University of Basel

DMI HPC

A Glimpse Into the **High Performance Computing** Group's **Research Activities** at DMI

hpc.dmi.unibas.ch | dmi-hpc@unibas.ch | @HPC_DMI_UniBas | @hpc_dmi_unibas

since 08/2015

Updated: 19.12.2023

University of Basel > Faculty of Science > Department of Mathematics and Computer Science > High Performance Computing

In a Nutshell (as of December 2023)

Team and Talents

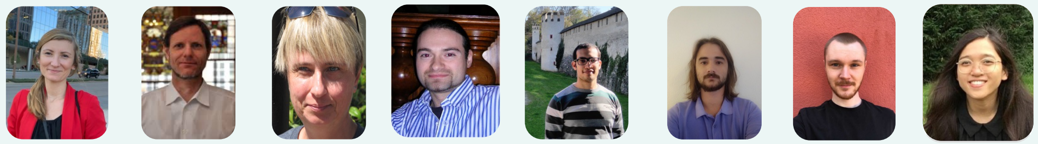

Florina Ciorba, Lead
Robert Frank, IT administration (shared)
Yvonne Walser, Secretary (shared)
Ruben Cabezon, Senior scientist (external, sciCORE)
Ahmed Eleliemy, Postdoc (alumnus)
Jonas Müller Korndörfer, Postdoc
Thomas Jakobsche, PhD student
Gabrielle Poerwawinata (alumnus)

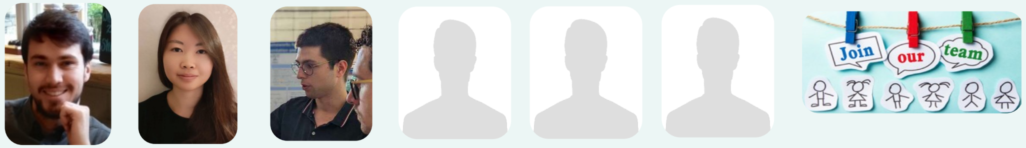

Osman Simsek, Postdoc
Yiqing Zhu, Scientific Software Developer
Quentin Guilloteau, Postdoc
PhD dissertations (3), BSc Theses (10+)
MSc Theses (10+), MSc Projects (12+)
Student workers (3)

PhD student Open Position

Lectures and Teaching

Operating Systems
BSc course, 8 CP
Spring 2022-present

High Performance Computing
MSc course, 4 CP
Spring 2016-present

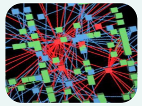

Foundations of Distributed Systems
MSc course, 8 CP
(with C. Tschudin, H. Schuldt)
Fall 2016-present

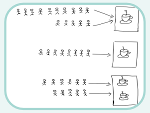

Algorithms, Concurrency, Parallelism (2022)
Privacy Enhancing Technologies and Performance (2023)
BSc seminar, 6 CP

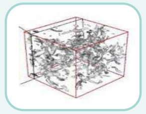

Applications of Computational Sciences
BSc course, 2 CP
(with others)
Fall 2019-present

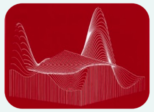

Advanced Methods in Computational Sciences
PhD course, 3 CP
(with others)
Spring 2020-present

Research and Projects

DAPHNE
Integrating Data Analytics Pipelines for Large-Scale DM, HPC, and ML
01.12.2020-30.11.2024

FNSNF
The DIALOGUE Study: Using digital health to improve care for families with predisposition to HBOC
01.11.2019-31.10.2023

PASC – Platform for Advanced Scientific Computing
SPH-EXA2: Smoothed Particle Hydrodynamics at Exascale
01.07.2021-30.06.2024

FNSNF
MLS2: Integrated Multilevel Scheduling for HPC
2025-2029
(in preparation)

FNSNF
MODA4HPC: Monitoring and Operational Data Analysis for Improving HPC Operations and Research
2025-2029
(in preparation)

Universität Basel
μ-Cluster: Visualizing Everything Parallel
01.03.2017-present

miniHPC: Research & teaching HPC cluster
01.01.2017-present

3BEARS: Broad Bundle of BEnchmARks for Scheduling in HPC, Big Data, and ML
01.01.2021-31.12.2022

SKACH
SKACH: Sky to Observations
01.09.2021-31.12.2024

SIREN: Software Identification & Recognition
31.03.2023 - 03.04.2024

High Performance Computing Group (DMI-HPC)

"How to exploit all parallelism efficiently across levels and devices?"

