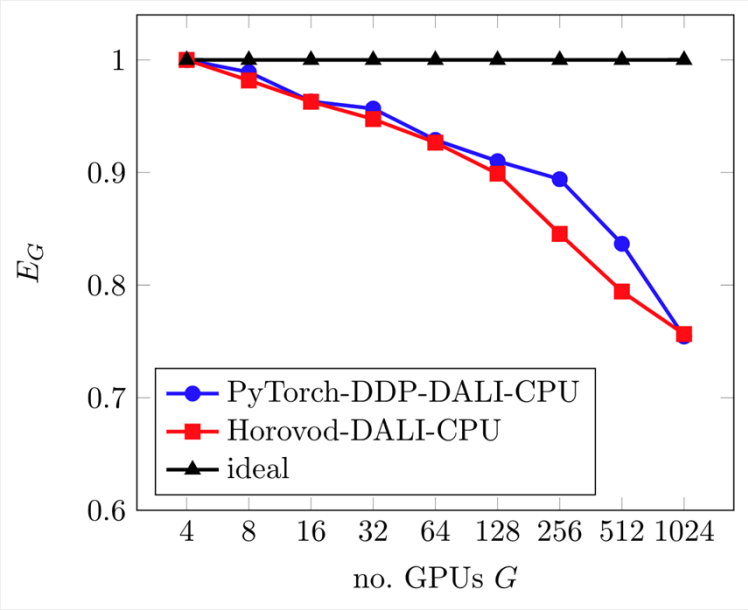

Parallel efficency over number of GPUs

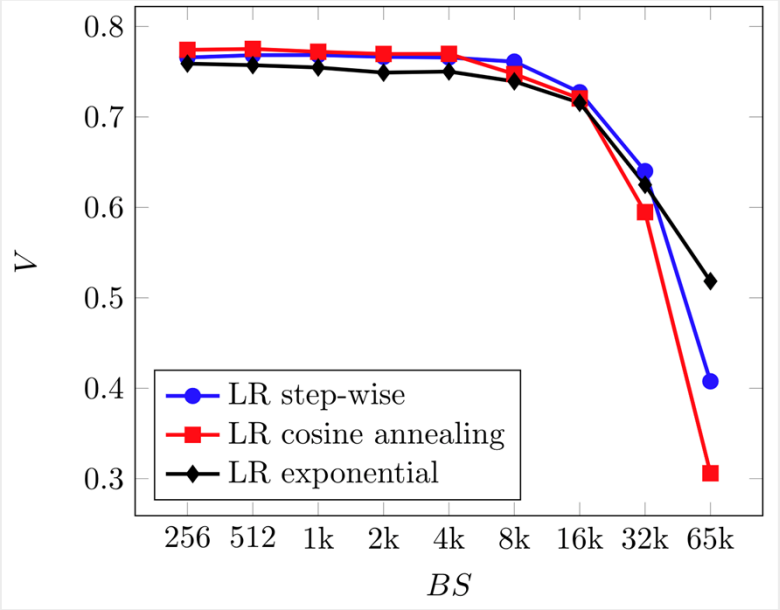

Validation accuracy over batch size showing impact of learning rate schedulers

Task 2.1 – AI Benchmarking: Batchsize Problem

Addressing the batch size problem by changing the batch size during the training: start small, then increase

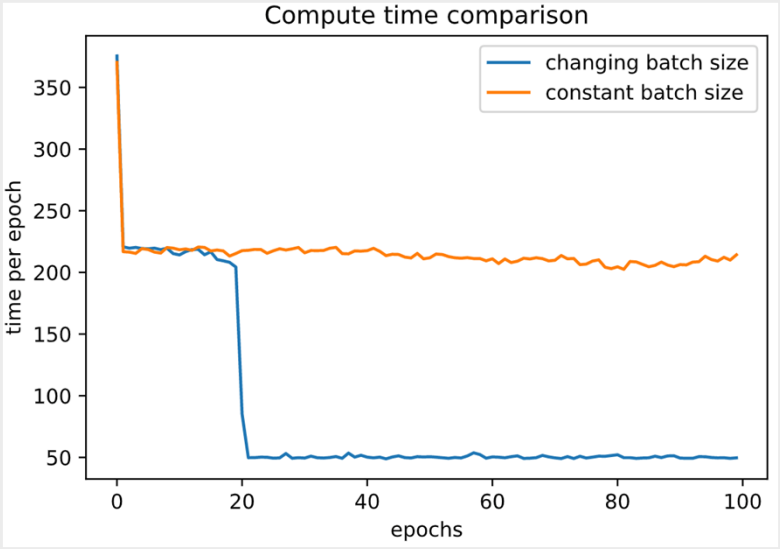

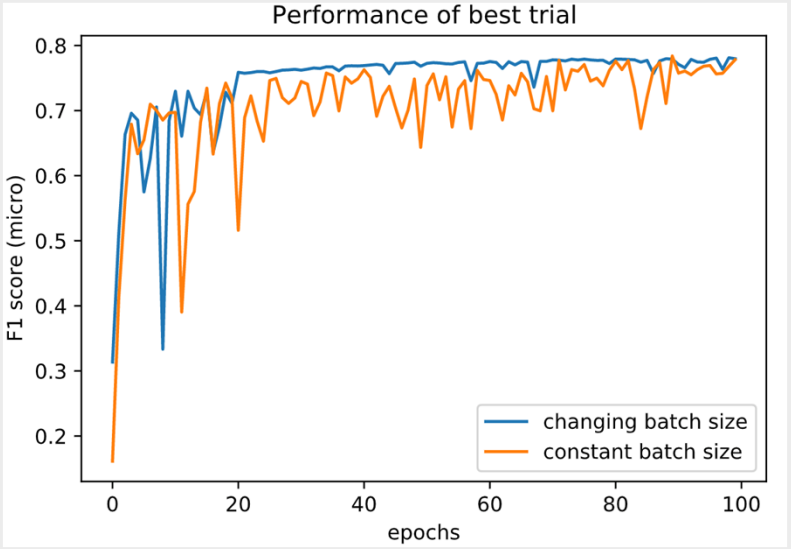

Task 2.1 – AI Benchmarking: Application Example (Task 4.2)

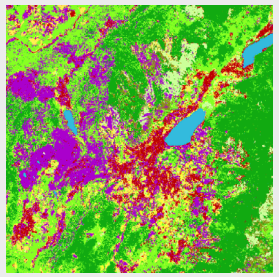

- Land Cover classification with ResNet50, TensorFlow + Horovod
- Intensive transformations are used as pre-processing
- Scaled on up to 256 GPUs on JUWELS-Booster with efficiency ~80%

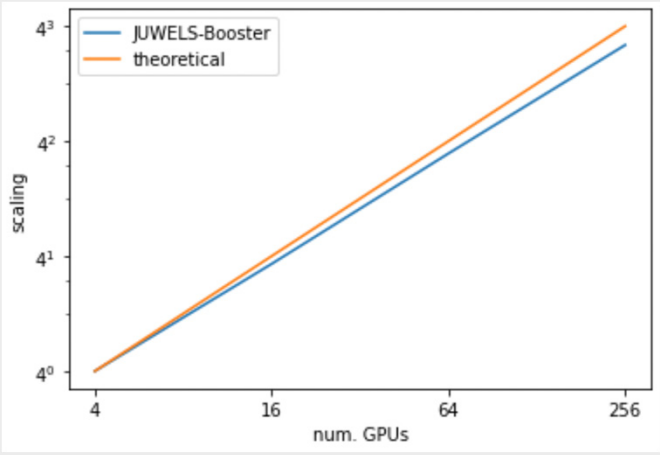

R. Sedona, C. Paris, L. Tian, M. Riedel, G. Cavallaro , AN AUTOMATIC APPROACH FOR THE PRODUCTION OF A TIME SERIES OF CONSISTENT LAND-COVER MAPS BASED ON LONG-SHORT TERM MEMORY, IGARSS2022

Task 2.2

Hardware prototypes

Task 2.2 – Status

Prototype systems:

- DEEP-EST (FZJ)
  - Intel CPU + NVIDIA GPU + FPGA
- CTE-AMD (BSC)
  - AMD CPU + AMD GPU
- JUAWEI (FZJ)
  - ARM + Intel CPU
- CTE-ARM (BSC)
  - ARM
- HUAWEI (BSC)
  - ARM

