UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING, MECHANICAL ENGINEERING AND COMPUTER SCIENCE

RAISE Center of Excellence

EURO

WP2 AI- & HPC-Cross Methods at Exascale – Monthly Meeting

Prof. Dr. – Ing. Morris Riedel et al.
School of Engineering & Natural Sciences, University of Iceland

2021-03-15, RAISE WP2 Monthly Meeting March 2021, Online

@ProfDrMorrisRiedel @Morris Riedel @MorrisRiedel @MorrisRiedel
https://www.youtube.com/channel/UCWC4VKHmL4NZgFfKoHtANKg
morris@hi.is

WP2 February Meeting – Welcome & Agenda

1. Approval of minutes from Monthly Meeting February 2021
   - (All), ~5 Min
2. WP2 Updates & Overview of WP2 Fact Sheet Process of RAISE & WP3/WP4
   - (Morris Riedel), ~15 Min
3. Fact Sheet Status Presentations & Discussions
   - (Morris Riedel & Use Case Teams), ~10 Min
4. Follow-up Resource Provisioning
   - (Guillaume Houzeaux), ~15 Min
5. Follow-Through Check: Meeting duration (1 hour monthly long enough?)
   - (All), ~5 Min
6. Next Monthly Meeting April & AOB
   - Doodle & Discussions, ~10 Min

Agenda Item (1) – Minutes Approval – Meeting Feb 2021

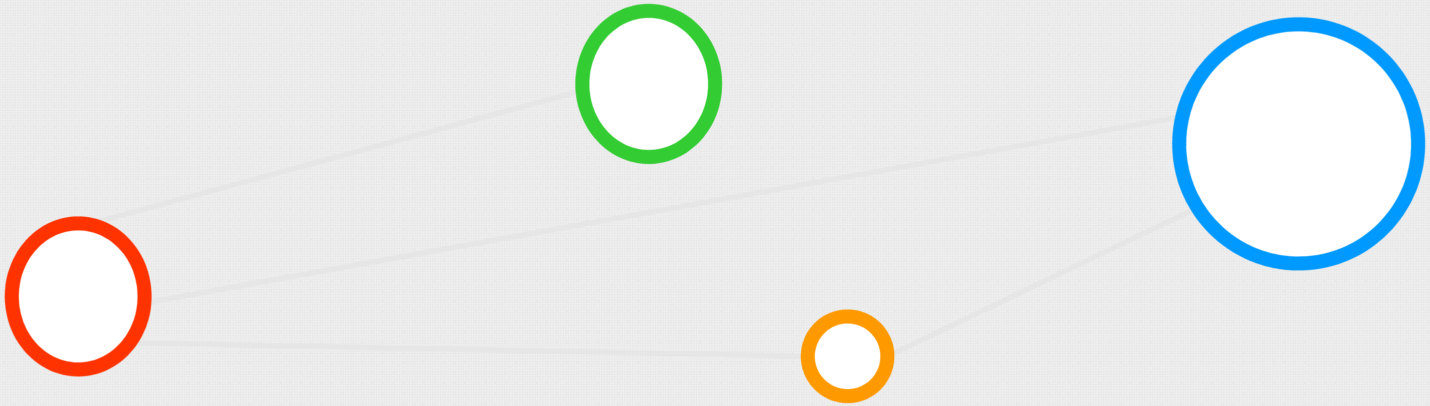

Minutes Approval – Monthly Meeting Feb 2021

1. Minutes available in BSCW
   - https://bscw.zam.kfa-juelich.de/bscw/bscw.cgi/d3349148/2021-02-11-Monthly-Meeting-Feb-2021-Minutes-v1.docx
   - TBD(all): Any objections/additions?
   - Only change since last time:
     Kristel Michielsen (FZJ) added apologies for not being able to attend
   - Action Items added to Action item tracker (see agenda item 2)
   - TBD(all): please read shortly minutes when available (lessons learned)

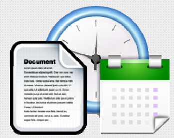

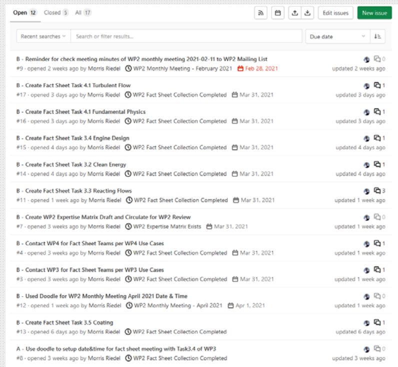

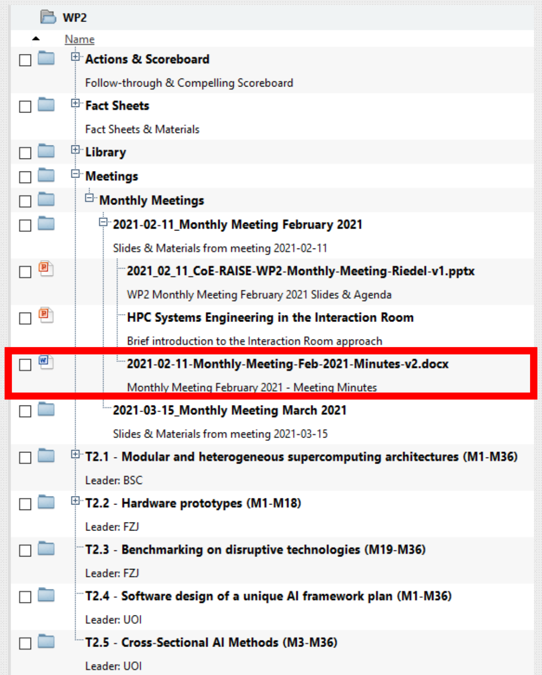

Agenda Item (2) – WP2 Updates & Fact Sheet Process

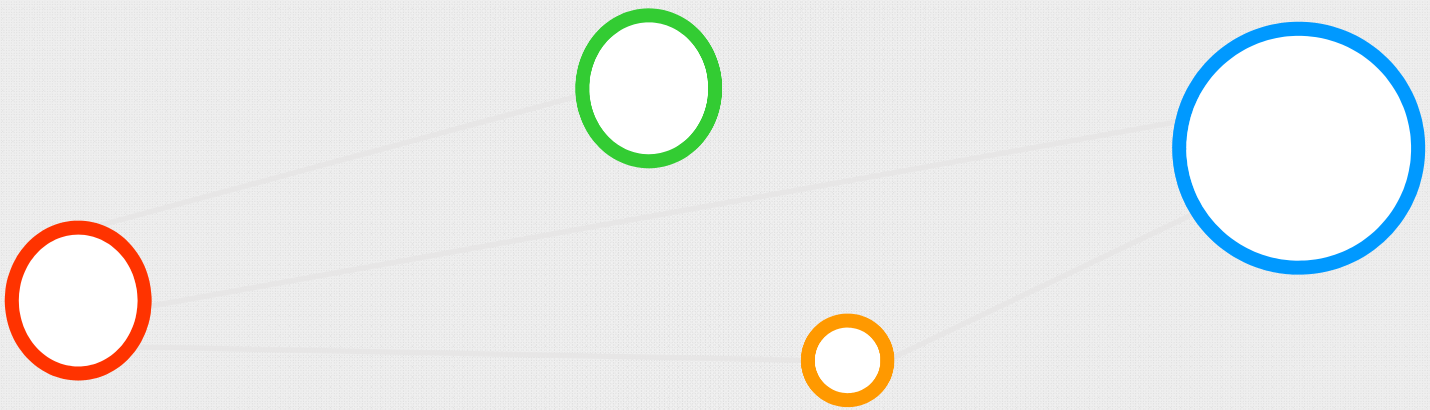

WP2 Updates – Action Items Tracker

- Tracking action items
  - Focus on WP2
  - Can be used by all
  - Milestones & Issues
  - Code (maybe later)
- Registration steps:
  1. Create JuDoor Account: https://judoor.fz-juelich.de/login
  2. Wait for Email to register
  3. Use my invitation link to get tracker access
  4. TBD(all): no access yet then please email morris@hi.is
- https://gitlab.version.fz-juelich.de/riedel1/raise-wp2/-/issues