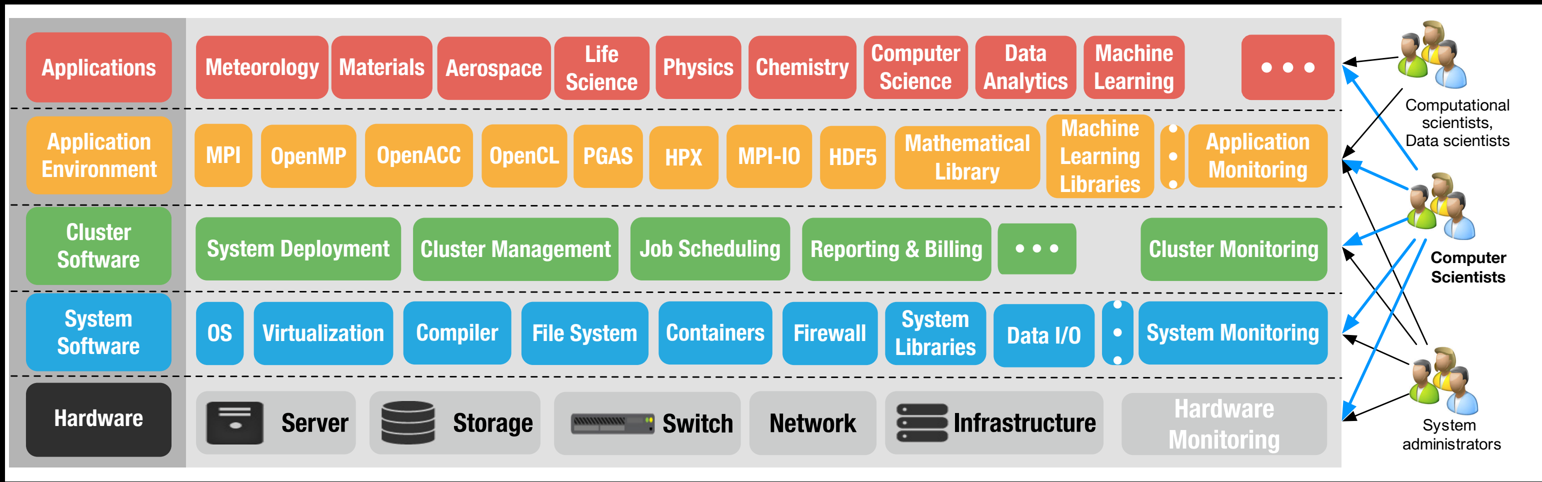

Questions & Goals

- Performance
- Resilience
- Security
- Sustainability
- Autonomous operation

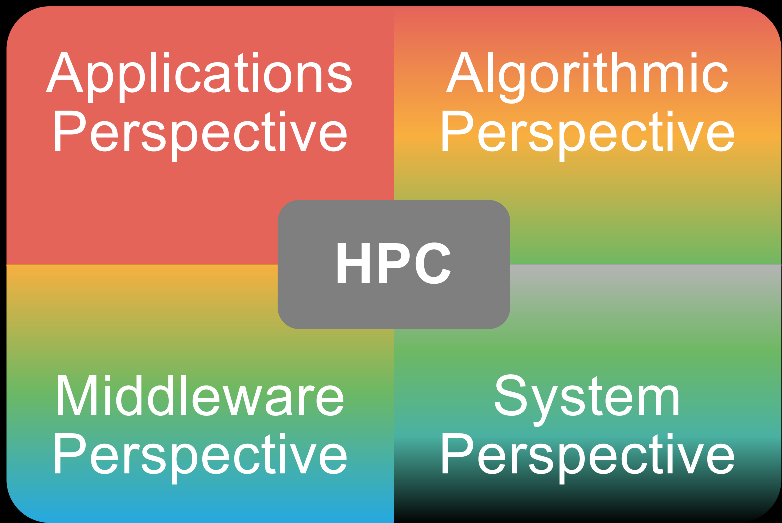

Challenges

- Multilevel parallelism
- Load imbalance
- Data locality
- Portability
- Holistic Monitoring
- Reproducibility

Fall 2023

Applications Perspective
Interdisciplinary Co-Design

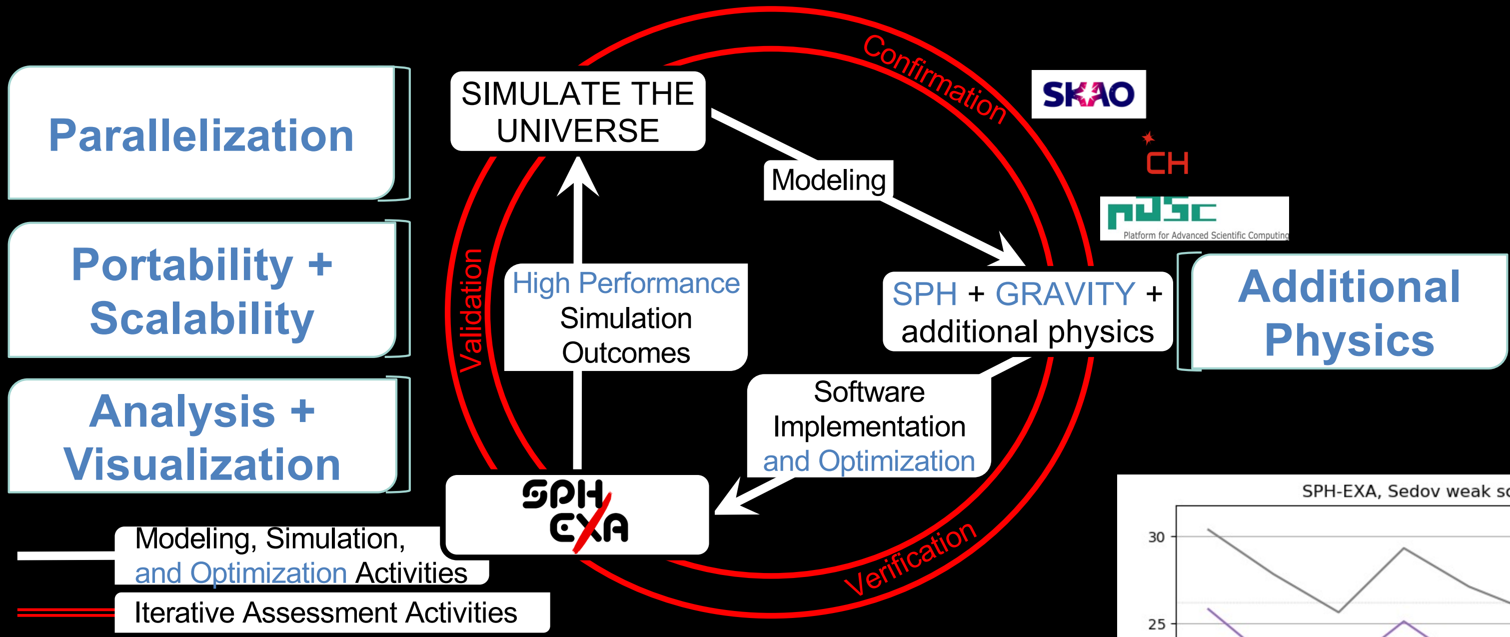

How to simulate and visualize the Universe efficiently & sustainably at most extreme resolutions & scales?

Interdisciplinary co-design

- Parallelization of all physics
- Portability + Scalability of each component
- Analysis and visualization of each simulation

Requirements

- State-of-the-art algorithms, libraries, system software
- Large-scale computing platform