Porting + Testing + Benchmark + Optimization

- Simulation codes:
  - m-AIA (Task 3.1)
  - Alya (Task 3.2)
  - Basilisk (Task 3.5)
- ML Frameworks:
  - PyTorch-DDP
  - Horovod
  - HeAT
  - DeepSpeed
  - Hyperparameter tuning

Task 2.2 – Simulation Code Benchmarking: Alya

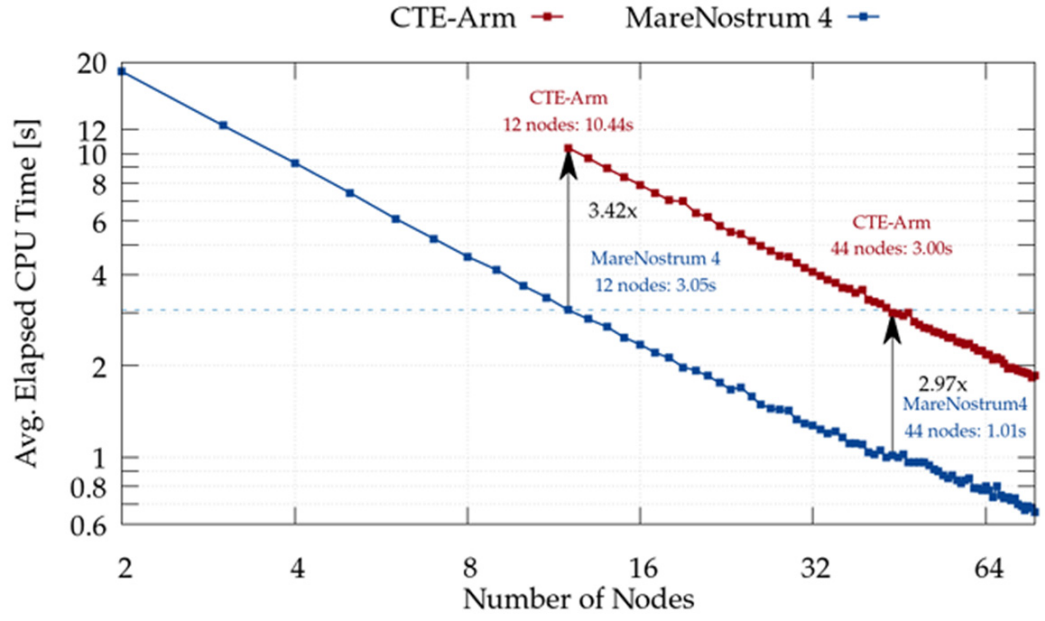

No. cores	Time [s]	Speed-up	Efficiency
128	1202.70	1.0	100%
256	658.85	1.8	91%
512	332.72	3.6	90%
1,024	182.57	6.6	82%