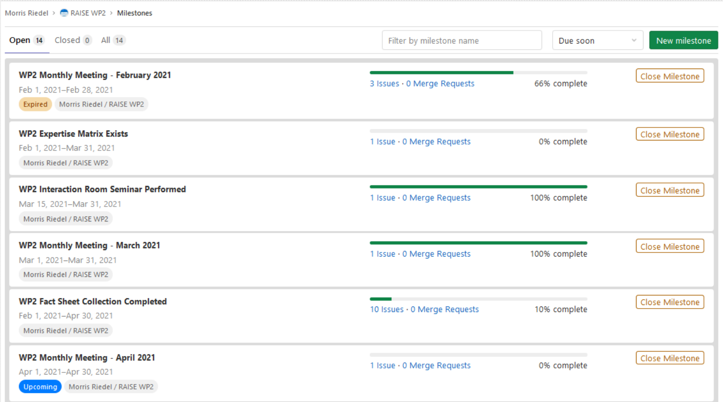

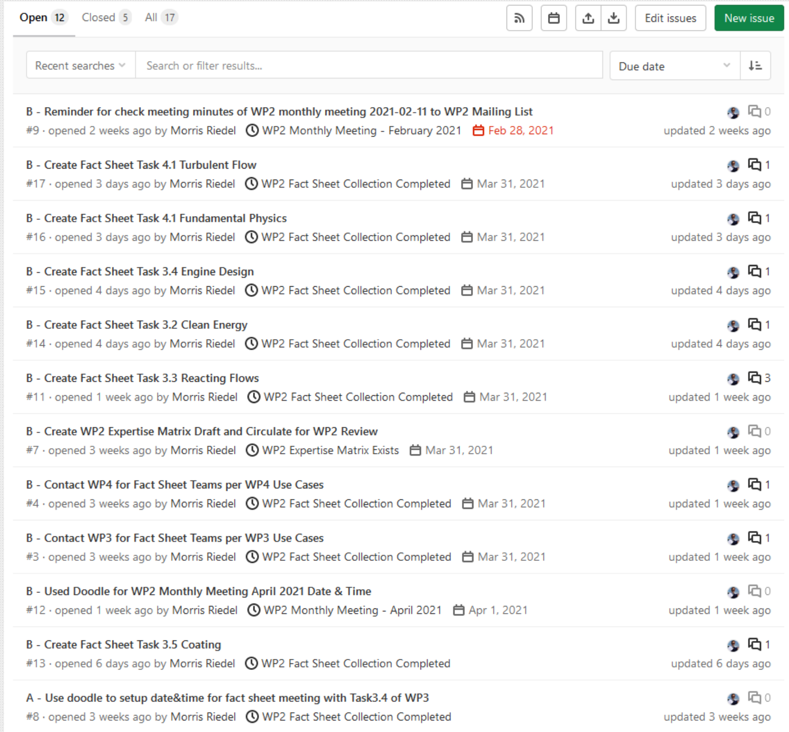

WP2 Updates – Library

- Collecting document relevant for WP2
  - Two major sections
  - https://bscw.zam.kfa-juelich.de/bscw/bscw.cgi/3373074
- RAISE Papers
  - Publications from us in RAISE!
  - Project outcome descriptions
  - Direct relevance for RAISE, e.g., RAISE acknowledgements in paper
- Related Work
  - Papers & Technical Reports
  - Policy documents
  - …

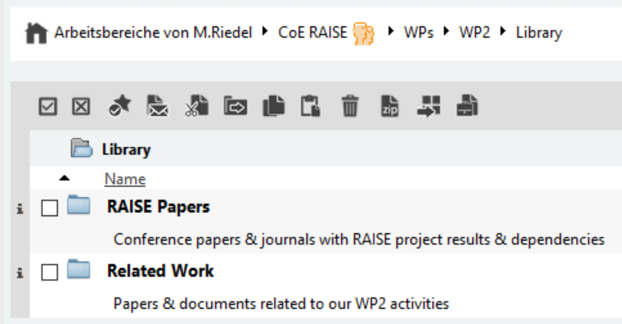

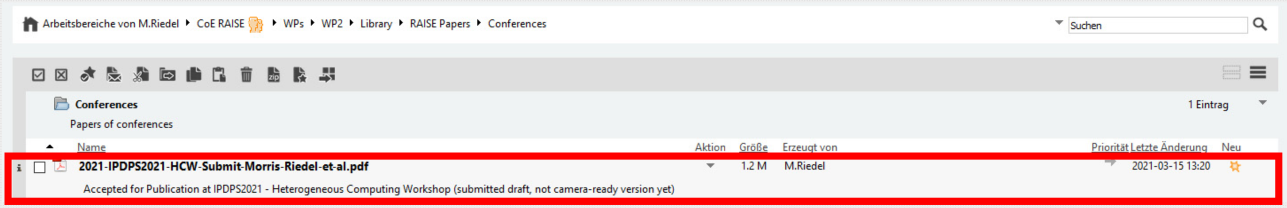

WP2 Process for our Discussion (Revisited from Feb Meeting)

- Fact Sheets
  - Foster initial understanding
  - Living document & each Fact Sheet per WP3/WP4 Use Case
  - *(Experience from many other EU projects)*

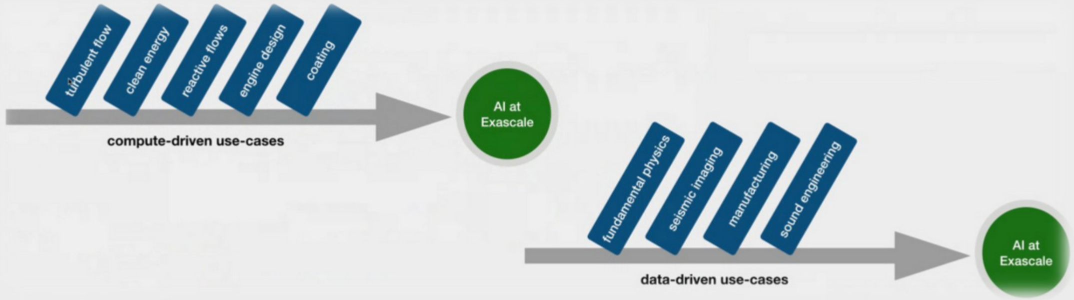

- Selected Contents
  - Short Application Introduction
  - Clarify Primary Contacts
  - Codes/Libraries/Executables
  - HPC System Usage Details
  - Specific Platforms & 'where is what data'?
  - Machine/Deep Learning Approaches of Interest (WP2 members in turn will then work on this)

initial steps

input to Interaction Room Process

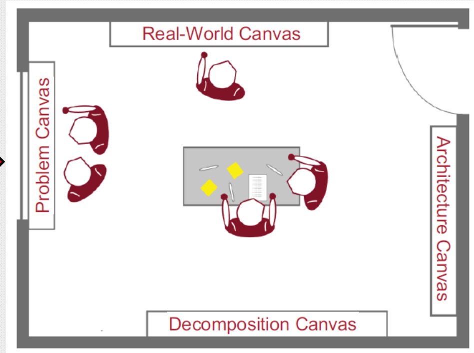

interaction room process

WP2 Fact Sheet Process of RAISE & WP3/WP4 – Idea

- Fact Sheets
  - Foster initial understanding → initial set should be completed at end of March 2021
  - Living document & each Fact Sheet per WP3/WP4 Use Case
  - *(Experience from many other EU projects)*
  - Input to Interaction Room process
  - Process started with initial use cases of WP3 & WP4
  - Potential updates throughout our project
  - Gradual improvement of fact sheet over time
  - For us as WP2 in RAISE: 'Where is what software running on which HPC system using what exact libraries and for what purpose …'