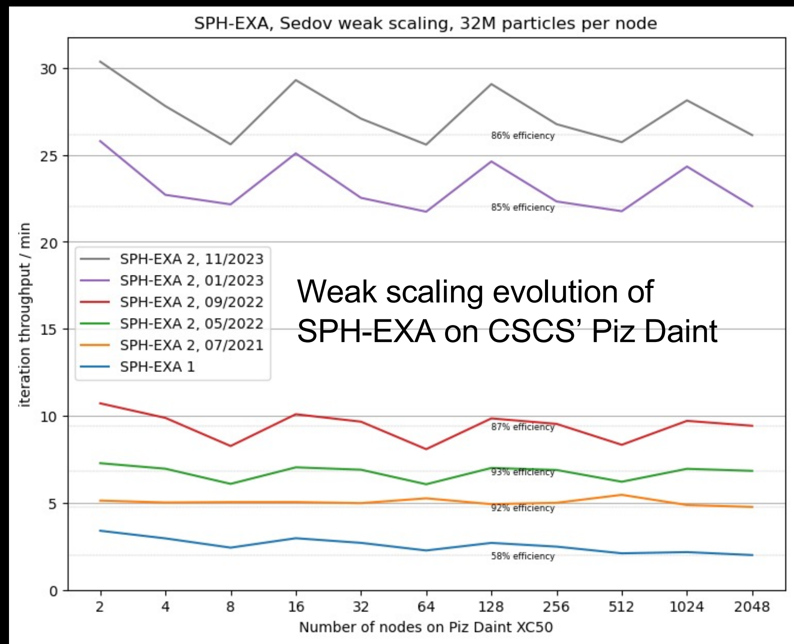

Algorithmic Perspective
Multilevel Scheduling

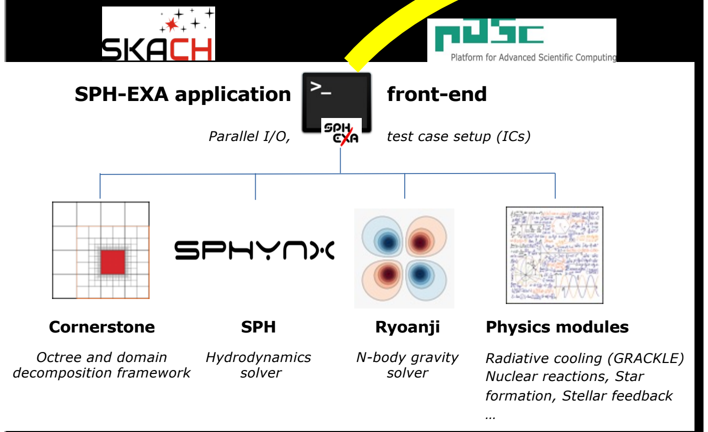

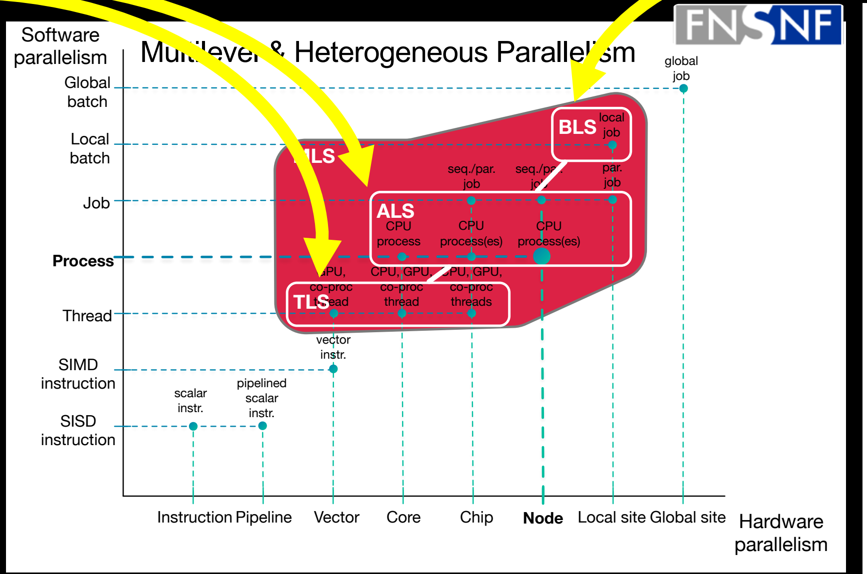

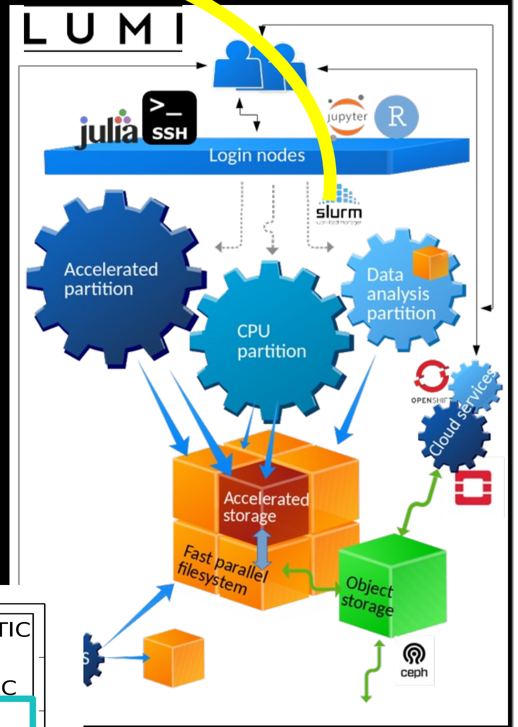

How to bridge competitive & cooperative scheduling?

- For various applications
- On different platforms
- With heterogeneous devices

We develop and/or bridge

- Domain decomposition
- **Cooperative scheduling algorithms**
- **Competitive OS and job schedulers**