Task 2.2 – MSA Prototype System DEEP-EST Examples

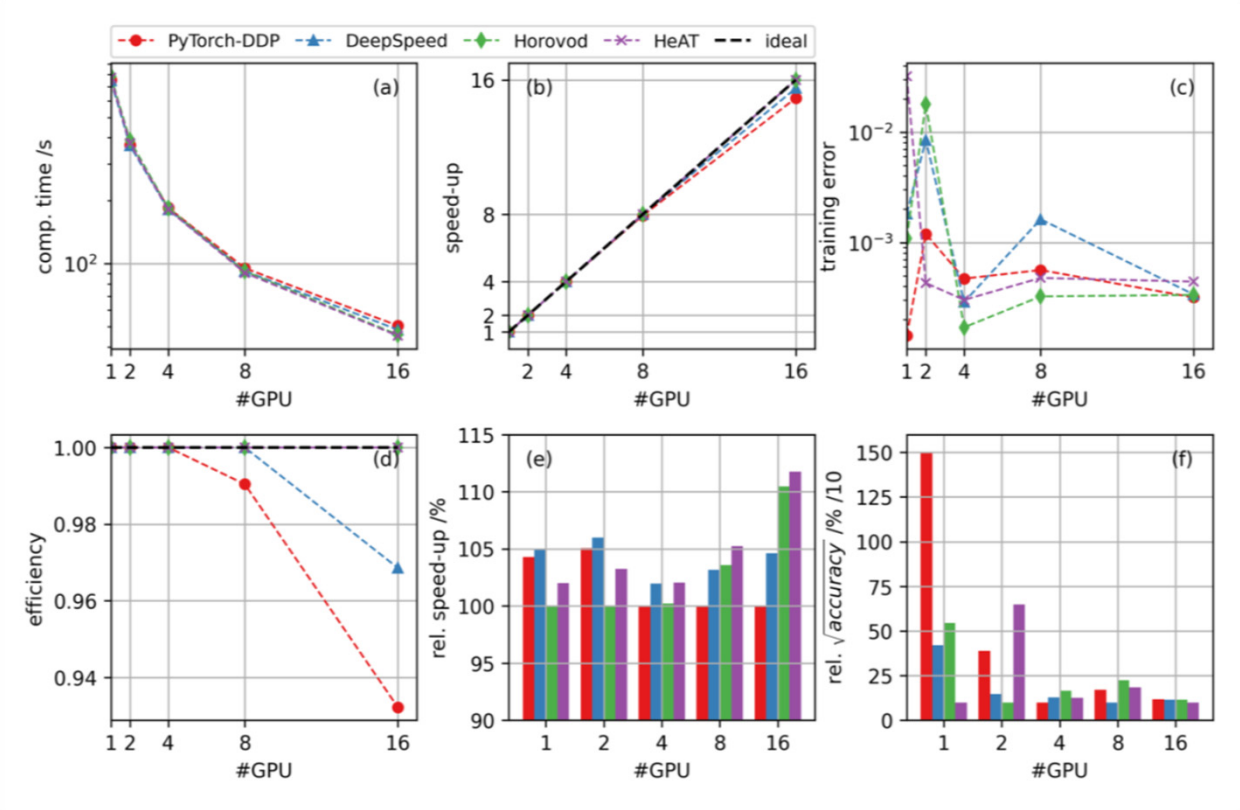

Actuated Turbulent Boundary Layer dataset

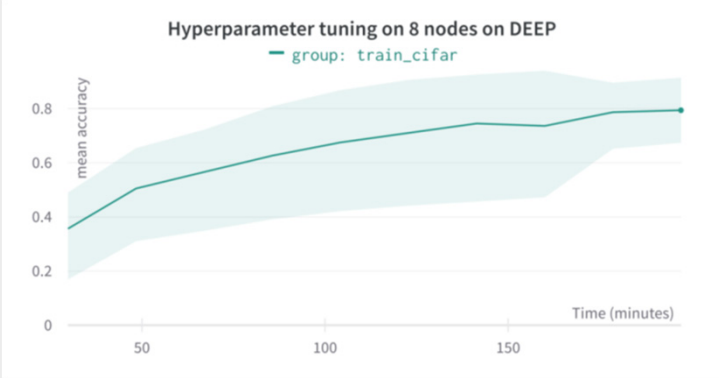

Hyperparameter – Tuning using Ray Tune

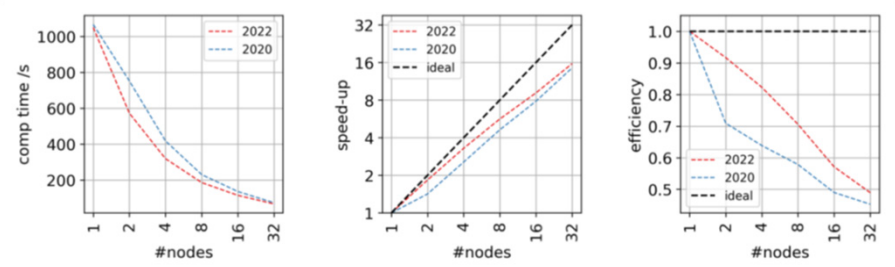

Basilisk Simulation Sciences Code

Task 2.2 – Prototype Systems CTE-AMD Examples

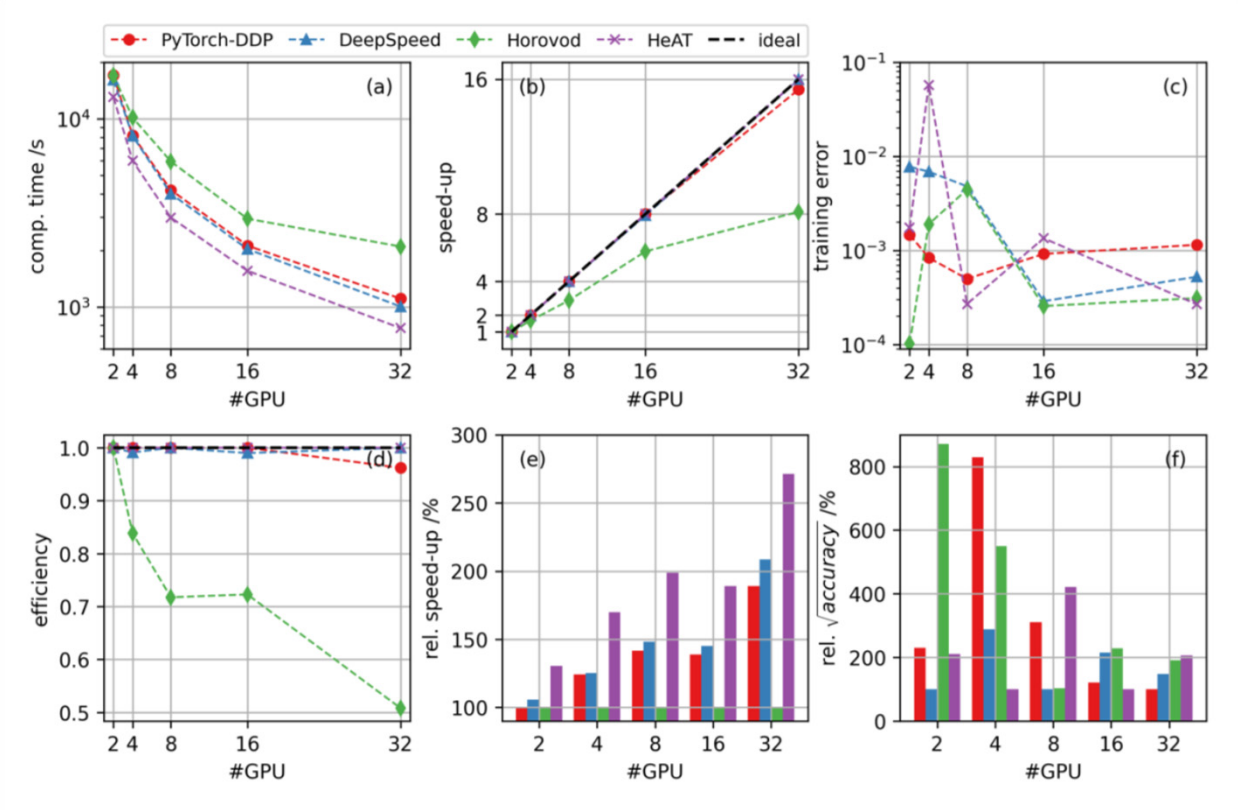

Actuated Turbulent Boundary Layer dataset

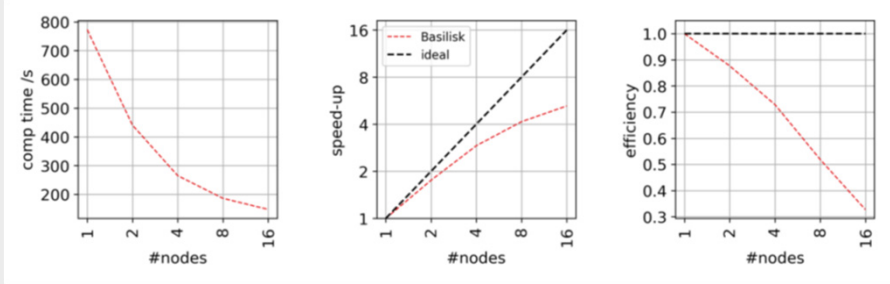

Basilisk Simulation Sciences Code

CTE-AMD

System Overview

CTE-AMD is a cluster based on AMD EPYC processors, with a Linux Operating System and an Infiniband interconnection network.

It has the following configuration:

- 1 login node and 33 compute nodes, each of them:
  - 1 x AMD EPYC 7742 @ 2.250GHz (64 cores and 2 threads/core, total 128 threads per node)
  - 1024GiB of main memory distributed in 16 dimms x 64GiB @ 3200MHz
  - 1 x SSD 480GB as local storage
  - 2 x GPU AMD Radeon Instinct MI50 with 32GB
  - Single Port Mellanox Infiniband HDR100
  - GPFS via two copper links 10 GBit

Task 2.3

Benchmarking on disruptive technologies

Task 2.3 – Benchmarking on disruptive technologies

First 5,000 Qubits Quantum Annealer (JUPSI) in Europe (FZJ)

- 5 hours of compute time for RAISE project
- problem size still very limited -> hybrid/modular supercomputing approach

Applications:

- Quantum accelerated hyperparameter tuning
- Quantum clustering of energy particles
- Quantum classification of satellite images

First results:

- predicting the learning curves of neural network training with Quantum Support Vector Regression