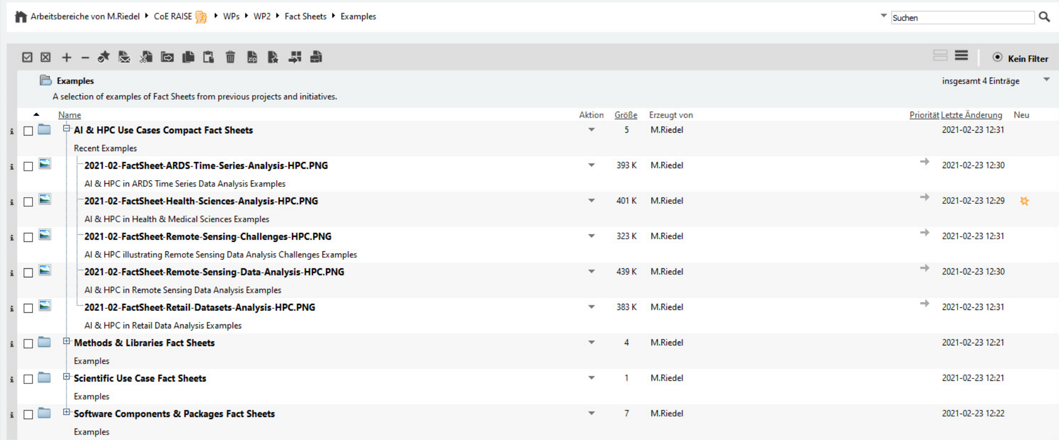

WP2 Fact Sheet Process of RAISE & WP3/WP4 – Examples

RAISE Center of Excellence

➤ Examples in BSCW: https://bscw.zam.kfa-juelich.de/bscw/bscw.cgi/3350073

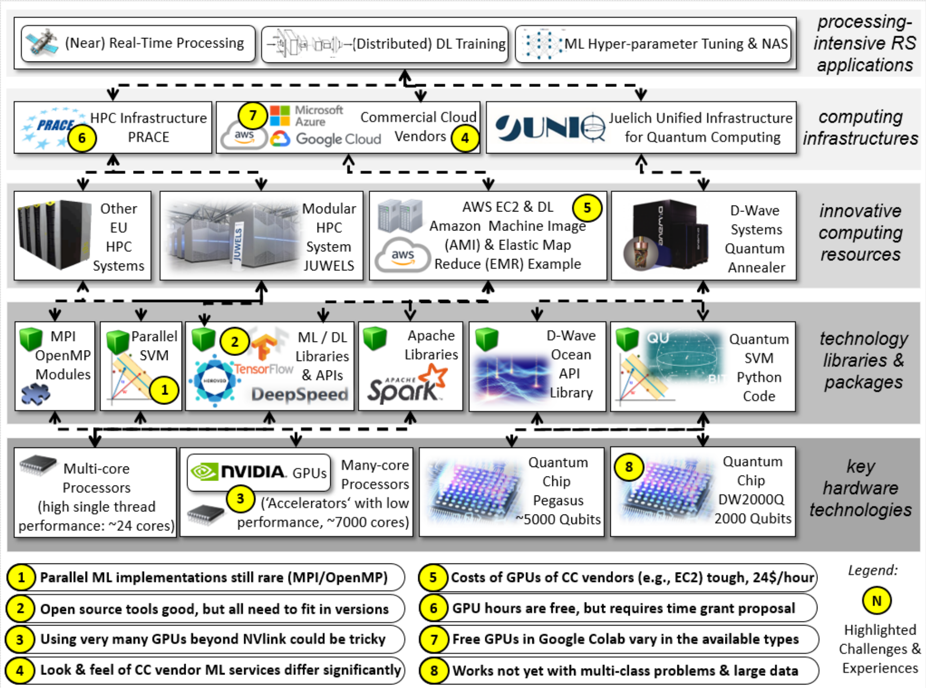

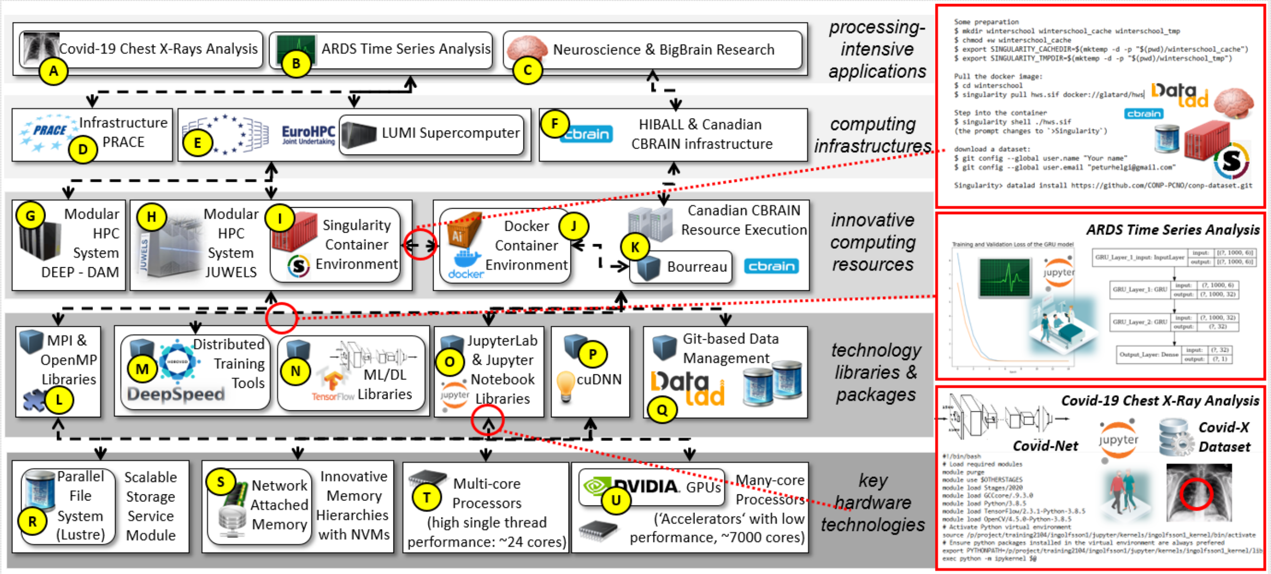

WP2 Fact Sheet Process of RAISE – Re-use Paper Example

RAISE Center of Excellence

- 'not a waste of time': Fact Sheets can be re-used for publications, project presentations & dissemination

Accepted for publication in IPDPS 2021, HCW workshop

Available in WP2 Library

PRACTICE AND EXPERIENCE IN USING PARALLEL AND SCALABLE MACHINE LEARNING WITH HETEROGENOUS MODULAR SUPERCOMPUTING ARCHITECTURES

Morris Riedel[1,2], Rocco Sedona[1,2], Chadi Barakat[1,2], Petur Einarsson[1,2], Gabriele Cavallaro[2]

[1] School of Engineering and Natural Sciences, University of Iceland, Iceland
[2] Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

ABSTRACT

We observe a continuously increased use of Deep Learning (DL) as a specific type of Machine Learning (ML) for data-intensive problems (i.e., 'big data') that requires powerful computing resources with equally increasing performance. Consequently, innovative heterogeneous High-Performance Computing (HPC) systems based on multi-core CPUs and many-core GPUs require an architectural design that addresses end user communities' requirements that take advantage of ML and DL. Still the workloads of end user communities of the simulation sciences (e.g., using numerical methods based on known physical laws) needs to be equally supported in those architectures. This paper offers insights into the Modular Supercomputer Architecture (MSA) developed in the Dynamic Exascale Entry Platform (DEEP) series of projects to address the requirements of both simulation sciences and data-intensive sciences such as High Performance Data Analytics (HPDA). It shares insights into implementing the MSA in the Jülich Supercomputing Centre (JSC) hosting Europe No. 1 Supercomputer Jülich Wizard for European Leadership Science (JUWELS). We augment the technical findings with experience and lessons learned from two application communities case studies (i.e., remote sensing and health sciences) using the MSA with JUWELS and the DEEP systems in practice. Thus, the paper provides details into specific MSA design elements that enable significant performance improvements of ML and DL algorithms. While this paper focuses on MSA-based HPC systems and application experience, we are not losing sight of the bigger picture, including Cloud Computing (CC) and Quantum Computing (QC) advances relevant for ML and DL.

Index Terms— High performance computing, cloud computing, quantum computing, machine learning, deep learning, parallel and distributed algorithms, remote sensing, health sciences, modular supercomputer architecture.

This work was performed in the Center of Excellence (CoE) Research on AI- and Simulation-Based Engineering at Exascale (RAISE), the Euro CC, and DEEP-EST projects receiving funding from EU's Horizon 2020 Research and Innovation Framework Programme under the grant agreement no. 951733, no. 951740 and no. 754304 respectively.