Requirements

- Algorithms, system software
- **Experimental large-scale computing platform**

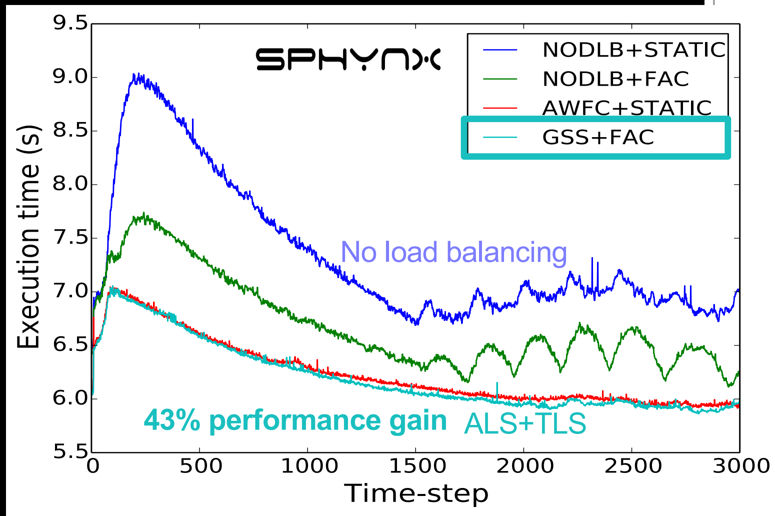

Middleware Perspective
Scheduling Algorithms Libraries & Automated Selection

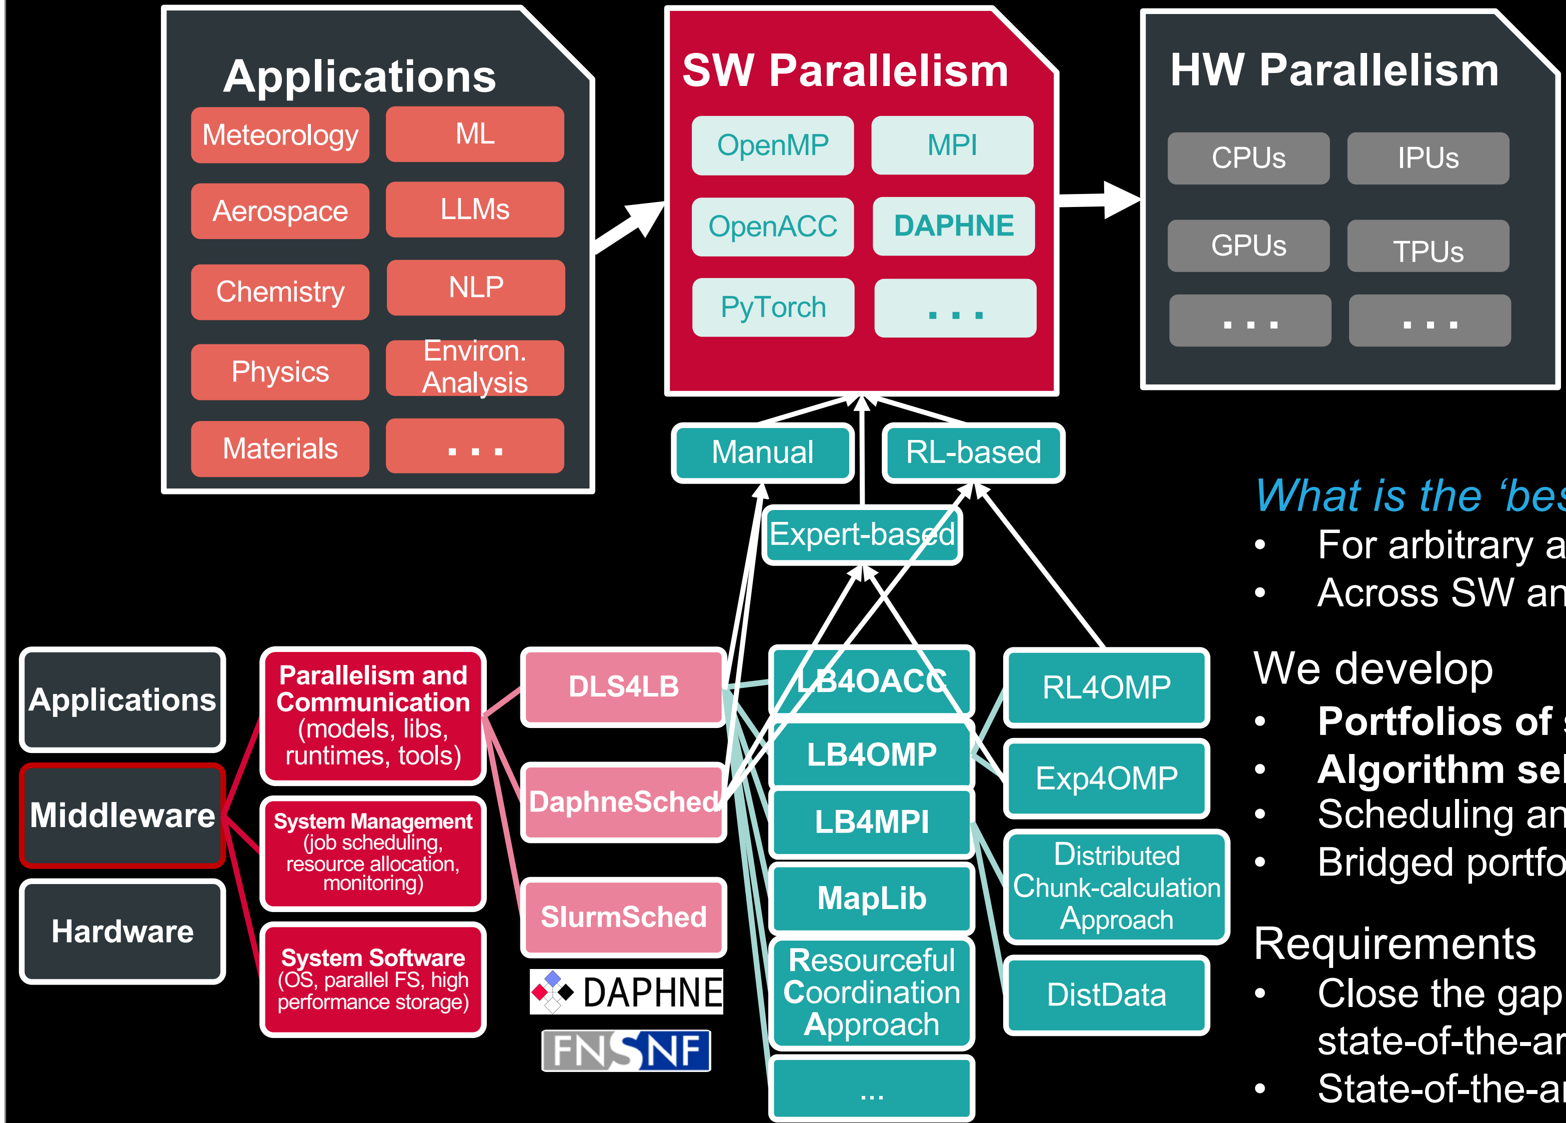

What is the 'best' scheduling algorithm?

- For arbitrary applications and phases
- Across SW and HW parallelism levels

We develop

- **Portfolios of scheduling algorithms**
- **Algorithm selection methods**
- Scheduling and selection parameters
- Bridged portfolio selections across levels

Requirements

- Close the gap between scheduling state-of-the-art and state-of-the-practice
- State-of-the-art AI/ML selection methods