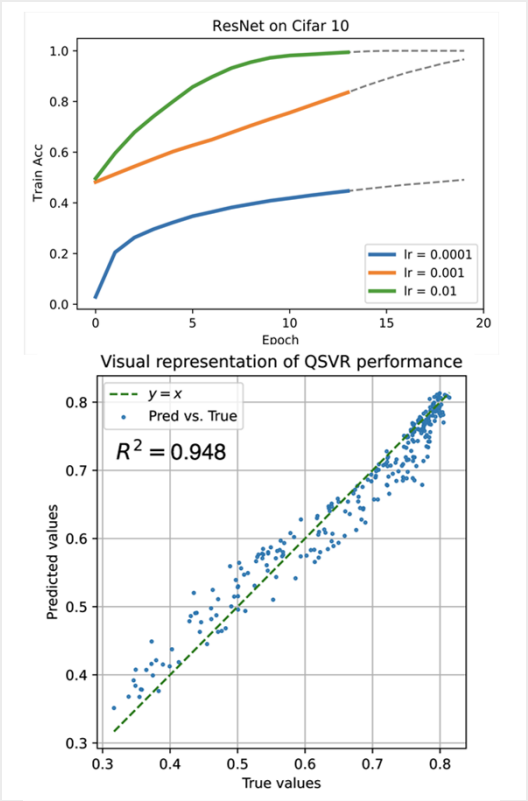

Solution

Data problem

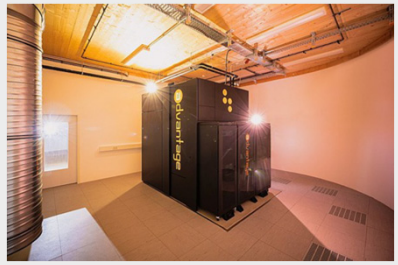

Task 2.4

Software design of a unique AI framework

Task 2.4 – Status: Process towards Framework Realization

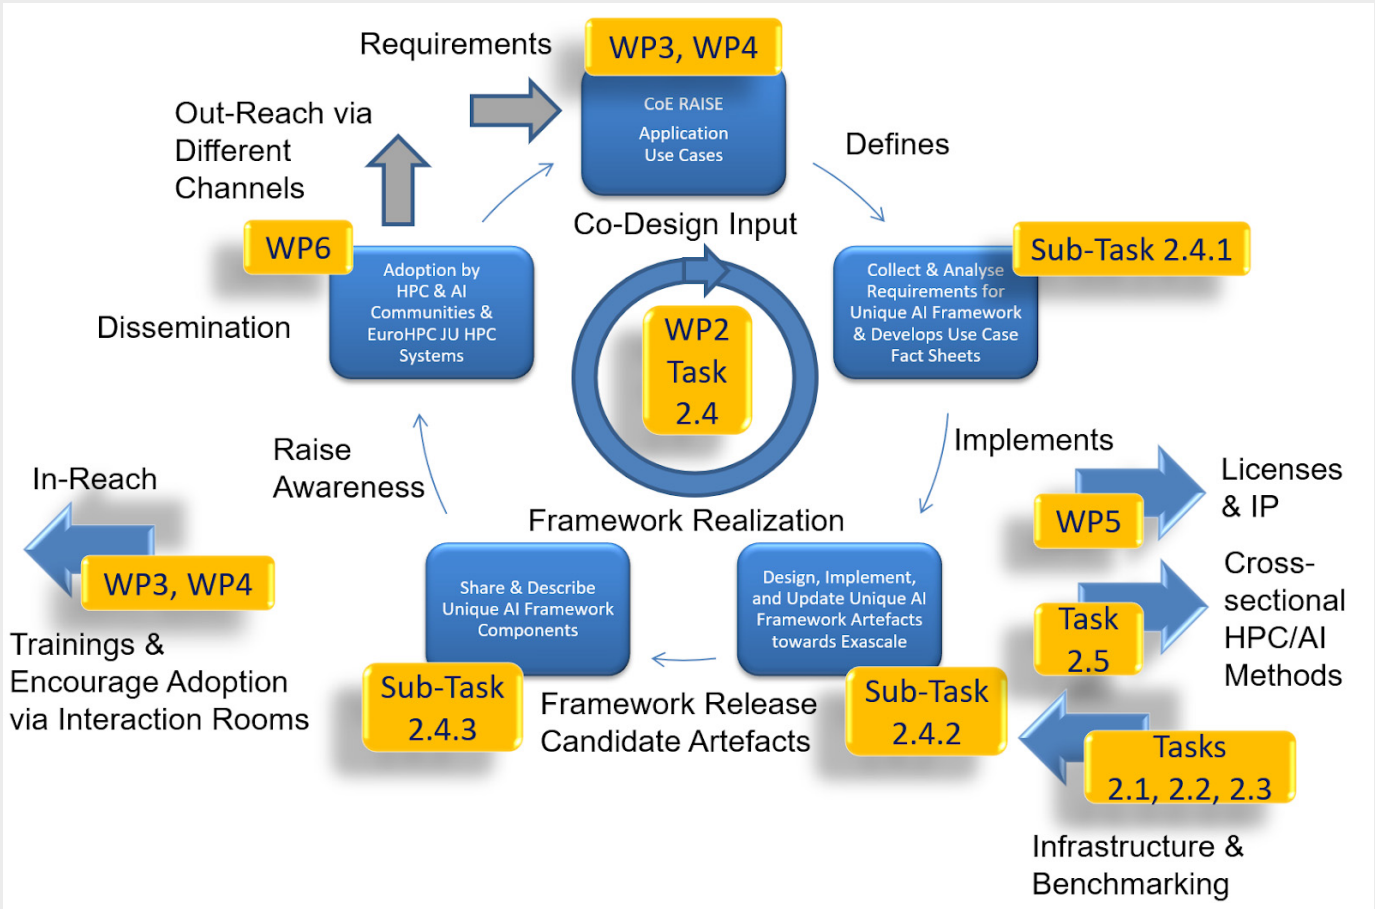

Task 2.4 – Co-Design Framework: Fact Sheets

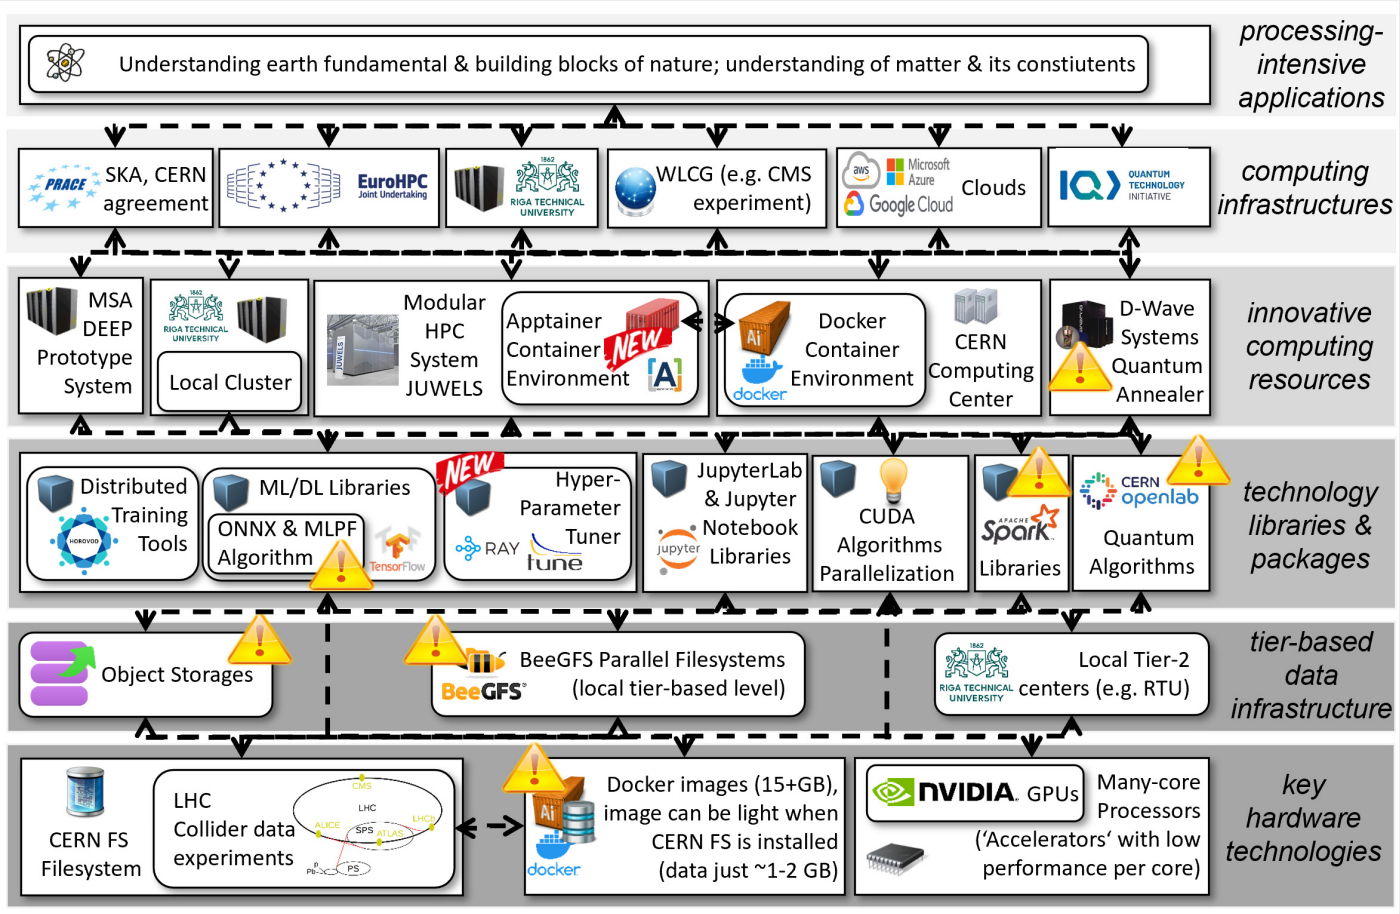

Example: Detailed Knowledge of Modules Necessary

- Modules
  - Vary heavily between different HPC systems
  - 2-3 Days/Months spend by researchers for getting the right environment / HPC system
  - Goal: UAIF simplify setup of components
  - E.g., automated job script generator for right module setup
  - E.g., re-usable scripts