1. INTRODUCTION

Today, an academically-driven supercomputing centre's (e.g., Juelich Supercomputing Centre[1], Barcelona Supercomputing Centre[2], or Finish IT Center for Science CSC[3]) application portfolio is highly multidisciplinary, raising diverse requirements for a HPC architecture that enables research for a wide variety of end-user communities [1]. Examples include but are not limited to astrophysics, computational biology and biophysics, chemistry, earth and environment, plasma physics, computational soft matter, fluid dynamics, elementary particle physics, computer science and numerical mathematics, condensed matter, and materials science. Not only the research approaches in these communities are diverse, but also the way how they employ scalable algorithms, numerical methods, and parallelisation strategies. Many of these are *'traditional HPC applications'* (i.e., modeling and simulation sciences) that use iterative methods and rely heavily on a small number of numerical algorithmic classes that operate on relatively small to moderate-sized data sets and accrue very high numbers of floating-point operations across iterations. But we observe that the complexity (e.g., using CPU in conjunction with GPUs) and memory requirements (e.g., using complex memory hierarchies) of HPC codes of those applications increases, leading to a dissonance with these traditional HPC system workloads.

More recently, new user communities add to the above mentioned diversity in the sense of using the HPC systems with HPDA using ML and DL in conjunction with containers [2] and interactive supercomputing (e.g., via Jupyter[4] notebooks) [3]. Those workloads (e.g., remote sensing or health sciences) are rapidly emerging and require a change in HPC systems architecture. They exhibit less arithmetic intensity and instead require additional classes of parallel and scalable algorithms to work well (e.g., DL networks with interconnections of GPUs to scale to extreme scale). Some end-user communities (e.g., neurosciences and earth sciences) make intertwined use of both traditional simulation sciences-based HPC

[1] https://www.fz-juelich.de/ias/jsc/EN/Home/home_node.html
[2] https://www.bsc.es/
[3] https://www.csc.fi/en/csc
[4] https://jupyter.org/

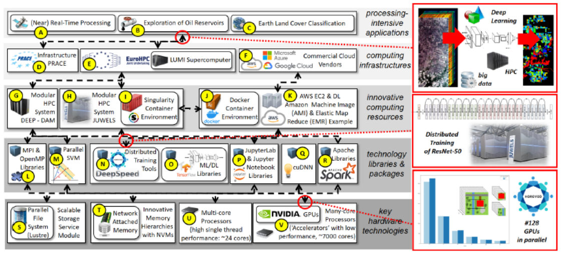

Fig. 2. Remote Sensing applications taking advantage of the MSA ensuring conceptual interoperability with Clouds.

D) is a necessary solution to train DL classifiers in a reasonable amount of time, providing RS researchers with a high-accuracy performance in the application recognition tasks. The same is true for the emerging HPC system landscape currently acquired by the EuroHPC Joint Undertaking such as the LUMI supercomputer in Finland (see Fig. 2 E).

Our RS case study mainly takes advantage of the MSA-based JUWELS system (see Fig. 2) at the JSC in Germany, representing the fastest EU supercomputer with 122,768 CPU cores only in its cluster module (cf. Section 2.1 H). While JUWELS and multi-core processors (see Fig. 2 U) offer tremendous performance, the particular challenge to exploit this data analysis performance for ML is that those systems require specific parallel and scalable techniques. In other words, using JUWELS cluster module CPUs with Remote Sensing (RS) data effectively requires parallel algorithm implementations opposed to using plain scikit-learn[15], R[16], or different serial algorithms. Parallel ML algorithms are typically programmed using the MPI standard, and OpenMP (see Fig. 2 L) that jointly leverage the power of shared memory and distributed memory via low latency interconnects (e.g., Infiniband[17]) and parallel filesystems (e.g., Lustre[18]).

Given our experience, the availability of open-source parallel and scalable machine learning implementations for the JUWELS cluster module CPUs that go beyond Artificial Neural Network (ANN)s or more recent DL networks (see Fig. 2 O) is still relatively rare. The reason is the complexity of parallel programming of ML and DL codes and thus using HPC with CPUs only can be a challenge when the amount of data

is relatively moderate (i.e., DL not always succesful). One example is using a more robust classifier such as a parallel and scalable Support Vector Machine (SVM) open-source package (see Fig. 2 M) that we developed with MPI for CPUs and used to speed up the classification of RS images [16].

3.1. Selected DL Experiences on MSA-based Systems

The many-core processor approach of the highly scalable JUWELS booster (see Section 2.2) with accelerators brings many advancements to both simulation sciences and data sciences, including innovative DL techniques. Using many numerous simpler processors with hundreds to thousands of independent processor cores enabled a high degree of parallel processing that fits very nicely to the demands of DL training whereby lots of matrix-matrix multiplications are performed. Today, hundreds to thousands of accelerators like Nvidia GPUs (see Fig. 2 V) are used in large-scale HPC systems, offering unprecedented processing power for RS data analysis. JUWELS Booster module offers 3744 GPUs of the most recent innovative type of Nvidia A100 tensor core[19] cards. Our experience on MSA-based systems such as DEEP[20] (see Fig. 2 G), JURECA [21], and JUWELS shows that open-source DL packages such as TensorFlow[22] (now including Keras[23]) or pyTorch[24] are powerful tools for large-scale RS data analysis.

We experienced that it can be quite challenging to have the right versions of python code matching the available DL

[15] https://scikit-learn.org/stable/
[16] https://www.r-project.org/
[17] https://www.mellanox.com/products/interconnect/infiniband-overview
[18] https://www.lustre.org/
[19] https://www.nvidia.com/en-us/data-center/a100/
[20] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/DEEP-EST/_node.html
[21] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JURECA/JURECA_node.html
[22] https://www.tensorflow.org/
[23] https://keras.io/
[24] https://pytorch.org/

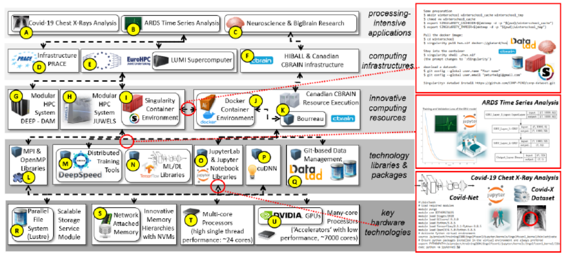

Fig. 3. Health science applications taking advantage of the MSA enabling seamless access for non-technical medical experts.

latest cuDNN support (see Fig. 3 P) the inference and training time of the Covid-Net model is significantly faster as with GPUs of the previous generation given its tensor cores.

We used several publicly available datasets of COVIDx [25] that is an open access benchmark dataset initially comprising of 13,975 CXR images across 13,870 patient patient cases. But in the last couple of month this dataset was extended numerous times with new datasets made available that in turn we used again with Covid-Net as well. The SSSM of the MSA systems and its parallel file system Lustre (see Fig. 3 R) provides a powerful storage mechanism to store the COVIDx datasets and its updates.

This module also stores additional data we obtained from a collaborating Pharma company that we in turn used to validate that Covid-Net is able to generalize well to unseen datasets. At the time of writing, the name and dataset of the collaborating pharma company can not be revealed, but will be made available to the workshop organizers during the workshop. Using the MSA-based systems JUWELS and DEEP seamlessly with Jupyter requires the definition of an own Kernel[47] using the module[48] environment of the MSA HPC systems (see Fig. 3 bottom right). Our experience on using own Kernels with Jupyter notebooks is extremely positive while at the same time offering a user interface with notebooks that are user-friendly enough for medical imaging experts.

4.2. Time Series Data Analysis of ARDS Patients

Our application case study *'ARDS Time Series Analysis'* (see Fig. 3 A) addresses the medical condition Acute Respiratory Distress Syndrome (ARDS) that affects on average 1-2% of mechanically-ventilated (MV) Intensive Care Unit (ICU) patients and has a 40% mortality rate [26, 27]. At present the leading protocol for diagnosing the condition is the Berlin definition that defines onset of ARDS as a prolonged ratio of arterial oxygen potential to fraction of inspired oxygen (P/F ratio) of less than 300 mmHg, and the lower this value is determined to be, the more severe the diagnosis is [28]. Several papers have determined a correlation between early detection of onset of ARDS and survival of the patient, which highlights the need of early detection and treatment of the condition, before onset of sepsis and subsequently multi-organ failure [27, 29, 30]. Hence, the goal of this case study is to develop an algorithmic approach that provides early warning and informs medical staff of mitigating procedures can be a beneficial tool for ICU personnel.