System Perspective
Holistic Monitoring and Operational Data Analysis

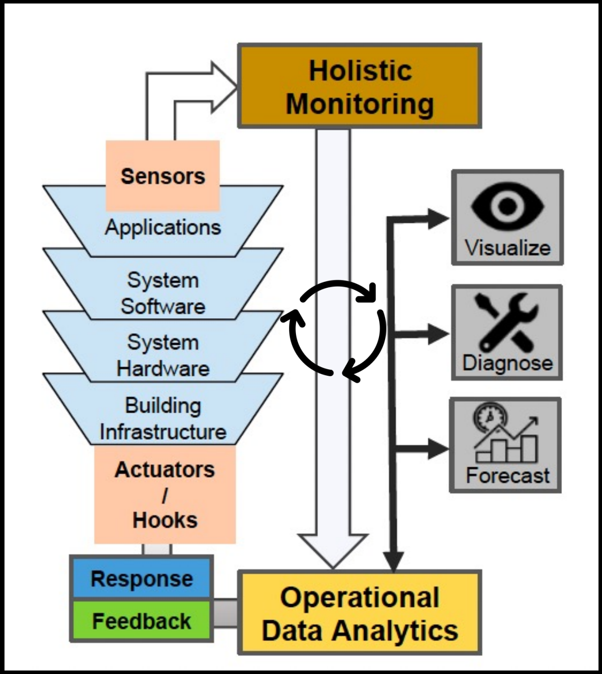

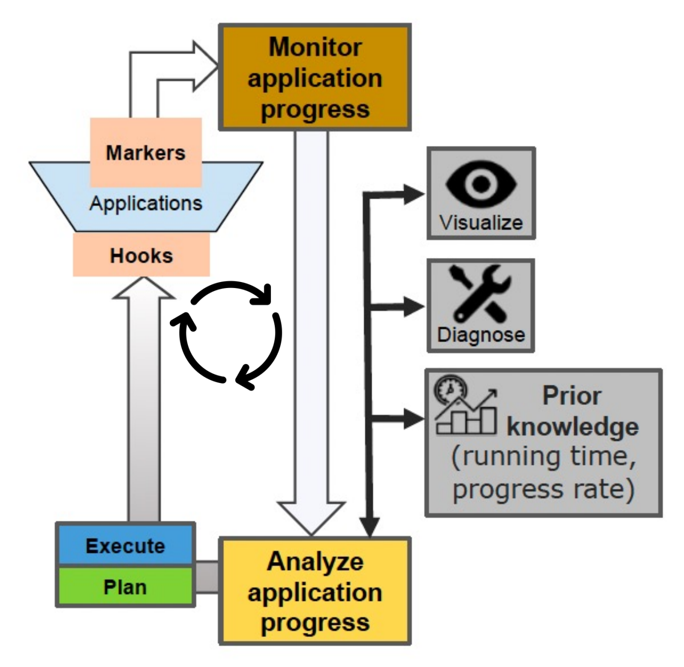

How to leverage, adapt, and extend feedback, analysis, and response for efficient HPC operations & research?

- Interactivity between applications, system software, and hardware
- True Co-Design of next generation HPC systems

We develop

- **Autonomous loops** using MAPE-K formalism (monitor, analyze, plan, execute, with knowledge)
- Recognize [malicious/inefficient] jobs to **predict resource needs**, attest workloads, **ensure security**, …

Requirements

- **Experimental large-scale computing platform**
- Cooperation with system administrators, operators, and users
- **Community building** (MODA, WAFVR, Dagstuhl 23171)

University of Basel

DMI HPC

Topics for Bachelor Theses (15 CP)

High Performance Computing Group | December 2023

Topics for Bachelor Theses
Selection (December 2023)

1. A Study of Job Scheduling Performance and Efficiency for CINECA's Marconi100
2. Automated Survey of Academic Papers through Local Large Language Models
3. Continuous Performance Evaluation of Astrophysical Applications Through CI/CD Pipelines
4. CPU Frequency Sensitivity Analysis of HPC Workloads
5. Porting and Measuring the Performance of SPH-EXA on SYCL
6. Enhancing Visualization of Cosmological Simulations through Virtual Reality
7. Performance Evaluation of In-situ Scientific Visualization Tools
8. Exploration of Mixed Precision Computations in Cosmological Simulations
9. Cyclomatic Complexity Analysis of HPC Application Codes
10. Verification and Validation of the OpenMP Standard Functionality of Scheduling Clauses

1. A Study of Job Scheduling Performance and Efficiency for CINECA's Marconi100

Context

Analyzing HPC data is crucial for performance and efficiency optimization
Usually, large HPC datasets are not publicly accessible

Motivation

Public M100 ExaData of CINECA's Marconi100 Tier-0 supercomputer
Understand and analyze a large HPC dataset to extract insights

Objectives

Understand and summarize the dataset
Implement a data analysis framework
Analysis on job-level, e.g. wait time, variability, energy
Extract insight for optimization of job management


scientific **data**

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature > scientific data > data descriptors > article

Data Descriptor | Open access | Published: 18 May 2023

M100 ExaData: a data collection campaign on the CINECA's Marconi100 Tier-0 supercomputer

Andrea Borghesi, Carmine Di Santi, Martin Molan, Mohsen Seyedkazemi Ardebili, Alessio Mauri, Massimiliano Guarrasi, Daniela Galetti, Mirko Cestari, Francesco Barchi, Luca Benini, Francesco Beneventi & Andrea Bartolini

2. Automated Survey of Academic Papers through Local Large Language Models (LLMs)

Context

- Large collection (thousands) of academic papers about scheduling in HPC
- Unclear/unknown a) gaps in existing research, b) trends, and c) potential future directions

Motivation

- Need for handling large data sets in surveys
- Era of AI and Natural Language Processing

Objectives

- Understand the landscape of LLMs
- Set up a local open-source LLM
- Develop and test survey prompts
- Automate the survey of papers
- Extract and analyze insights

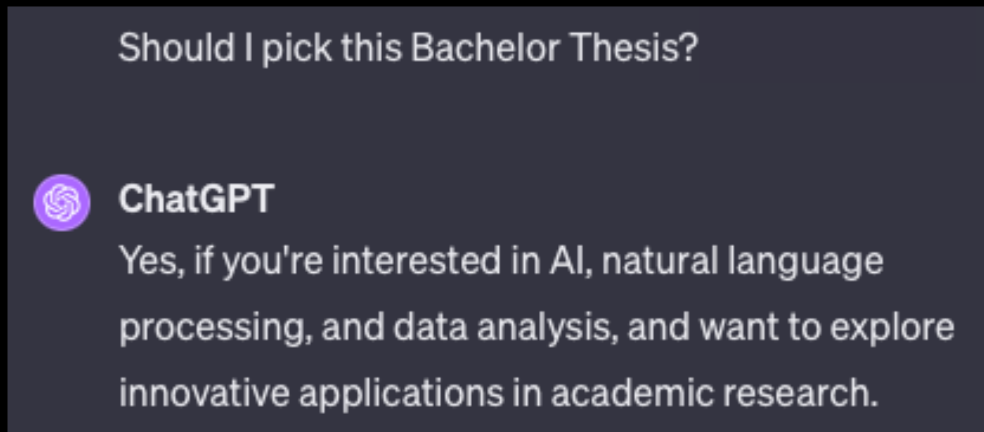

3. Continuous Performance Evaluation of Astrophysical Applications Through CI/CD Pipelines

Context

- CI/CD pipelines are useful for delivering high quality software faster
- Automation of building and testing software exists, but performance measurements do not

Motivation

- Software verification from performance aspect
- Continuous measurement of performance improvements

Objectives

- Understand the use of CI/CD pipelines
- Learn about performance measurement of software
- Integrate performance monitoring to CI/CD
- Visualize software performance for each build