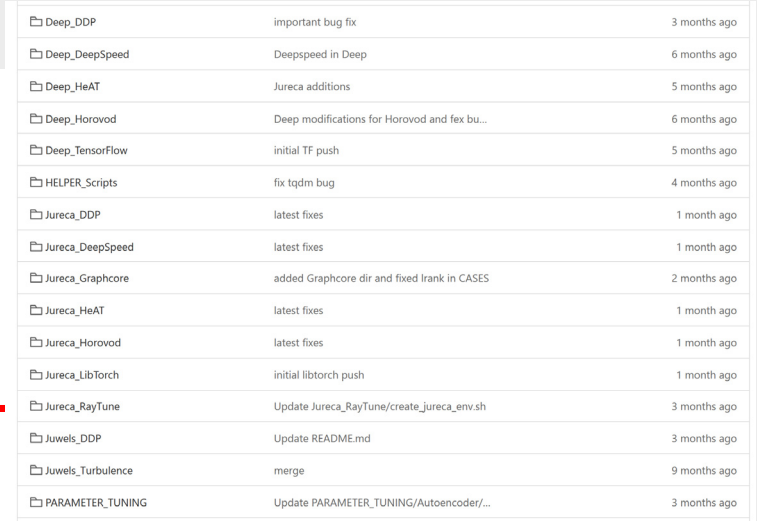

Deep_DDP	important bug fix	3 months ago
Deep_DeepSpeed	Deepspeed in Deep	6 months ago
Deep_HeAT	Jureca additions	5 months ago
Deep_Horovod	Deep modifications for Horovod and fex bu...	6 months ago
Deep_TensorFlow	initial TF push	5 months ago
HELPER_Scripts	fix tqdm bug	4 months ago
Jureca_DDP	latest fixes	1 month ago
Jureca_DeepSpeed	latest fixes	1 month ago
Jureca_Graphcore	added Graphcore dir and fixed Irank in CASES	2 months ago
Jureca_HeAT	latest fixes	1 month ago
Jureca_Horovod	latest fixes	1 month ago
Jureca_LibTorch	initial libtorch push	1 month ago
Jureca_RayTune	Update Jureca_RayTune/create_jureca_env.sh	3 months ago
Juwels_DDP	Update README.md	3 months ago
Juwels_Turbulence	merge	9 months ago
PARAMETER_TUNING	Update PARAMETER_TUNING/Autoencoder/...	3 months ago

```
#!/usr/bin/env bash

# Slurm job configuration
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=4
#SBATCH --cpus-per-gpu=20
#SBATCH --account=hai_so2sat
#SBATCH --output=output.out
#SBATCH --error=error.er
#SBATCH --time=6:00:00
#SBATCH --job-name=BENTF2
#SBATCH --gres=gpu:1 --partition=booster

#load modules
ml Stages/2020  GCC/9.3.0  OpenMPI/4.1.0rc1
ml Horovod/0.20.3-Python-3.8.5
ml TensorFlow/2.3.1-Python-3.8.5
#activate my virtualenv
#source /p/project/joaiml/remote_sensing/rocco_sedona/ben_TF2/scripts/env_tf2_juwels_booster/bin/activate

#export relevant env variables
#export CUDA_VISIBLE_DEVICES="0,1,2,3"

#run Python program
srun --cpu-bind=none python -u train_hvd_keras_aug.py
```

Already available for the community: https://gitlab.jsc.fz-juelich.de/CoE-RAISE/FZJ/ai-for-hpc-oa

Task 2.4 – Realization Framework: Requirements & Revisions

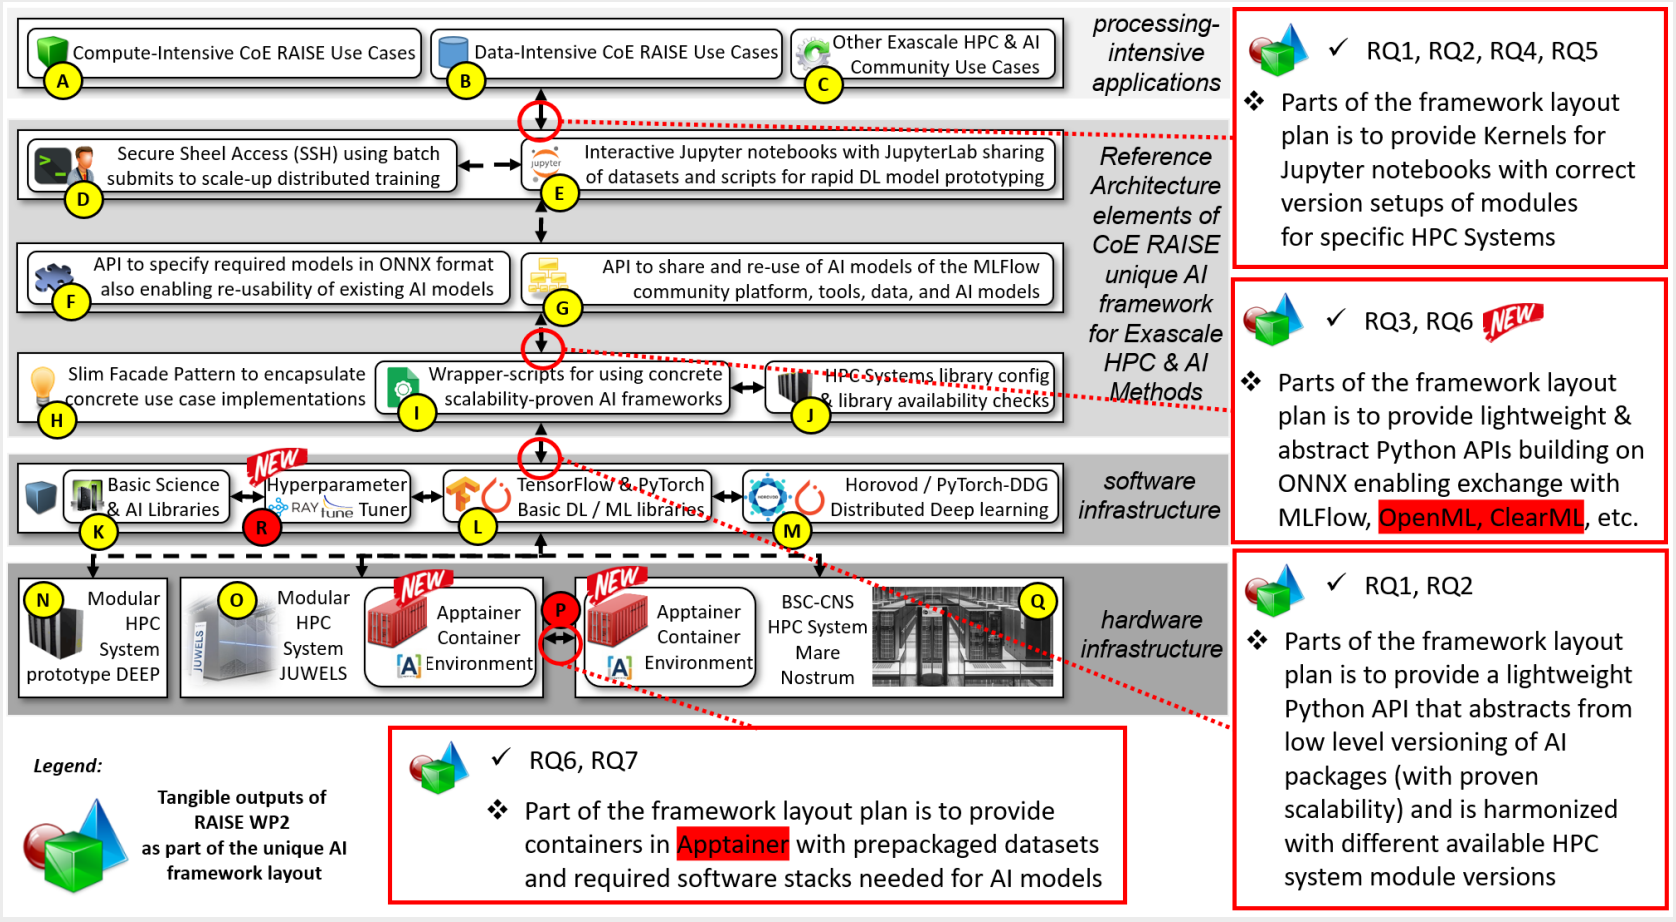

Continuously updating!