We take advantage of the freely available ICU patient data provided in the Medical Information Mart for Intensive Care - III (MIMIC-III) database, compiled between 2001 and 2012 from admissions to the Beth Israel Deaconess Medical Center in Boston, MA [31]. The procedure thus, is to build and test our models using patient data from the MIMIC-III database, then verify our results using patient data collected from hospital participating in our German Smart Medical Information Technology for Healthcare (SMITH) project consortium[49] with real hospitals, and finally roll out the developed model

[47] https://jupyter-jsc.fz-juelich.de/nbviewer/github/FZJ-JSC/jupyter-jsc-notebooks/blob/master/001-Jupyter/Create_JupyterKernel_general.ipynb
[48] https://hpc-wiki.info/hpc/Modules
[49] https://www.smith.care/home-2/

Agenda Item (3) – Fact Sheet Status & Discussions

RAISE
Center of Excellence

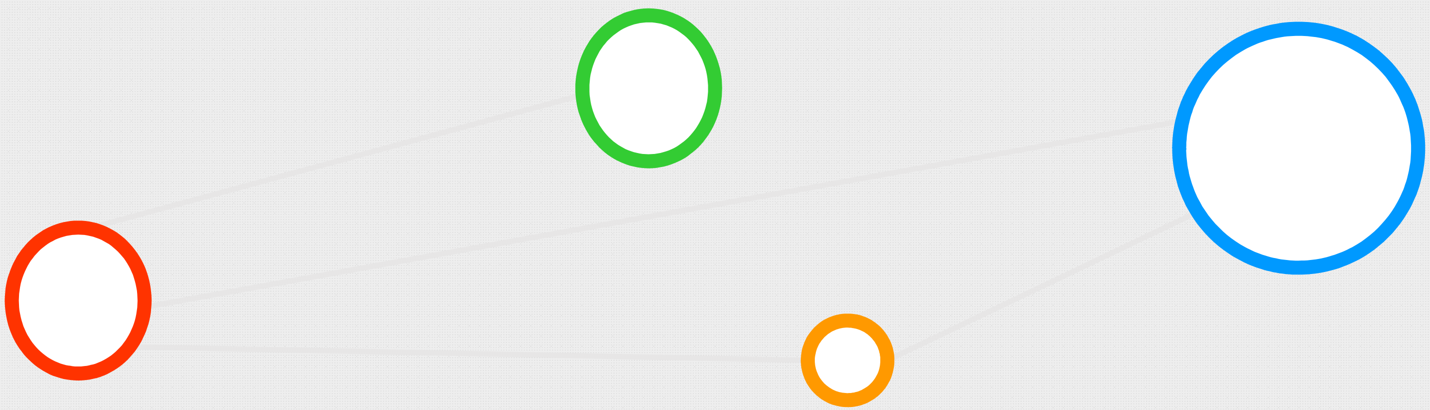

Fact Sheet Status & Discussions – WP3/WP4 Tasks Overview

- WP3
  - T3.1: Turbulent Flow (planned)
  - T3.2: Clean Energy (started)
  - T3.3: Reactive Flows (started)
  - T3.4: Engine design (partly covered by T3.4)
  - T3.5: Coating (started)
- WP4
  - T4.1: Fundamental physics (planning)
  - T4.2: Seismic imaging (planning)
  - T4.3: Manufacturing (planning)
  - T4.4: Sound engineering (started)

Table 6: Use-case vs. AI-methods matrix.

Use-Case vs. AI-Methods	DA	NAS	AE	TL	PF	PIDL	LSTM
Turbulent boundary layers	x	x	x	x	x	x	
Wind farm layout optimization	x			x		x	
AI for data-driven models in reacting flows				x		x	
Smart models for next-generation aircraft engine design	x	x		x		x	
Wetting hydrodynamics		x	x			x	x
Event reconstruction and classification at the CERN HL-LHC		x		x			x
Seismic imaging with remote sensing - oil and gas exploration and well maintenance	x	x		x			
Defect-free metal additive manufacturing		x				x	x
Sound engineering	x	x		x			x

- Lessons Learned
  - Hard to find time
  - But doable, sometimes also w/o 'big doodle'
  - So far, all found Fact Sheet process useful
  - One more iteration at least necessary:
    Dive more into potential AI topics

Fact Sheet Status & Discussions – T3.2

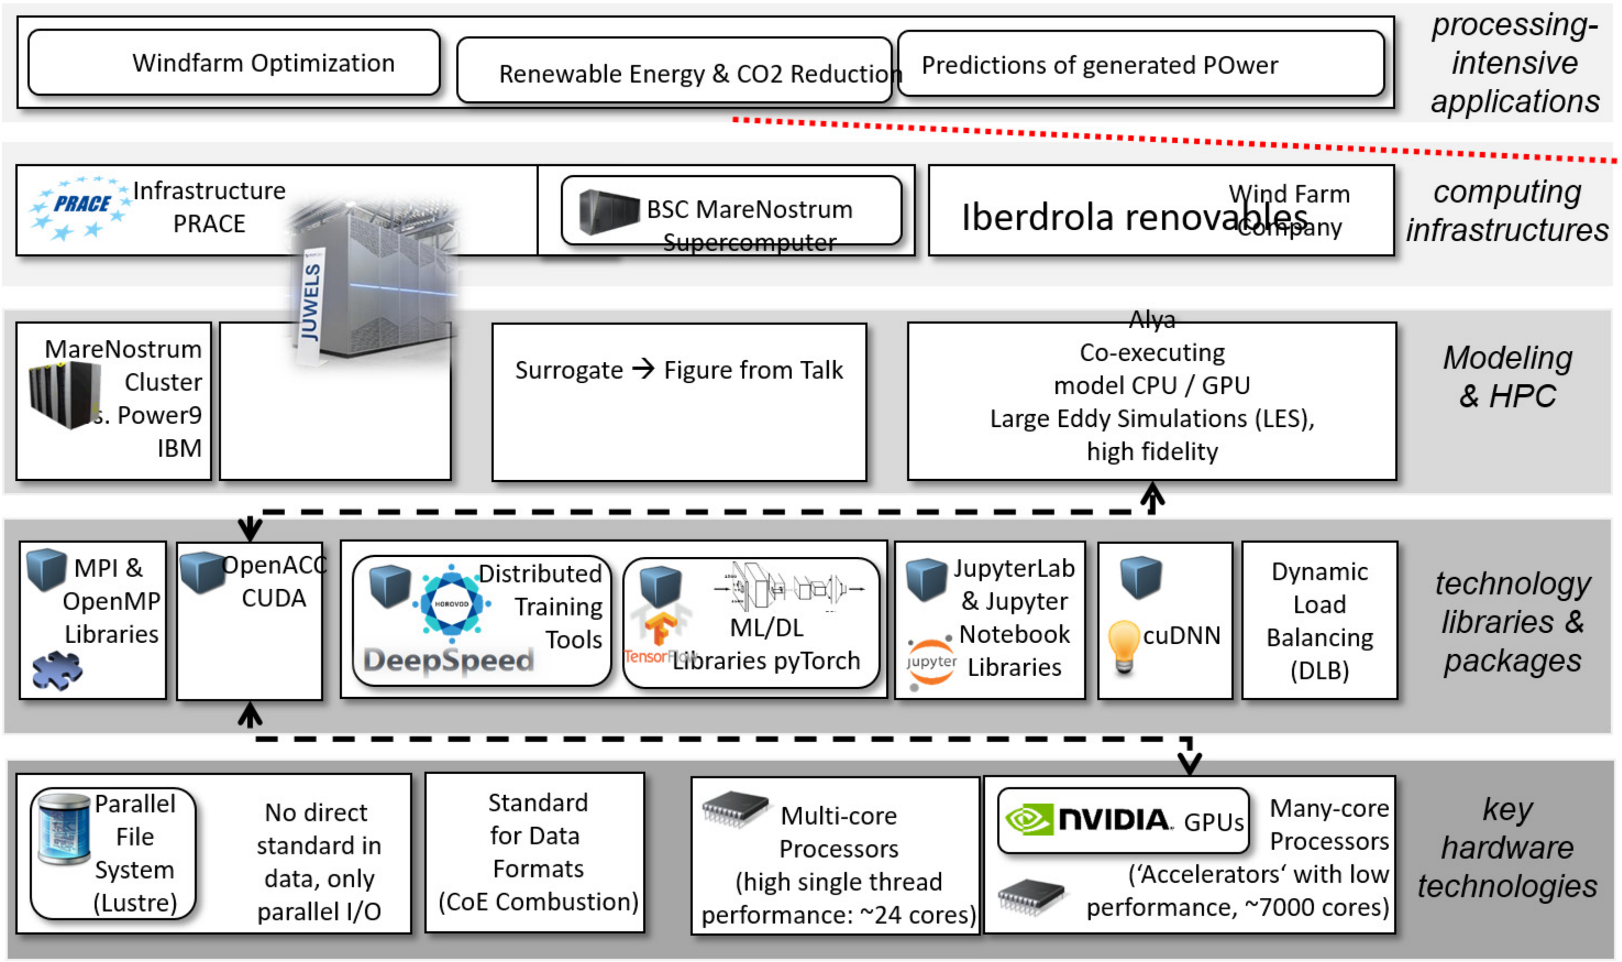

WORK IN PROGRESS

Fact Sheet Status & Discussions – T3.3 (& T3.4)