4. CPU Frequency Sensitivity Analysis of HPC Workloads

Context

- Energy-efficiency is a major concern for HPC
- DVFS method can decrease consumption, but reduces performance

Motivation

- Exploring which workloads can benefit from DVFS more without sacrificing performance
- How do CPU and GPU workload behave when DVFS is applied

Objectives

- Perform frequency sensitivity analysis
- Identify frequency sensitive workloads
- Identify and model the effect of frequency changes to voltage output
- Establish a frequency sensitivity threshold

5. Porting and Measuring the Performance of SPH-EXA on SYCL

Context

SPH-EXA is a highly scalable code that can run on GPUs using CUDA and HIP
In order to provide portability across all GPU vendors, porting to SYCL is required

Motivation

Porting the code to SYCL is necessary to run on Aurora
SYCL can compile code to all GPUs, not just Intel

Objectives

Porting SPH-EXA to SYCL
Evaluating the performance of SYCL-ported SPH-EXA
Profiling the SYCL code to find performance deficiencie
Evaluate possible performance improvements

6. Enhancing Visualization of Cosmological Simulations through Virtual Reality

Context

Immersive visualization helps scientists in deepening understanding of intricate cosmological structures.
Gaps exists between high-performance simulations and visualization

Motivation

Need for high-performance immersive visualization
Need for visualization of complex physical structure

Objectives

Explore open-source VR solutions
Set up a VR scientific visualization workflow
Explore visualization methodologies for cosmological data

7. Performance Evaluation of In-situ Scientific Visualization Tools

Context

In-situ visualization optimizes I/O efficiency and offers real-time insights.
Lack of benchmarking for in-situ visualization tools on high-performance clusters

Motivation

Understand in-situ visualization workflow in HPC
Offer insights into the design of visualization architectures.

Objectives

Set up in-situ visualization frameworks in HPC context
Implement a data analysis framework
Define and implement performance analysis for visualization

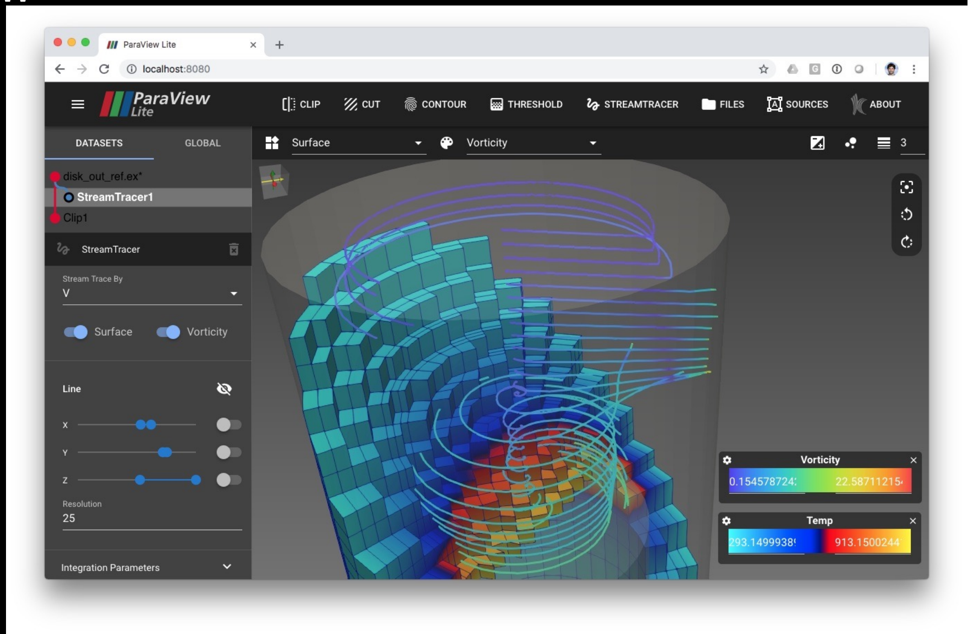

8. Exploration of Mixed Precision Computations in Cosmological Simulations

Context: Classical HPC simulations typically require 64-bit DP. Recent AI models typically requires 16-bit HP or 32-bit SP.
Motivation: Evaluate *performance* vs. *accuracy* of mixed precision computations in cosmological simulations (extremely large & long running, typically use FP64 precision)
Objectives:

- Create configurations for changing the precision of different fields in the simulation code
- Prepare and run experiments with multiple mixed-precision configurations
- Compare against the performance and accuracy of original FP64 runs and find the best performing mixed-precision configurations within acceptable errors