M. Riedel, M. Book, H. Neukirchen, G. Cavallaro, A. Lintermann, 'Practice and Experience using High Performance Computing and Quantum Computing to Speed-up Data Science Methods, MIPRO 2022, Croatia

Unique AI Framework Overview

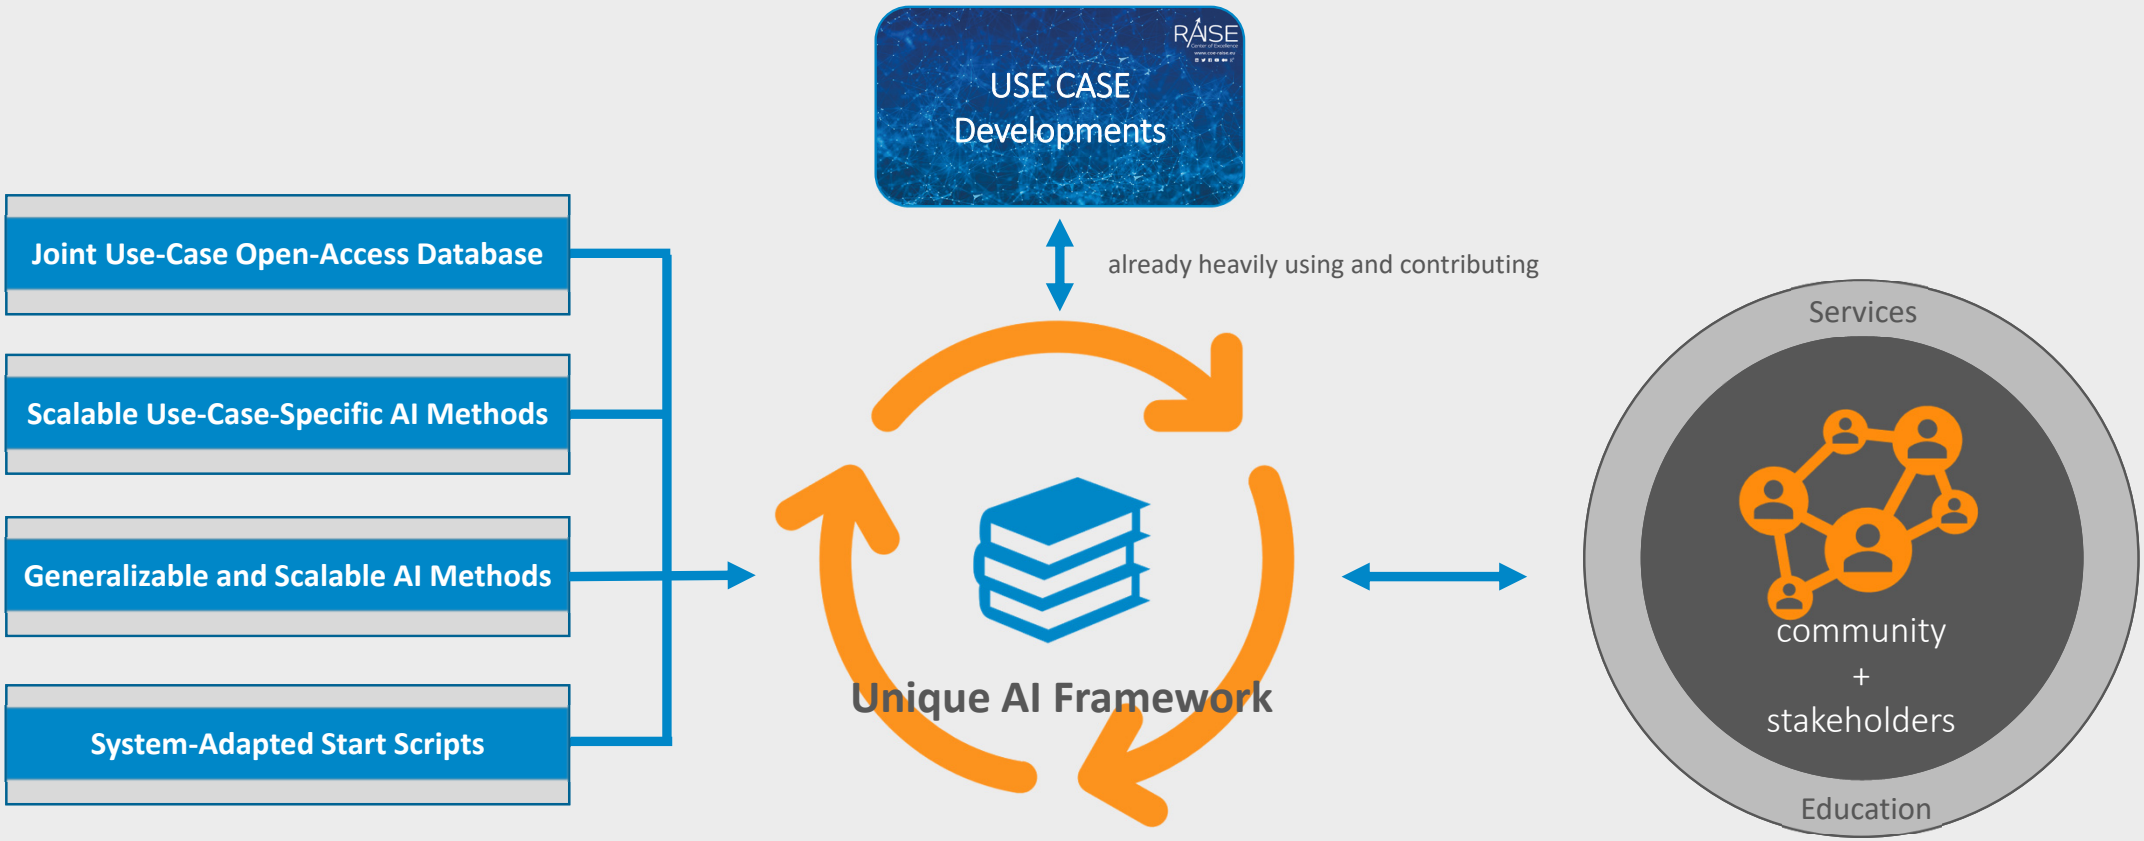

Task 2.5

Cross-Sectional AI Methods

Task 2.5 – Status: Process building AI/HPC Methods

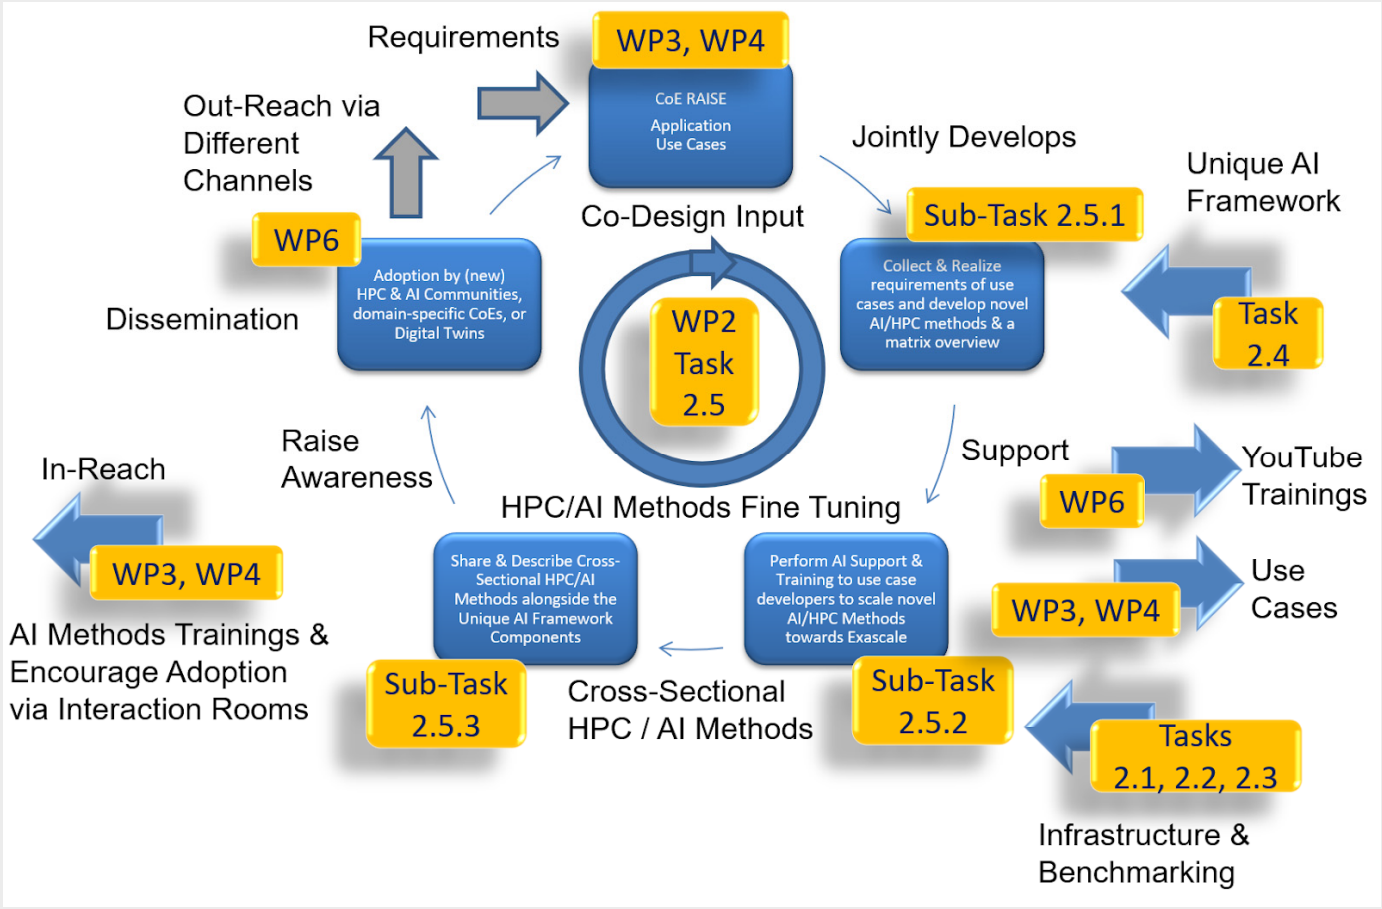

Task 2.5 – Interaction Rooms with WP3/WP4 Teams

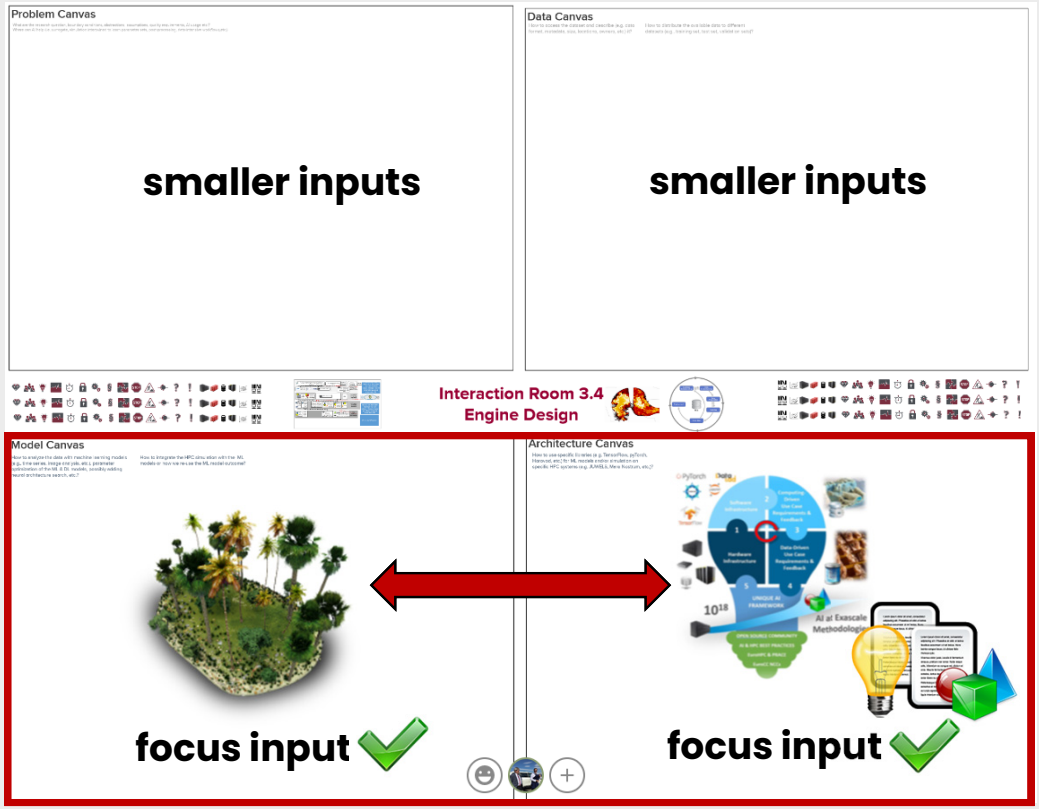

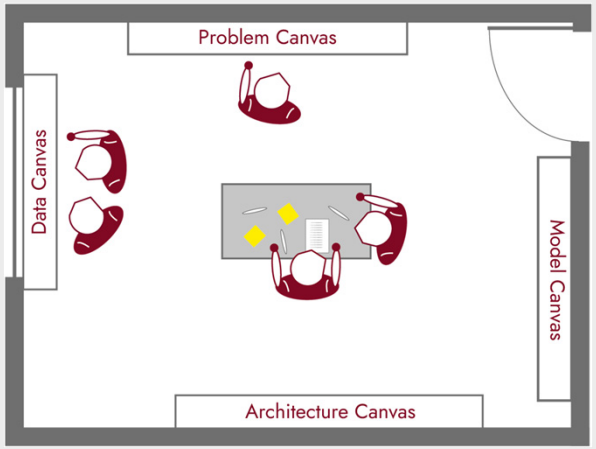

IR Mural Links

IR3.1 Turbulent Flow: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621377866397/8613c384d54f66fb5e78599ff307a4ce8a9090c0?sender=u15e3008bb41d6628a5bb5701

IR3.2 Clean Energy: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621377887905/cb44cca3eedd3bb9964fbfa36af16b1bfcce085f?sender=u15e3008bb41d6628a5bb5701

IR3.3 Reactive Flows: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621377959022/0c363886f24833ecb19b025d87324b57fd50e2db?sender=u15e3008bb41d6628a5bb5701

IR3.4 Engine Design: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621377976343/8d7aba6be09af3b2ffd305d2f709c53661ac889d?sender=u15e3008bb41d6628a5bb5701

IR3.5 Coating: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621377991014/7a5d7e1eaf230178342d1e1d4a84d656d9055d52?sender=u15e3008bb41d6628a5bb5701

IR4.1 Fundamental Physics: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621378007555/6f0d5285fcaec5eafa515bd6676e84d8b4879d39?sender=u15e3008bb41d6628a5bb5701

IR4.2 Seismic Imaging: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621378023838/a0b9503abb837ae3e28af4bb8d9adbec33874998?sender=u15e3008bb41d6628a5bb5701

IR4.3 Manufacturing: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621378038069/93df6fa7a41093f4eaae7be9d72979de2ba42b9d?sender=u15e3008bb41d6628a5bb5701

IR4.4 Sound Engineering: https://app.mural.co/t/matthiasbook8855/m/matthiasbook8855/1621378050431/b5fa122190024040 59f90a4bbb0101fa379a8503?sender=u15e3008bb41d6628a5bb5701

In addition to WP2 Monthly Meetings & Monthly Trainings

Book, Riedel, Neukirchen, Erlingsson: Facilitating Collaboration in Machine Learning and High-Performance Computing Projects with an Interaction Room, International Workshop on Software Engineering for eScience, IEEE 2022