RAISE Center of Excellence

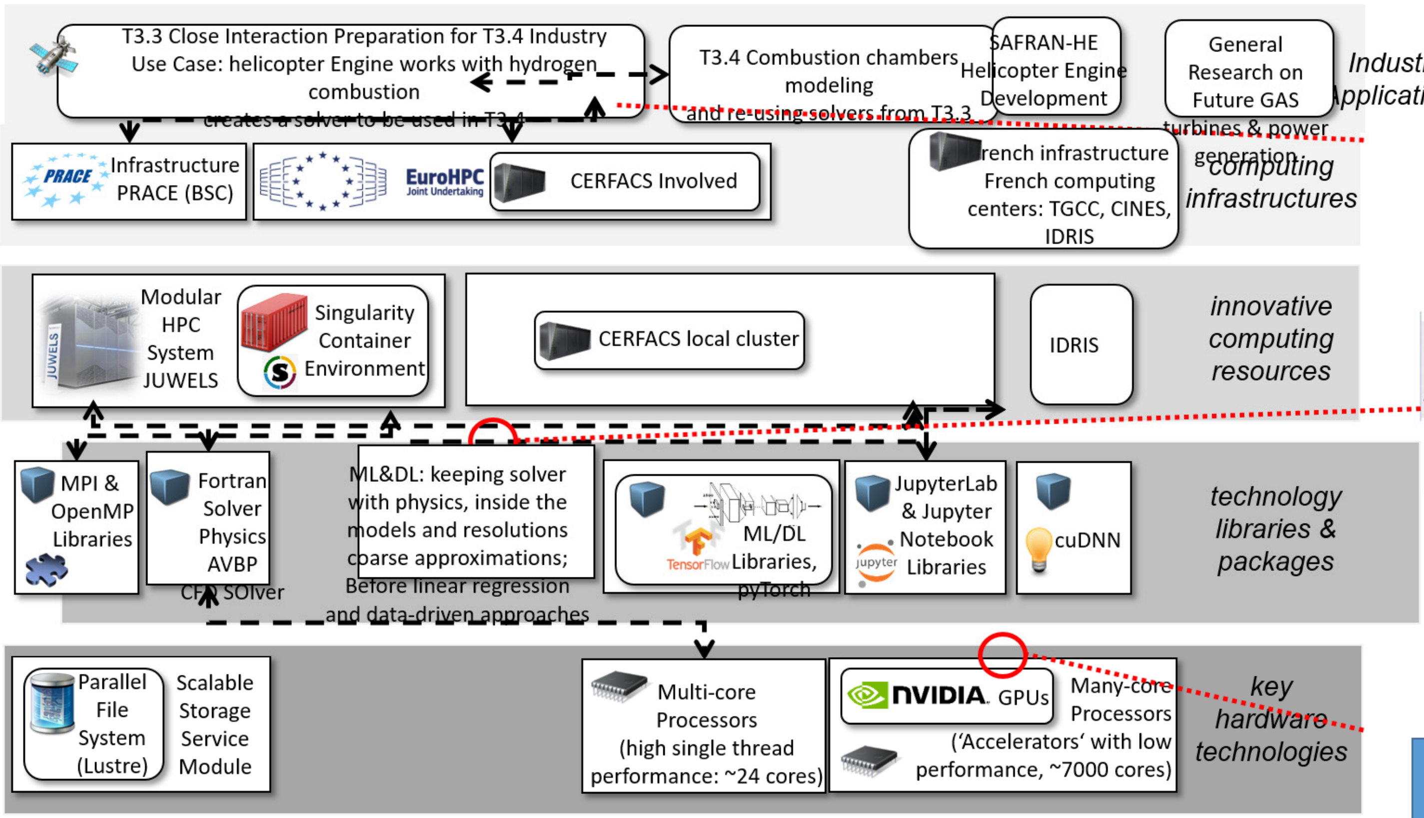

SAFRAN-CERFACS Interactions: helicopter might use hydrogen quicker, helicopter engine more faster designed, airbus not really

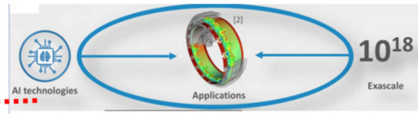

FACTS

WORK IN PROGRESS

SAFRAN-CERFACS Interactions: helicopter might use hydrogen quicker, helicopter engine more faster designed, airbus not really

GPU ai and CPU solver

Fact Sheet Status & Discussions – T3.5

RAISE Center of Excellence

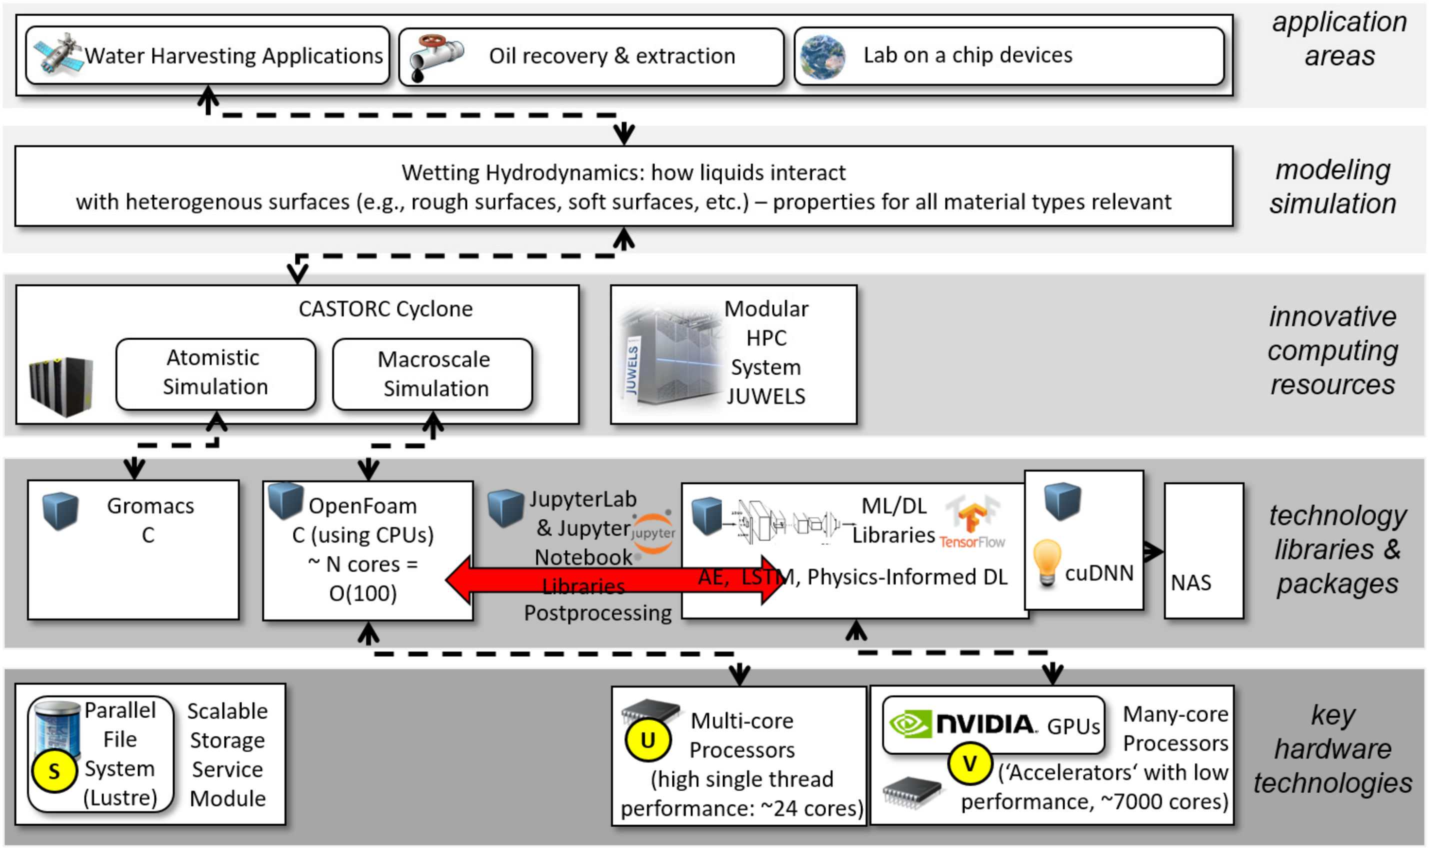

Lab on a chip devices: experimentation in biology lab: these devices change how the droplet moves, batch processing of bio processes, droplet, biology experiments

No Cloud

WORK IN PROGRESS

No containers

S

Fact Sheet Status & Discussions – T4.4

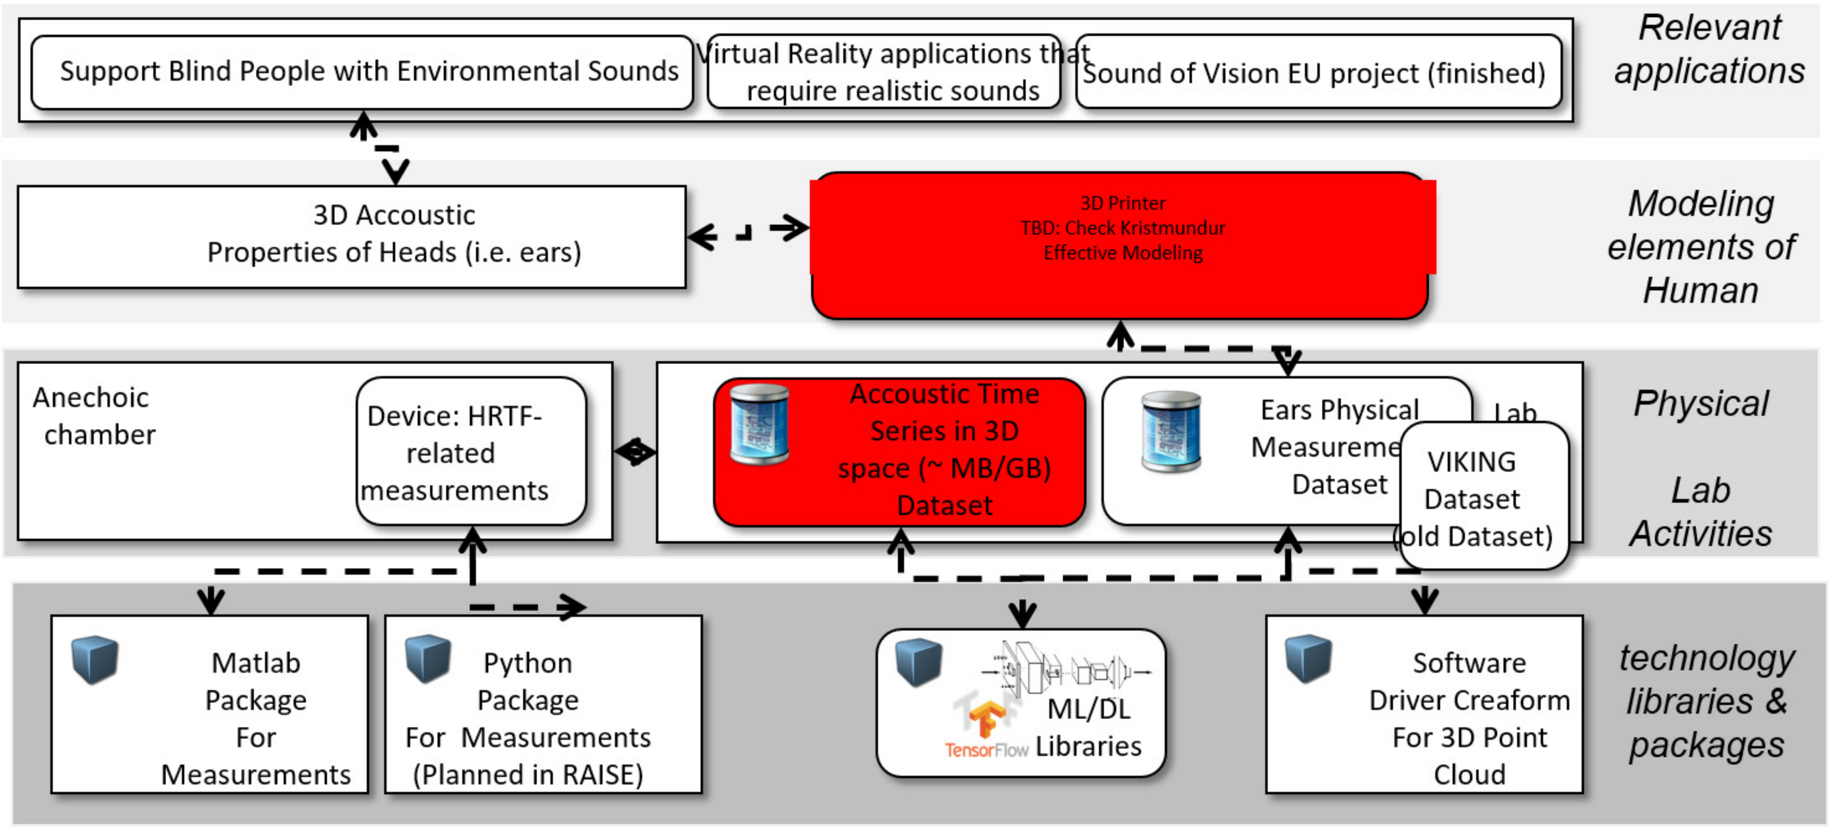

WORK IN PROGRESS

Agenda Item (4) – Follow-up Resource Provisioning

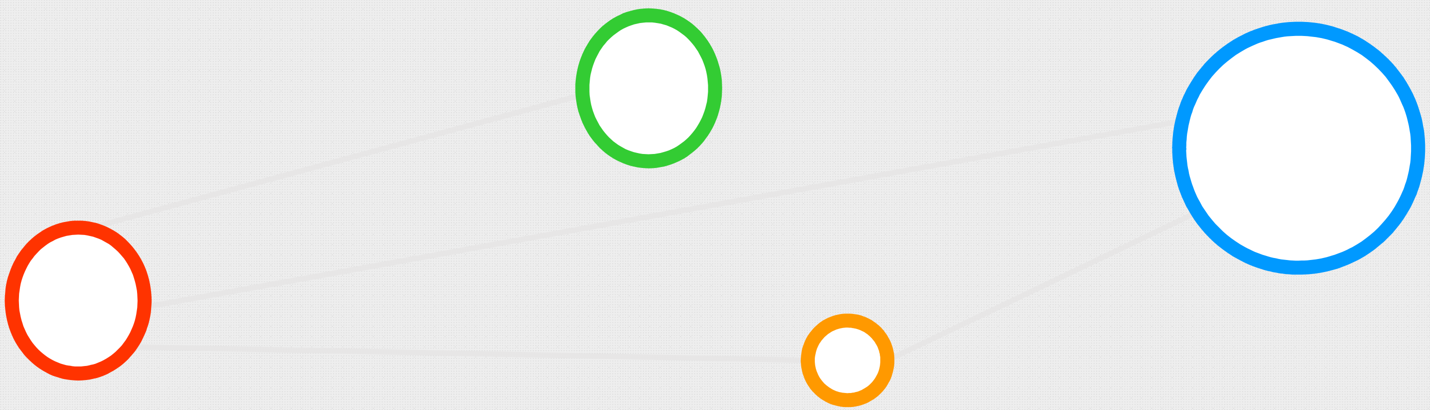

Follow-up Resource Provisioning (Guillaume)

- Follow-up Resource Provisioning
  - Slides (Guillaume Houzeaux), ~15 Min

Agenda Item (5) – Follow-Through Check: Meeting duration

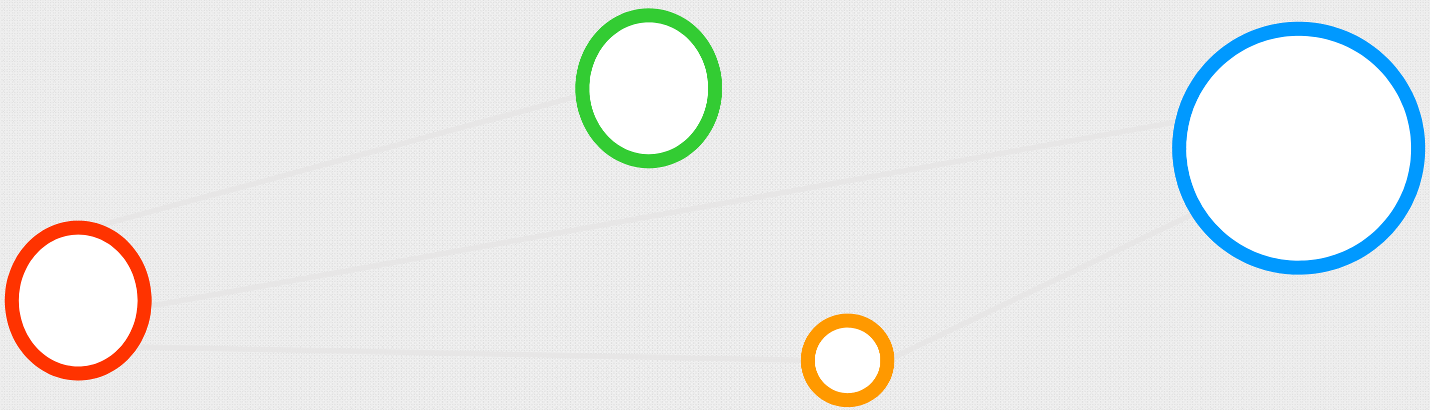

Follow-Through Check: Meeting duration

- Follow-Through Check:
  Meeting duration (1 hour monthly long enough?)
  - (All), ~5 Min
  - Feedback: no improvement suggestion, all seem to be fine with it now
  - Outcome: April/May will be the same way structured
  - In addition, monthly seminars on specific topics:
    Interaction room (April), DataLad (May), AI model XYZ depending no WP3/WP4 needs, etc.

Agenda Item (6) – Next Monthly Meeting April & AOB

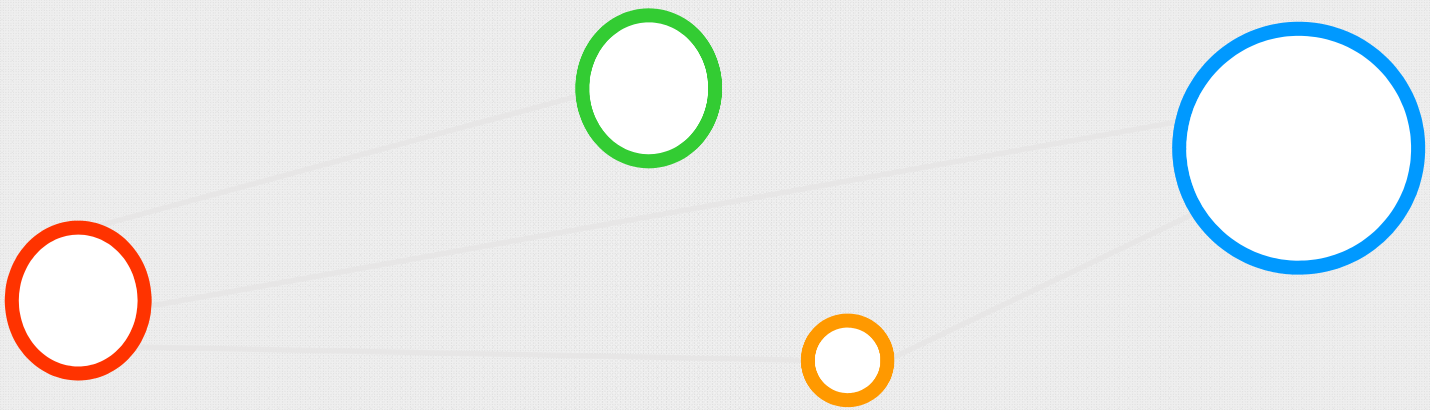

REMINDER: INTERACTION ROOM SEMINAR ON 2021-04-08

RAISE Center of Excellence

- Start: Interaction Room Process 10:00 – 12:00 GMT!
  - Supports the proper software engineering design of the unique AI framework blueprint
  - Expecting to work with WP3 & WP4 experts in an open minded way
  - Process will be guided by **Prof. Dr. Matthias Book** (Software Engineering, University of Iceland)
  - Supported by Software Engineering & testing expert **Prof. Dr. Helmut Neukirchen** (University of Iceland)

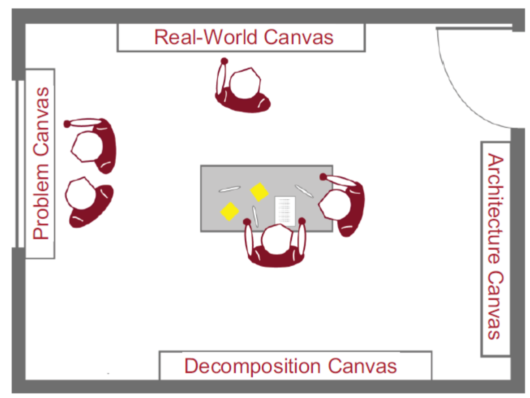

Book, M., **Riedel, M.**, **Neukirchen, H.**, Goetz, M.: Facilitating Collaboration in High-Performance Computing Projects with an Interaction Room, in conference proceedings of the 4th ACM SIGPLAN International Workshop on Software Engineering for Parallel Systems (SEPS 2017), October 22-27, 2017, Vancouver, Canada

Compelling Scoreboard: Summary & Next Month

✓ Call for Monthly Meeting February & check missing people in WP2 (Kick-off)

✓ Monthly Meeting February

✓ Final D1.5 (M2) & D1.2 (M2)

✓ Monthly Meeting March

✓ Interaction Room Seminar scheduled in April 2021-04-08

✓ Planing of DataLad Seminar

01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36

- D2.12 Software layout plan for a unique AI framework (M10)

❖ Fact Sheet Process & Interaction Room (started): revision and refinement process

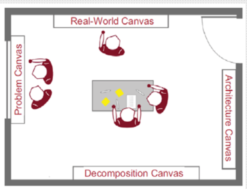

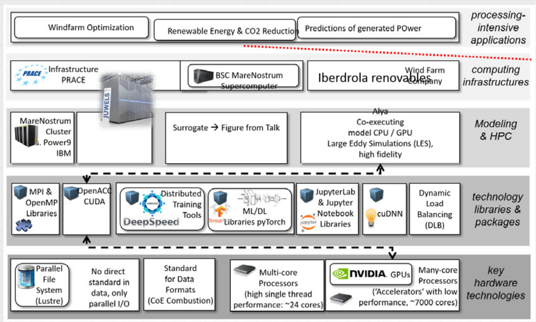

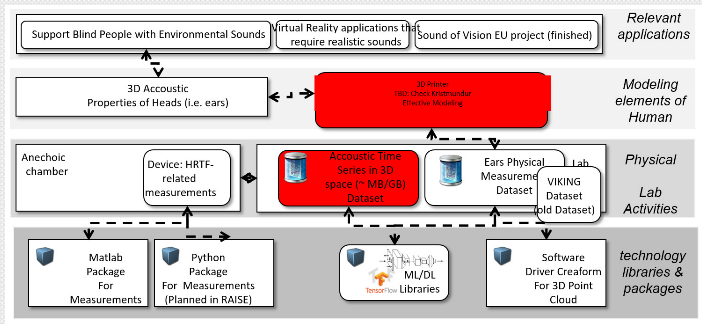

➢ Next Steps:

- ➢ Finish Fact Sheets Initial Versions
- ➢ Prepare Interaction Room Seminars
- ➢ Prepare May Seminar on DataLad

Agenda Item (6) – WP2 Next Monthly Meeting & AOB

RAISE Center of Excellence

1. Next Monthly Meeting(s)
   - April, May: Doodle will be available soon
   - Topics: Selected progress with Factsheets & Interaction Room presentations

- TBD(All): Discussions & AOB?
  - Good feedback from Project Director Andreas Lintermann for WP2 progress, process & materials
  - TBD(Morris): Re-Invite all not yet members of WP2 Action Item Tracker

RAISE
Center of Excellence

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