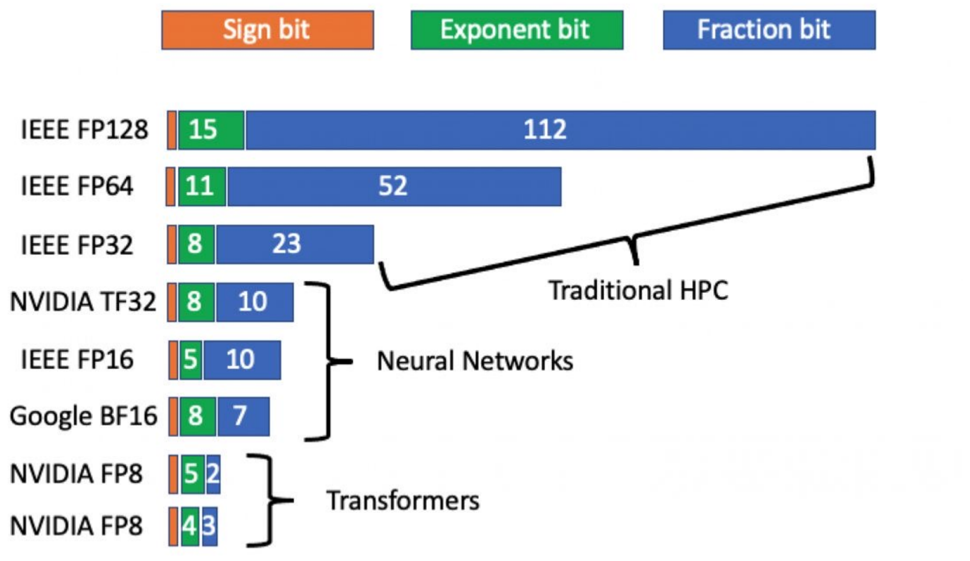

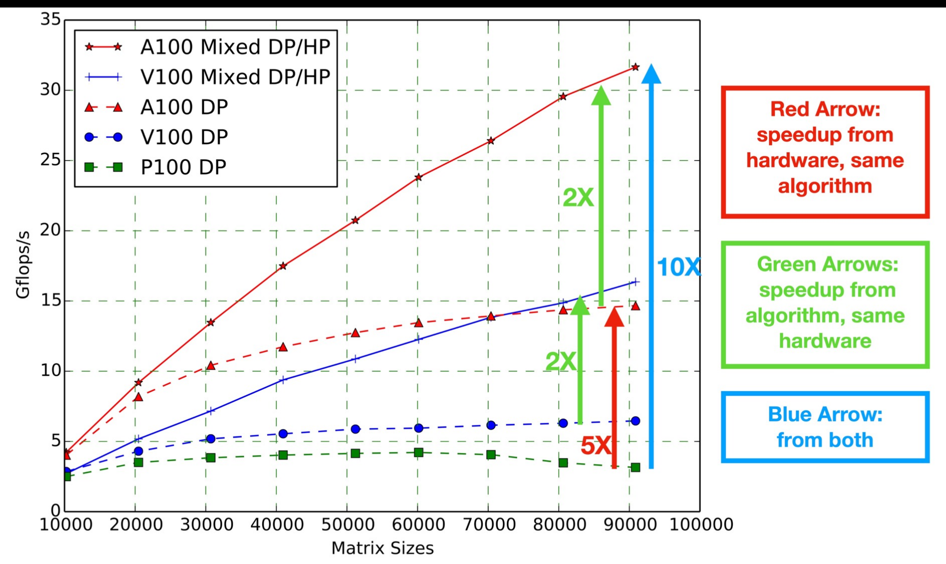

9. Cyclomatic Complexity Analysis of HPC Applications

Context: Can we infer the complexity of parallel programs by analyzing their source-code?
Goal: Analyze the *cyclomatic complexity* (# of code paths) of HPC codes.
Objectives

- *Compute cyclomatic complexity** for a large number of applications
- Investigate how *average cyclomatic complexity* compares for various applications
- Compare cyclomatic complexity for applications across programming languages & paradigms
- Quantify analysis insights through statistics and appropriate visualization

*with the Lizard[1] cyclomatic complexity analyzer.

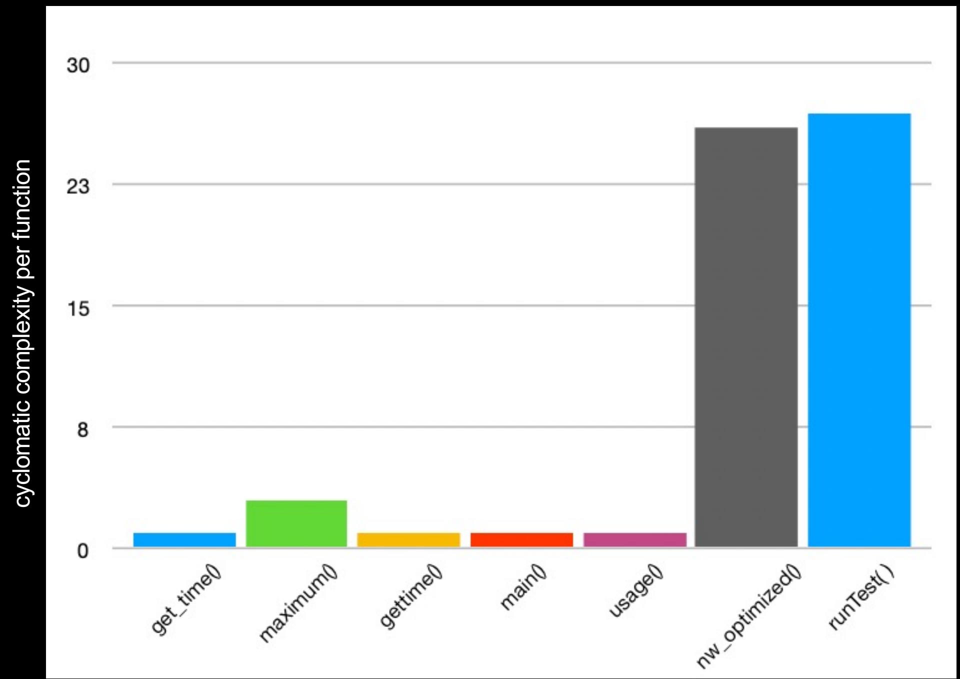

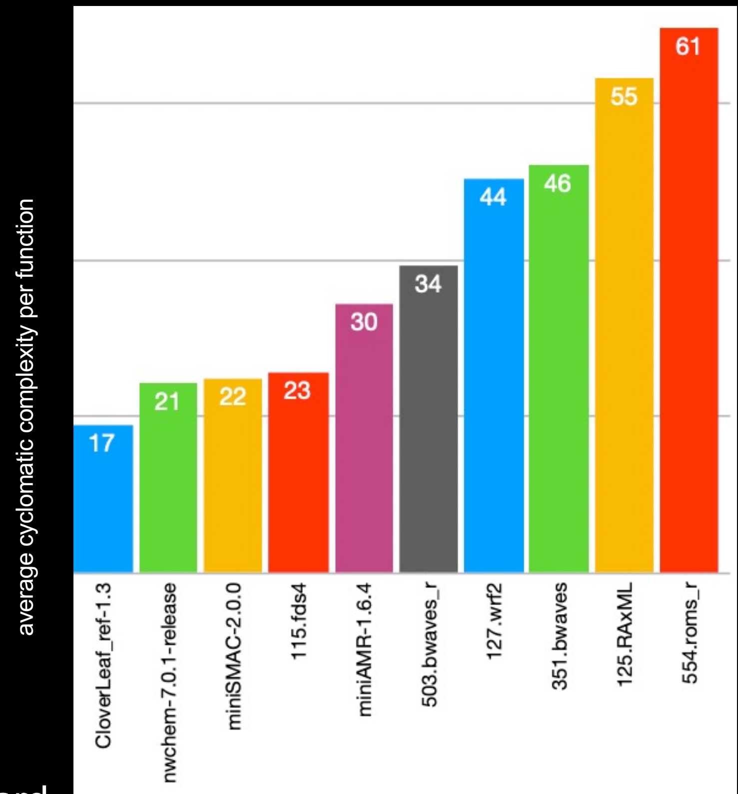

[1]https://github.com/terryyin/lizard

10. Verification and Validation of the OpenMP Standard Functionality of Scheduling Clauses

Context: OpenMP is a parallel-based programming model for performance optimization in C, C++ & Fortran. Its “directives” features listed in the OpenMP specification (spec) benefit from a Verification & Validation (V&V) testsuite, which tests various OpenMP directives to evaluate system & vendor compliance on various systems.

Motivation: The OpenMP V&V testsuite ensures that the OpenMP specification, compiler vendors & system operators are implementing OpenMP effectively. Currently, no test exists for scheduling directives and clauses.

Objective: Create unit tests to validate and verify compilers' implementation of the OpenMP specification.

- Reveal ambiguities in the OpenMP specification
- Determine missing implementation of a feature, especially scheduling
- Highlight unmentioned restriction of a feature
- Evaluate implementations for multiple platforms
- Identify and report compiler bugs

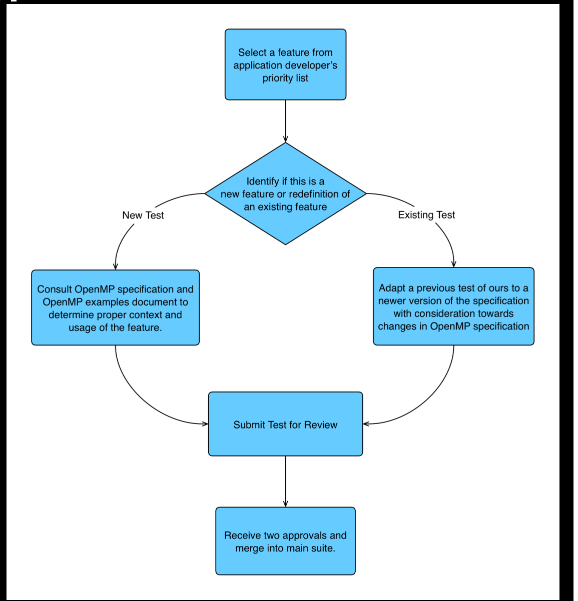

Topics for Bachelor Theses

Selection (December 2023)

1. A Study of Job Scheduling Performance and Efficiency for CINECA's Marconi100
2. Automated Survey of Academic Papers through Local Large Language Models
3. Continuous Performance Evaluation of Astrophysical Applications Through CI/CD Pipelines
4. CPU Frequency Sensitivity Analysis of HPC Workloads
5. Porting and Measuring the Performance of SPH-EXA on SYCL
6. Enhancing Visualization of Cosmological Simulations through Virtual Reality
7. Performance Evaluation of In-situ Scientific Visualization Tools
8. Exploration of Mixed Precision Computations in Cosmological Simulations
9. Cyclomatic Complexity Analysis of HPC Application Codes
10. Verification and Validation of the OpenMP Standard Functionality of Scheduling Clauses

High Performance Computing Group
Your own topic?

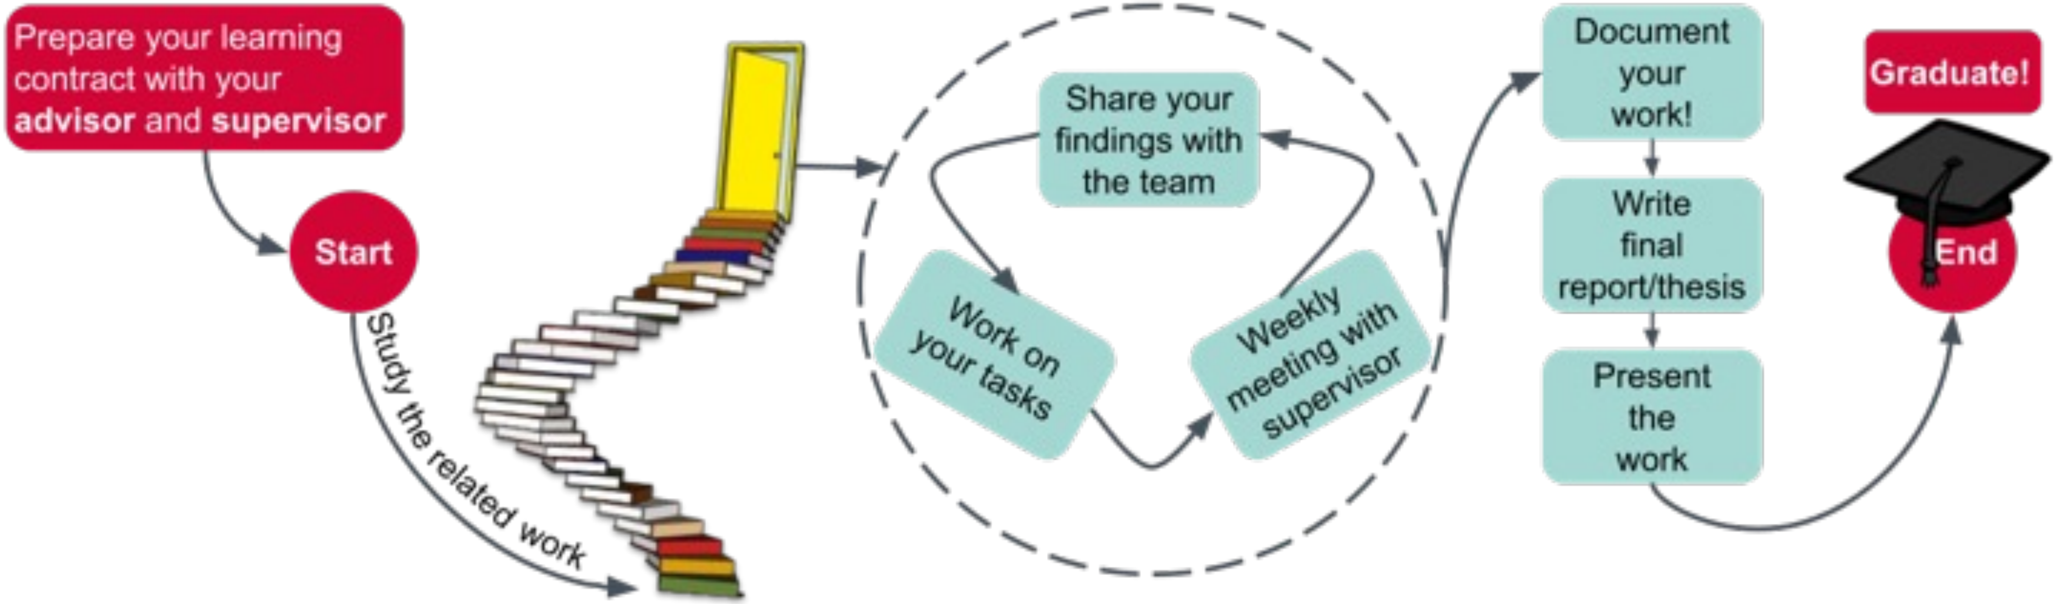

https://hpc.dmi.unibas.ch

dmi-hpc@unibas.ch

@HPC_DMI_UniBas

@hpc_dmi_unibas

Selected topics online https://hpc.dmi.unibas.ch/en/theses/

Fall 2023