Task 2.5 – Interaction Room Example: Task 3.1

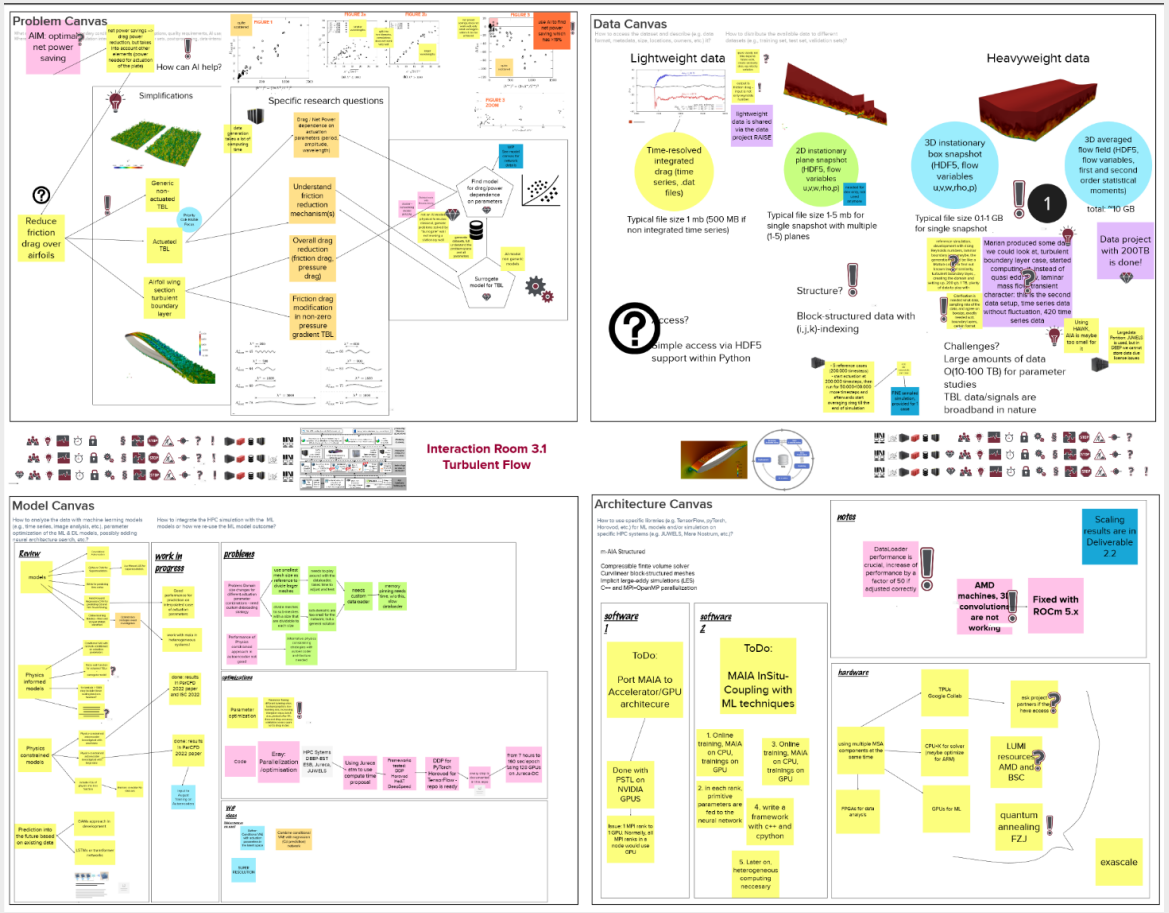

Task 2.5 – Cross HPC/AI Methods Table

Interaction Room results:

- Update of Matrix
- Components relatively constant & common
- Methods change & new methods added (e.g., Transformers, RFs)

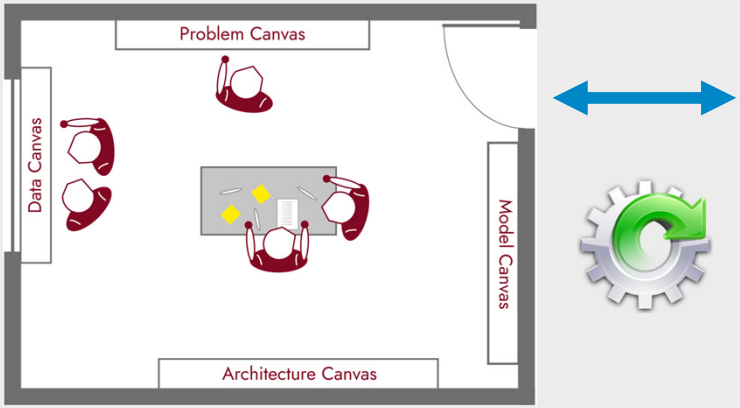

Use Case	AE	PINN	ANNs		CNN		NO	GNN		RNN		GAN	TF				SVM	RF
Details	*CAE*		*ANN*	*RBF-ANN*	*U-Net*	*RES NET*	*FNO*	*MLPF*	*GAT*	*LSTM*	*GRU*	*WGAN*	*MViT*	*ViViT*	*Swin*			
AI for turbulent boundary layers	X	X	X									X						
AI for wind farm layout optimization				X												X		
AI for data-driven models in reacting flows					X				X									
Smart models for next generation aircraft engine design					X				X									
AI for wetting hydrodynamics	X	X					X			X								
Event reconstruction and classification at the CERN HL-LHC use case								X										
Seismic imaging with remote sensing for energy applications	X	X				X	X			X	X					X	X	X
Detect-free metal additive manufacturing	X		X									X	X	X	X			
Sound Engineering	X		X															

WP 2 – Summary, Conclusions, and Outlook

AL0

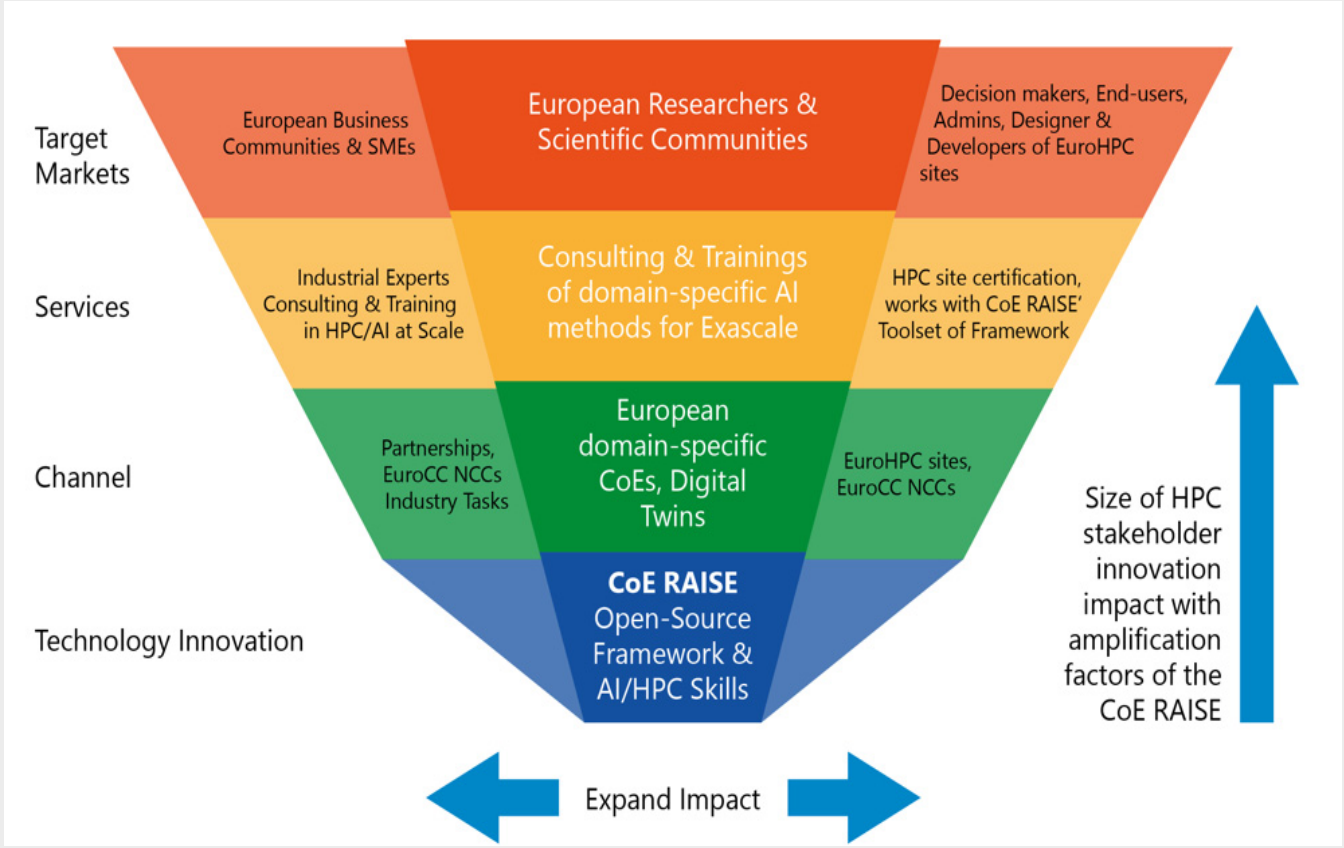

Slide 33

AL0 Where is the summary and the conclusions?
Andreas Lintermann; 2022-09-06T10:11:48.862

drive. enable. innovate.

RAISE
Center of Excellence

The CoE RAISE project has received funding from
the European Union's Horizon 2020 –
Research and Innovation Framework Programme
H2020-INFRAEDI-2019-1 under grant agreement no. 951733

Follow us: