H2020-FETHPC-3-2017 - Exascale HPC ecosystem development

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D3.4
First report on the organisation of WP3 workshops during HPC Summit Week 2019

Final

Version: 1.1
Author(s): Mathis Bode (RWTH Aachen University), Giovanni Aloisio (CMCC), Sandro Fiore (CMCC)
Date: 24/09/2019

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-800957**	
	Project Title: European eXtreme Data and Computing Initiative - 2	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: < **D3.**>	
	Deliverable Nature: <DOC_TYPE: Report>	
	Dissemination Level: PU	**Contractual Date of Delivery:** 24 / June / 2019
		Actual Date of Delivery: 21 / June / 2019 (Revised version resubmitted 30 /09 /2019)
	EC Project Officer: Evangelos Floros	

* - The dissemination level are indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	**Title: First report on the organisation of WP3 workshops during HPC Summit Week 2019**	
	ID: D3.4	
	Version: <1.1>	**Status:** ***Final***
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D3.4-v1.1.docx	
Authorship	**Written by:**	Mathis Bode
	Contributors:	Stéphane Requena, GENCI/PRACE
	Reviewed by:	Paul Carpenter, BSC Katia Soto, PRACE aisbl
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	05/06/2019	Draft	Before internal review
0.2	12/06/2019	Draft	After internal review 1
0.3	18/06/2019	Draft	After internal review 2
1.0	21/06/2019	Final version	Approved by TB/MB
1.1	24/09/2019	Final version	Revised version following Interim Review.

Document Keywords

Keywords:	PRACE, Research Infrastructure, HPDA, Big data, HPC, Workflows.

Table of Contents

List of Figures

References and Applicable Documents

[1] http://www.exdci.eu
[2] http://www.prace-project.eu

List of Acronyms and Abbreviations

AI	Artificial Intelligence
BDEC	Big Data and Extreme-scale Computing
CFD	Computational Fluids Dynamics
CoE	Centres of Excellence for Computing Applications
D	Deliverable
DSL	Domain Specific Language
EC	European Commission
EIP	European Innovation Partnership
FETHPC	Future and Emerging Technologies High Performance Computing
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
HPDA	High Performance Data Analytics
M	Month
ML	Machine Learning
PRACE	Partnership for Advanced Computing in Europe

D3.4	1st report on the organisation of WP3 workshops during HPC Summit Week 2019
WP	Work Package

Executive Summary

This deliverable summarizes the most relevant information about the EXDCI-2-WP3-organized "*1st European Communities Workshop on Exascale Computing*" focusing this year on "*High Performance Data Analytics*" (HPDA) held on 16 May 2019 during the EuroHPC Summit Week 2019.

The workshop was organized in two sessions, one dedicated to the Centres of Excellence for Computing Applications (CoEs) and another one to the FETHPC projects. In both sessions, the representatives of the different projects reported about the major challenges and issues faced with respect to the HPDA topic, the outcomes achieved so far as well as plans for the next years.

The CoE session was chaired by Guy Lonsdale (head of the FocusCoE project) whereas the FETHPC one was led by Giovanni Aloisio (EXDCI-2 [1] – Task 3.2 leader). Mathis Bode (EXDCI-2 [1] – WP3 leader) chaired the workshop in total.

Of the ten existing CoEs, representatives of six CoEs and FocusCoE were attending the meeting: ESiWACE2, EXCELLERAT, HiDALGO, MaX, CompBioMed, BioExcel, and FocusCoE. The FETHPC projects were represented by MAESTRO, ASPIDE, VESTEC, ExaNeSt, and Sage2, for a total of eleven presentations. The presentation covered various aspects of HPDA from big data processing and parallel data analysis, heterogeneous infrastructure to deal with different aspects of HPDA, to provenance and reproducibility of simulation results.

A wrap-up discussion was also held at the end of the workshop to better clarify some aspects not addressed in the previous question time and to summarize some conclusions. In this respect, the discussions and questions showed that many European communities are facing similar challenges and would be thankful for more joint efforts and combined solutions as well as networking activities like this workshop to ease experiences, results, solutions, and best practices sharing. For WP3, it was important to get this feedback, which helps to produce a more relevant High Performance Computing (HPC) roadmap.

Overall the workshop was successful, with about 50 people attending the two sessions. It is planned to organize a "2nd European Communities Workshop on Exascale Computing" during the next EuroHPC Summit Week in 2020, with a different focus, most likely (based on preliminary thoughts) on Artificial Intelligence (AI) and Machine Learning (ML) topics.

As next actions, the collaboration with the BDEC (Big Data and Extreme-scale Computing) Application Group will be intensified to see the needs with respect to HPDA in more applications. Furthermore, this will also allow to find more already existing solutions, which can be shared with the workshop participants as well as the European community.

1 Introduction

Work Package 3 (WP3) “Excellence in HPC applications and usages” of EXDCI-2 [1] focuses on applications and best practice usage in the context of potential requirements towards Exascale platforms. This concerns applications from both classical High Performance Computing (HPC) and High Performance Data Analytics (HPDA). As ascertained during the EXDCI project (i.e. the forerunner to the current project EXDCI-2), Europe is developing a significant fraction of the applications used in the world and is the biggest producer of data. Therefore, this effort must be continued as science opportunities are evolving very quickly with the expected availability of Exascale supercomputers (and corresponding HPDA tools). For example, the following three points can be named:

- Large research infrastructure designs are evolving with new capabilities of HPC and Big Data
- New tools and approaches are increasingly needed to take into account new technical realities
- As users get a better understanding of the potential of HPC and Big Data applications, possibly combined, science goals evolve rapidly and new application domains appear

For scientific applications, WP3 relies on the PRACE (Partnership for Advanced Computing in Europe, [2]) Scientific Case and interacts with the PRACE user communities, and the HPC Centres of Excellence for Computing Applications (CoEs) to focus on the Exascale aspects, as well as on the influence of specific technological or algorithmic innovations. WP3 roadmaps HPC applications and usages and coordinates with HPC user communities and CoEs.

Similar considerations apply to industrial applications, but in this case, it is often required to fit the application in a complex proprietary workflow, for instance a “digital twin” of a product or manufacturing setup. This requires engaging and working directly with the users or their collective organizations including European Innovation Partnerships (EIPs) and national initiatives. The roadmap results are of importance to avoid gaps in the value chain, but also to evaluate scenarios where “disruptive innovation” entails changes in the value chain or permit different entry strategies.

To fulfil this role within the EXDCI-2 project, WP3 is divided into three tasks:

- Task 3.1 – Roadmap of HPC applications and usages
- Task 3.2 – Engagement with HPC users communities and CoEs
- Task 3.3 – Preparation of industrial codes to exascale

Mathis Bode (RWTH Aachen University) is WP3 leader and Stéphane Requena (GENCI/PRACE) is WP3 co-leader.

This deliverable entitled “D3.4 – First report on the organisation of WP3 workshops during HPC Summit Week 2019” was postponed from M15 to M16 due to the fact that the EuroHPC Summit Week 2019 in Poznan, Poland took place in M15 and was essential for this deliverable. Consequently, the deliverable outlines the work carried out in WP3 until and including M15 of the EXDCI-2 project. Relevant milestones during this period for this deliverable are:

- MS32 – 1st brainstorming session
- MS34 – Workshop during HPC Summit Week 2019 (Workshop held)

A detailed description of MS32 can be found in the deliverable “D3.2 – First Report on joint brainstorming sessions among scientific and industrial users communities” also due in M16 (postponed from M15). This deliverable provides a summary of the workshop organized by WP3 during the EuroHPC Summit Week 2019. The aimed target of the workshop is explained and details about different contributions to the workshop are summarized. Conclusions with respect to the focus topic of the workshop “HPDA” are given.

2 Summary of Deliverable

This delivery summarizes the organization of the WP3-organized workshop during the EuroHPC Summit Week 2019. The description is separated into three parts:

- Summary of the Workshop “1st European Communities Workshop on Exascale Computing” (cf. Chapter 3)
- Summary of CoE Presentations (cf. Chapter 4)
- Summary of FETHPC Presentations (cf. Chapter 5)

The report finishes with some conclusions and a summary of next steps.

3 Summary of the Workshop

WP3 organized a workshop during the EuroHPC Summit Week 2019 (for details on the overall week, see EXDCI-2-D6.1) with the target to discuss important Exascale topics among various scientific communities. The workshop was designed as the first of a series of workshops and was therefore denoted as “1st European Communities Workshop on Exascale Computing“. It is planned to organize a second version during the EuroHPC Summit Week 2020. In order to support fruitful discussions, HPDA was chosen as target topic and 11 speakers were invited to present their ideas/experiences/solutions/questions with respect to this topic.

As it is important to get a detailed idea about the complete European Exascale landscape, the EXDCI-2-WP3 workshop was organized in close collaboration with Guy Lonsdale from FocusCoE. Finally, six speakers representing different CoEs and five speakers representing different Future and Emerging Technologies High Performance Computing (FETHPC) projects were selected. The agenda of the workshop is shown in Figure 1 and Figure 2.

The workshop took place on Thursday, May 16th, 2019 from 2:30pm to 7:00pm, which was after the official end of the PRACEdays 2019. About 50 people attended the workshop in total, which was a success. The CoE session was chaired by Guy Lonsdale (head of the FocusCoE project) whereas the FETHPC one was led by Giovanni Aloisio (EXDCI-2 [1] – Task3.2 leader). Overall, the workshop was chaired by Mathis Bode (EXDCI-2 [1] – WP3 leader). The presentation covered various aspects of HPDA from Big Data processing and parallel data analysis, to provenance and reproducibility of simulation results. The discussions and questions showed that many European communities face similar challenges and would be thankful for more joint efforts and combined solutions. For WP3, it was important to get this feedback, which helps to produce a more relevant HPC roadmap.

1st European Communities Workshop on Exascale Computing

Focus on High Performance Data Analytics

May 16, 2019, Poznan, Poland

hosted by EuroHPC Summit Week 2019, May 13 - 17, 2019

Agenda May 16, 2019

14:30	**Welcome and brief introduction**	**Giovanni Aloisio, Mathis Bode**
14:30-16:30	**Session 1 – CoEs & HPDA**	**Chair: Guy Lonsdale**
14:30-14:50	ESiWACE	Sandro Fiore (CMCC)
14:50-15:10	EXCELLERAT	Bastian Koller - Dimitris Liparas (HLRS)
15:10-15:30	HiDALGO	Francisco Javier Nieto de Santos (ATOS)
15:30-15:50	MaX	Sebastiaan Huber (EPFL)
15:50-16:10	CompBioMed	Hector Martinez (Univ. of Oxford)
16:10-16:30	BioExcel	Rossen Apostolov (KTH)
16:30-17:00	*Coffee break*	

Figure 1 Agenda of the 1st European Communities Workshop on Exascale Computing (1/2)

17:00-19:00	**Session 2 – FETHPC & HPDA**	**Chair: Giovanni Aloisio**
17:00-17:20	Maestro: Towards a Memory- and Data-aware Middleware	Dirk Pleiter (Jülich Supercomputing Centre) MAESTRO Project
17:20-17:40	Towards data intensive aware programming models for Exascale systems	Francisco Javier Garcia-Blas (Universidad Carlos III de Madrid) ASPIDE Project
17:40-18:00	The VESTEC project: Fusing HPC and real-time sensor data for urgent decision making	Gordon Gibb (EPCC, Univ. of Edinburgh) VESTEC Project
18:00-18:20	ExaNeSt: Low-Latency Communication and Acceleration in a liquid-cooled energy-efficient Prototype Rack	Manolis G.H. Katevenis (FORTH, Crete, GR) ExaNeSt Project
18:20-18:40	Sage2: Architecting a storage platform for the extreme data era	Sai Narasimhamurthy (Seagate Technology, LLC) Sage2 Project
18:40-18:55	Final Discussion	
18:55-19:00	**Wrap up and closing session**	**Giovanni Aloisio**

Program Committee
Giovanni Aloisio
Guy Lonsdale
Mathis Bode
Stefan Krieg
Jean-Claude André
Sandro Fiore

This event is funded by EXDCI2: the EXDCI2 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 800957 http://www.exdci.eu

Figure 2 Agenda of the 1st European Communities Workshop on Exascale Computing (2/2)

4 Summary of CoE Presentations

In the next sub-sections, the HPDA contributions from the following CoE projects are presented: ESiWACE2, EXCELLERAT, HiDALGO, MaX, CompBioMed, and BioExcel.

4.1 ESiWACE2

Speaker: Sandro Fiore (CMCC)

Abstract

ESiWACE is the Centre of Excellence in Simulation of Weather and Climate in Europe. It is now at the end of the first phase, while the second one just started at the beginning of 2019. To improve efficiency and productivity of numerical weather and climate simulation on high performance computing platforms ESiWACE has improved and supported scalability of models, tools and data management on state-of-the-art supercomputer systems, usability of models and tools throughout the European HPC eco-system, and the exploitability of the huge amount of resulting data. In its second phase, ESiWACE will enable leading European weather and climate models to leverage the performance of pre-exascale systems with regard to both compute and data capacity as soon as possible and prepare the weather and climate community to be able to make use of exascale systems when they become available. With respect to the workshop topic, ESiWACE will improve the data management tool chain from weather and climate simulations to address data challenges related to I/O, data analytics, post-processing and visualization at scale, through the development of specific middleware components (ESDM). The talk addresses in particular a set of activities in the data management area, from I/O to HPDA. It gives an overview about what has been already accomplished and what is planned as a future work until the end of the second phase of the project (December 2022). A test case related to the Ophidia HPDA framework is also presented in a real HPC settings on the Athena Cluster available at the CMCC SuperComputing Centre.

Talk highlights

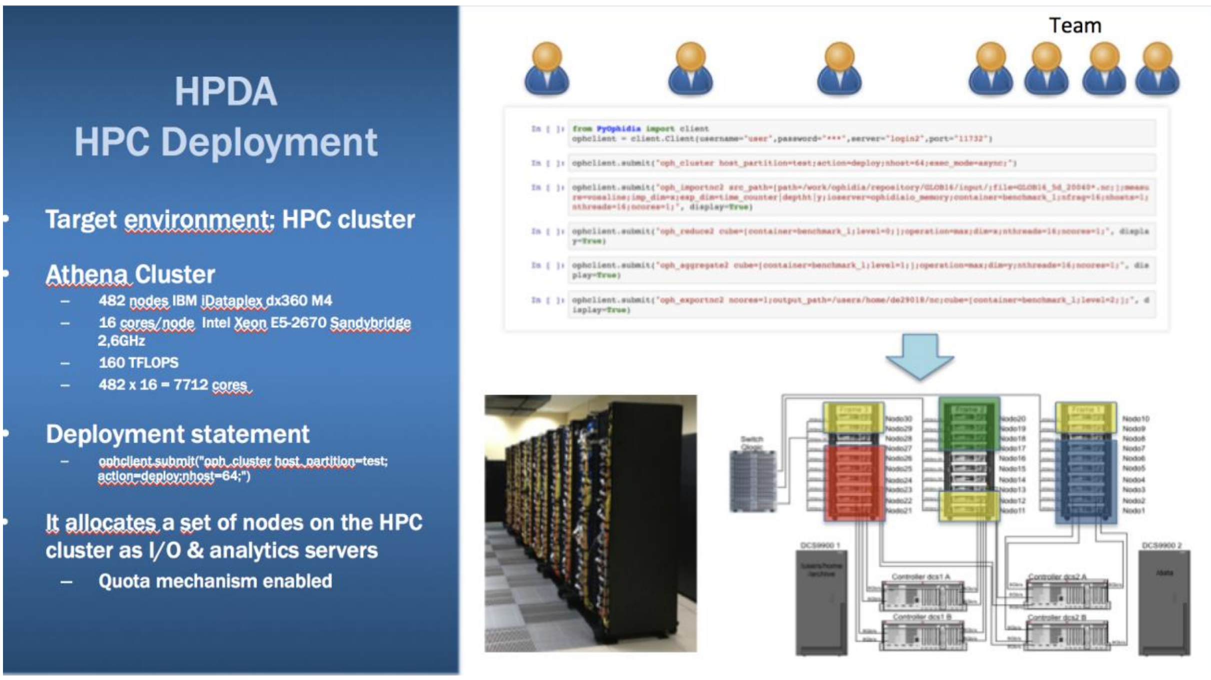

Figure 3 The Ophidia HPDA framework deployment on a HPC infrastructure at the CMCC SuperComputing Centre. The deployment can be enacted via Python API and it is completely isolated, giving each user his/her own HPDA instance.

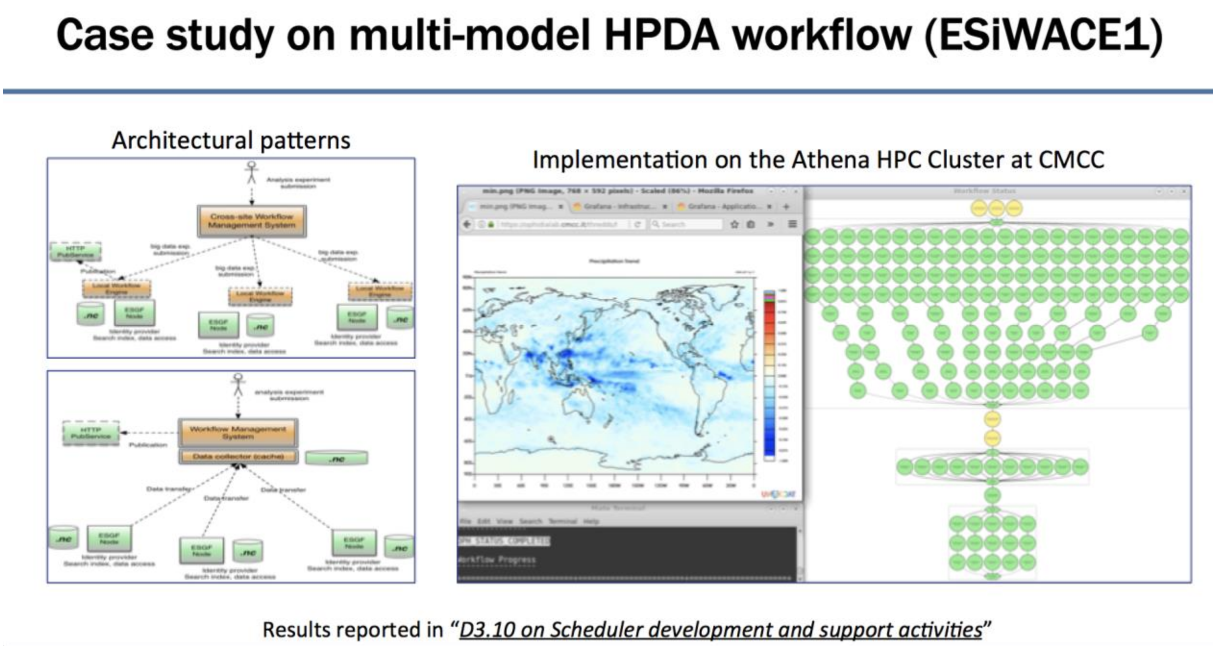

Figure 4 Multi-model HPDA workflow use case in the ESiWACE2 CoE running on the Athena HPC cluster at CMCC. The workflow has been executed and validated on eleven models from CMIP5.

4.2 EXCELLERAT

Speaker: Bastian Koller - Dimitris Liparas (HLRS)

Abstract

The EXCELLERAT consortium brings together the necessary expertise in Europe to set up a Centre of Excellence (CoE) in engineering, offering a broad service portfolio to its customers. All of this aims at supporting strongly the engineering community and paving the way for the evolution of engineering applications towards best use of Exascale technologies. Thus, EXCELLERAT will work on a set of reference applications (namely Nek5000, Alya, AVBP, Fluidity, FEniCS, Flucs). Those were selected as key representatives of challenges faced by broad parts of the engineering domain, when moving to Exascale systems. Overall, the selected six reference applications will drive the identification and implementation of services to be provided by the CoE. Work on these applications has started from a pre-analysis, done before the start of the project and now enters a phase of further evolution, by applying mechanisms for maintenance, optimization and scaling. All data-related aspects will be part of the Centre's services, including secured and optimized Data Transfer, High Performance Data Analytics (HPDA) and Machine Learning, thereby enhancing simulation runs and results.

HPDA activities within EXCELLERAT will focus on the problems arising from the analytics of Exascale simulation data. The value of such data lies only to a small extent in accurate integral values, e.g. forces or heat loads. It will be of much higher importance for the engineer to understand complex physical mechanisms, sensitivities of the performance with respect to design parameters and uncertainties connected to the simulation. In EXCELLERAT, in-situ post-processing and visualization techniques will be established as standard monitoring mechanisms, such that: (i) results are already available during the runtime of a simulation, avoiding large amount of disk I/O and (ii) trends and sensitivities can be indicated earlier during the simulation and thus, computing time can be reduced. Of special importance will be the development of new and adaptation of existing advanced analytics tools, which are indispensable for an efficient work flow, as well as machine learning/AI approaches, as an alternative to long existing response surface methods and physical-based models.

Talk highlights

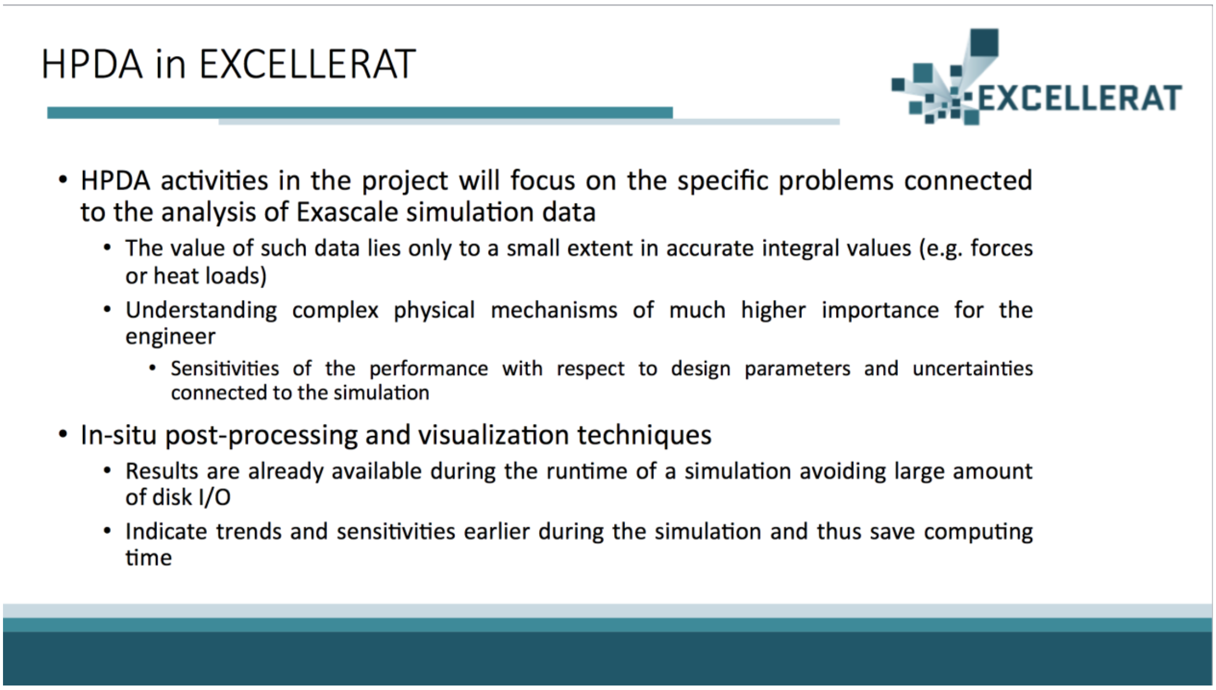

Figure 5 HPDA in EXCELLERAT will focus on the analysis of Exascale simulation data. In-situ techniques are adopted to streamline the data pipeline through the processing and visualization steps.

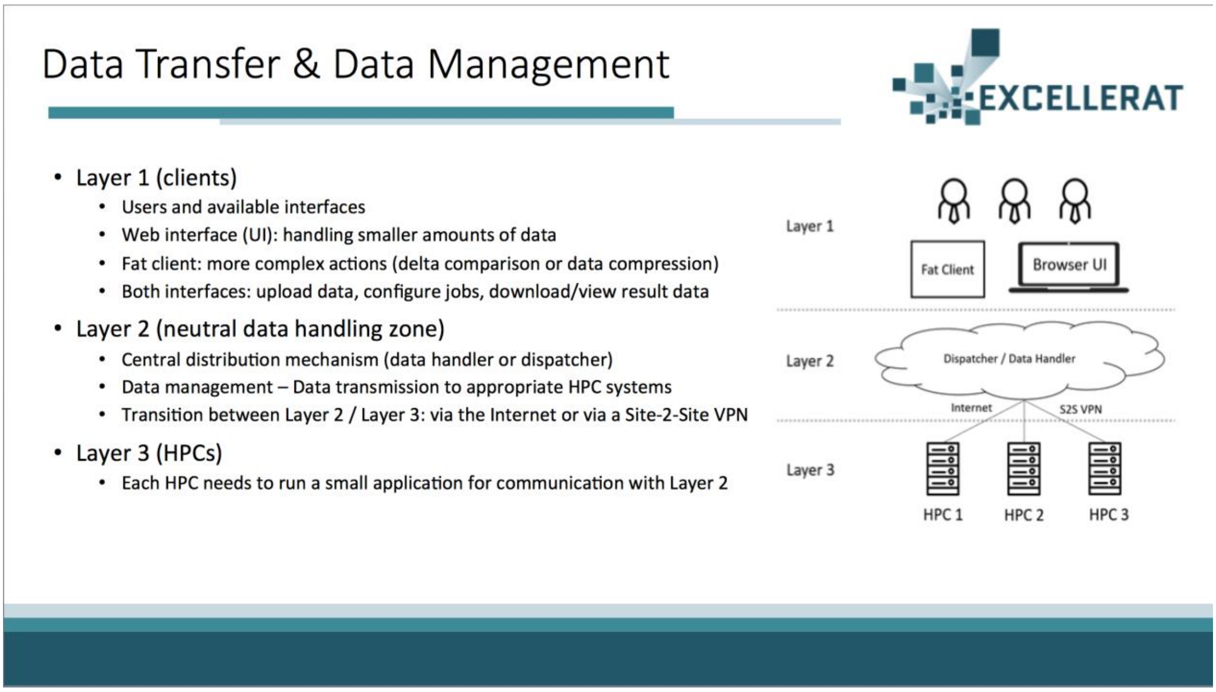

Figure 6 EXCELLERAT 3-layer design. Layer 1 is about the client applications. Layer 2 is about the neutral data handling zone which will implement a set of data management strategies/approaches to efficiently handle data, including HPDA. Layer 3 is about the HPC resources that are connected to Layer 2 via a lightweight application.

4.3 HiDALGO

Speaker: Francisco Javier Nieto de Santos (HLRS)

Abstract

The presentation introduces the HiDALGO project, as one of the new CoEs for HPC, focused on using such technology for solving global challenges. It shows the kind of workflows that the applications in the project have and how HPDA and AI are applied. It also provides an overview of the ongoing tasks related to HPDA, now focused mainly on analyzing the tools that will be used (i.e. Apache Spark, Flink, Dask, etc…) and on performing some benchmarking in order to understand how they scale up and the potential bottlenecks.

Talk highlights

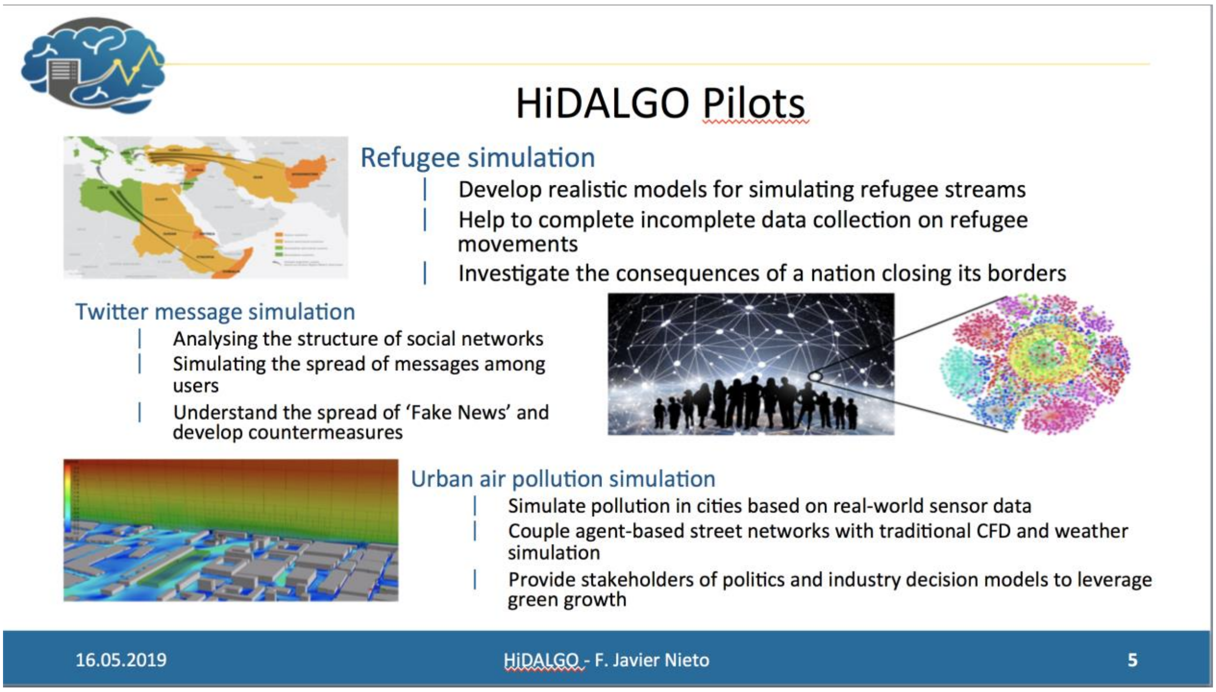

Figure 7 HiDALGO pilots include refugee simulation, Twitter message simulation, and urban air pollution simulation which all need HPDA. As discussed at the workshop, security and privacy aspects must be and are properly addressed in the use cases.

Figure 8 Technologies for HPDA exploited/evaluated in the HiDALGO CoE. Some of them address the analysis and mining of static datasets, other ones relate to streaming data analysis and some more are at the storage layer.

4.4 MaX

Speaker: Sebastiaan Huber (EPFL)

Abstract

In recent years, there has been a great increase in the performance and capabilities of computers, and preparations for the Exascale transition are in motion. Materials science has greatly benefited from this computational boom, which is continuously boosting research, the discovery of new materials and the development of simulation codes. The "materials by design" approach has become very powerful, but requires running large numbers of simulations and building databases of computed properties. A key challenge is the need to automatically prepare, execute and monitor workflows of calculations, and then retrieve and store the results in a format that is easy to browse and query. The AiiDA open-source framework and the Materials Cloud dissemination platform, provide researchers with the tools that fulfil these conditions. The AiiDA and Materials Cloud combination has matured into an ecosystem with multiple ORM (Object Relational Mapper) backend options for increased performance and flexibility, a powerful graph querying tool for easy result analysis, a redesigned plugin system to simplify external user contributions and an improved workflow architecture.

Talk highlights

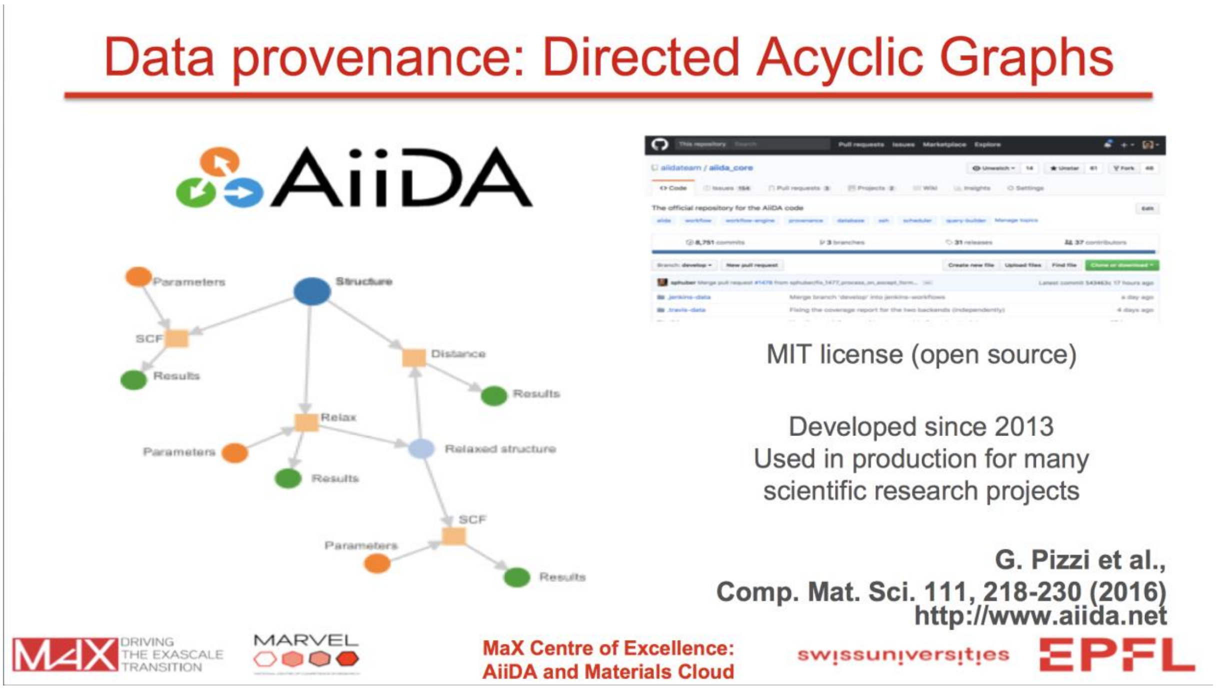

Figure 9 Data provenance tracking with the AiiDA framework in the MaX CoE. Used in production in different contexts, AiiDA provides a tool to track huge lineage information in the form of Directed Acyclic Graphs (DAGs).

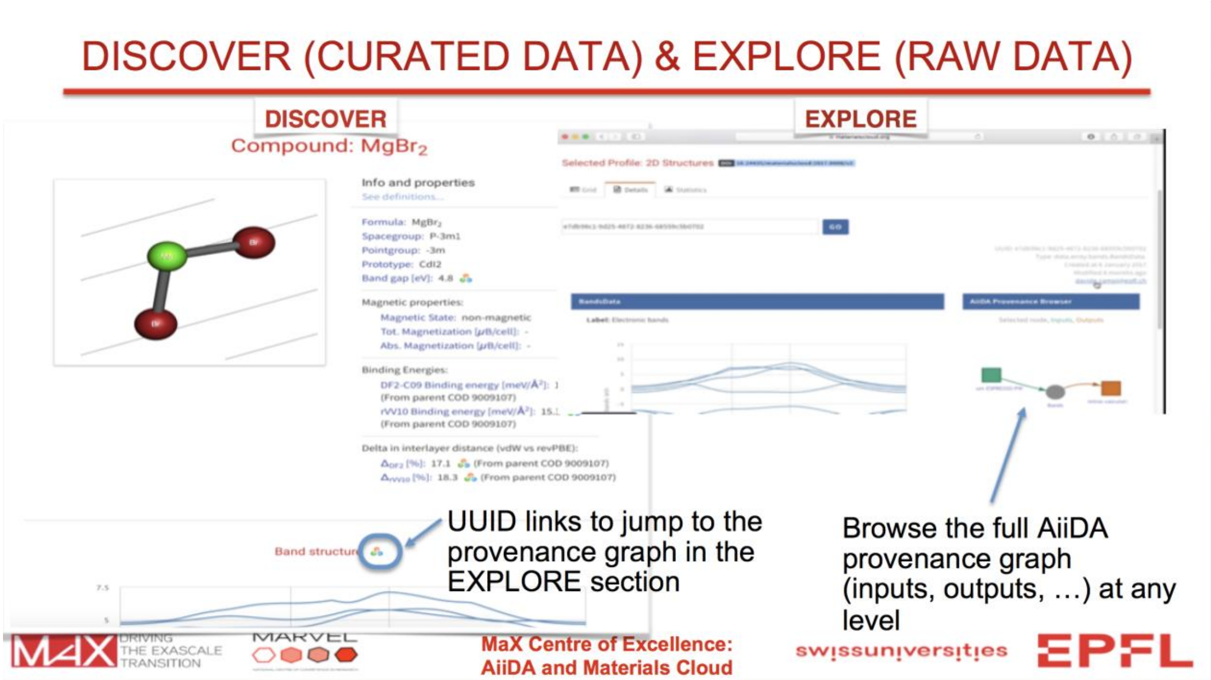

Figure 10 Discovery and exploration workflow in the MaX CoE. Users can navigate from compounds to the provenance graph getting access to the full AiiDA DAG.

4.5 CompBioMed

Speaker: Hector Martinez (Univ. of Oxford)

Abstract

Acute myocardial ischemia is a major cause of sudden arrhythmic death. Anti-arrhythmic treatments or side-effects associated to cancer therapies can produce cardiotoxic effects increasing the occurrence of adverse cardiac events, especially in patients with coronary artery disease. In-vivo and in-vitro drug trials embed several complications regarding ethics and costs. If experiments include other species, the results are not necessarily translational to humans. High performance computing simulations were conducted using a computational multiscale model of acute ischemic human ventricles embedded in the torso, validated with extensive experimental and clinical data. Tissue, cellular and subcellular dynamics were represented to enable simulating the effects of drugs and disease. From the simulations, biomarkers obtained from the simulated 12-lead ECG were quantified and compared to arrhythmic risk computed as the vulnerability window (VW) for re-entry. The effects of pharmacological action were simulated by altering sodium, potassium and calcium ionic conductance.
High performance computing simulations were in agreement with clinical data on the effect of ischemia and drug action on the ECG and arrhythmic risk. Drug-induced changes in ion channels resulted in substantial changes in ischemia-induced arrhythmic risk. The simulation of class I antiarrhythmic drug effects increased arrhythmic risk by reducing tissue excitability. Results are consistent with the high mortality in patients under ischemic risk using flecainide/encainide (CAST trial). Mild potassium block (class III antiarrhythmic drugs) reduced dispersion in refractoriness and hence risk, but stronger block facilitated other pro-arrhythmic mechanism associated to this type of drug which prolongs the QT interval (SWORD trial). In this study, high spatio-temporal resolution of High Performance Computing simulations data enables mechanistic insights into the pro-arrhythmic effects paradoxically

triggered by classic anti-arrhythmic drugs, such as flecainide and sotalol. This technology can also point out to new directions towards new anti-arrhythmic therapies.

Talk highlights

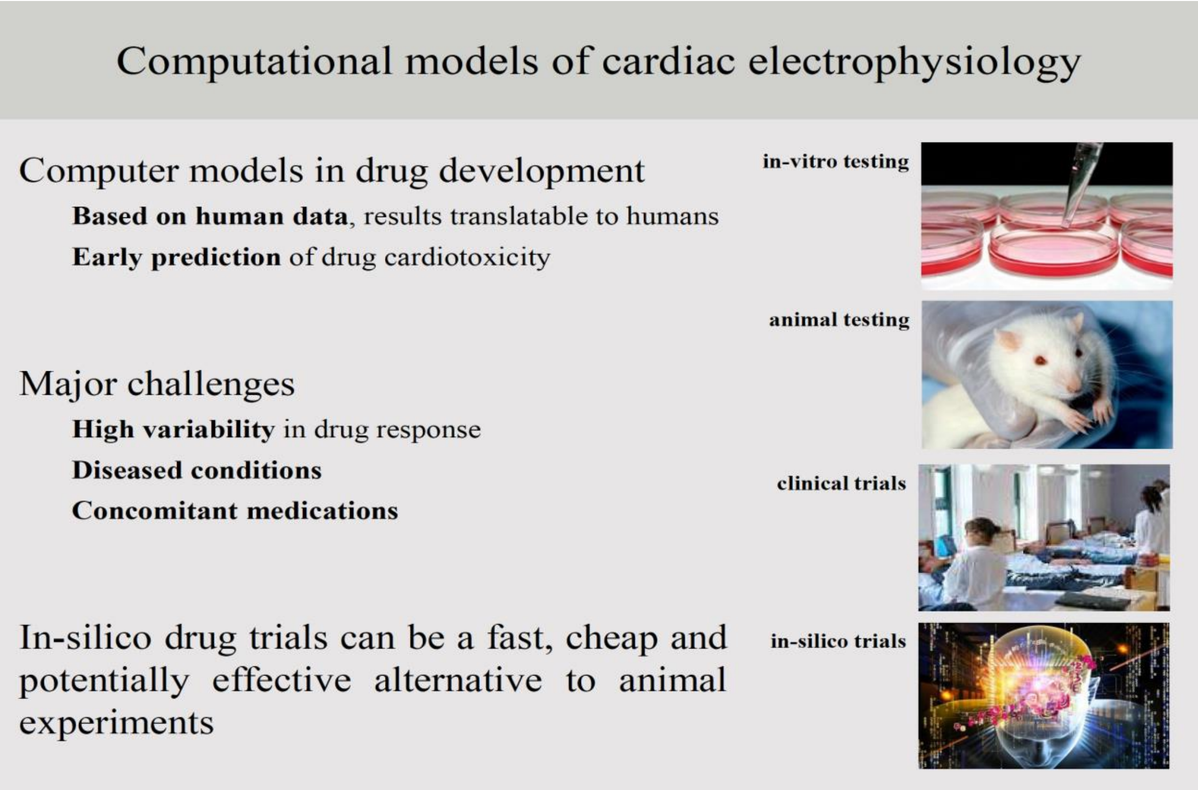

Figure 11 Computational models of cardiac electrophysiology in the CompBioMed CoE. The in-silico drug trials can result in being a very effective alternative to animal experiments.

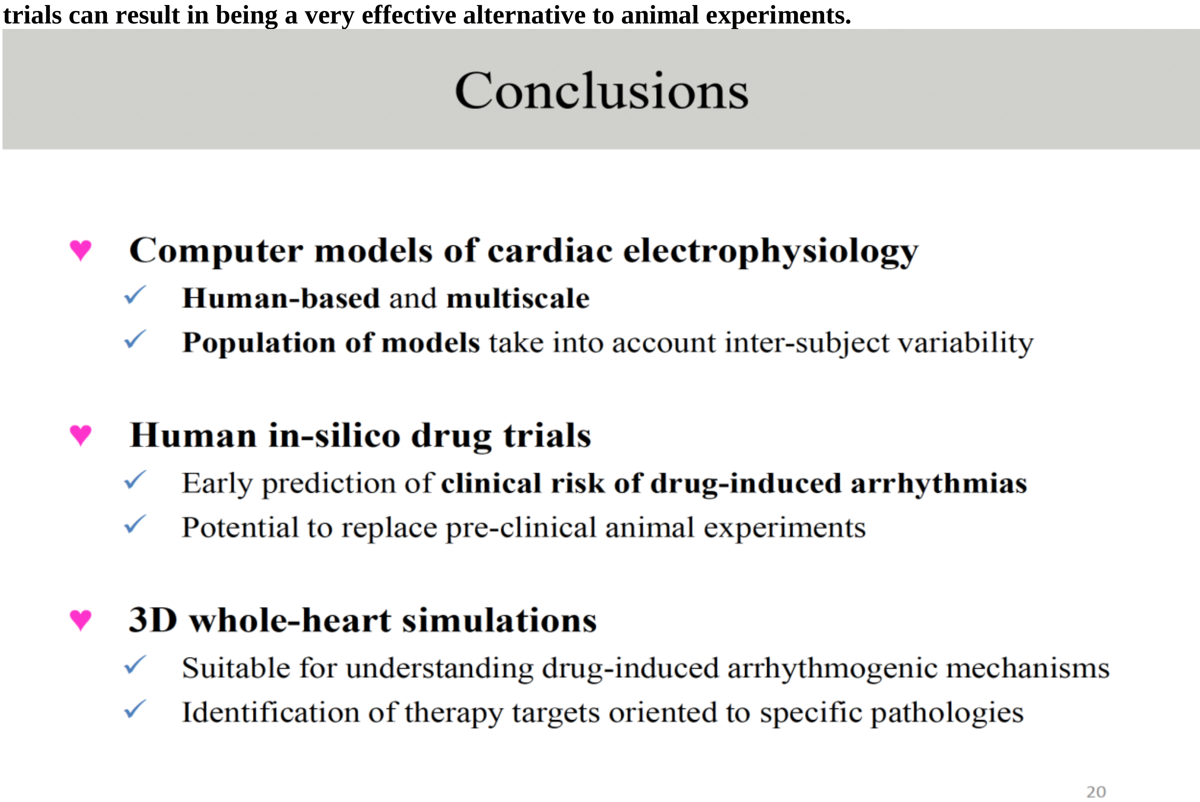

Figure 12 Main conclusions form the CompBioMed CoE related to the (i) Computer models of cardiac electrophysiology, (ii) human in-silico drug trials, and (iii) 3D whole-heart simulations.

4.6 BioExcel

Speaker: Rossen Apostolov (KTH)

Abstract

Biomolecular modelling and simulations are becoming an ever larger user of modern HPC and High Throughput Computing (HTC) resources. Accurate predictions of complex molecular structures, their dynamics and function can be dramatically improved by incorporating experimental data in the computational analysis. The importance of data analytics techniques in the field is growing. In this talk we present ongoing research work in BioExcel Centre of Excellence and our efforts towards convergence of HPC, HTC and HPDA in the area of Life Sciences.

Talk highlights

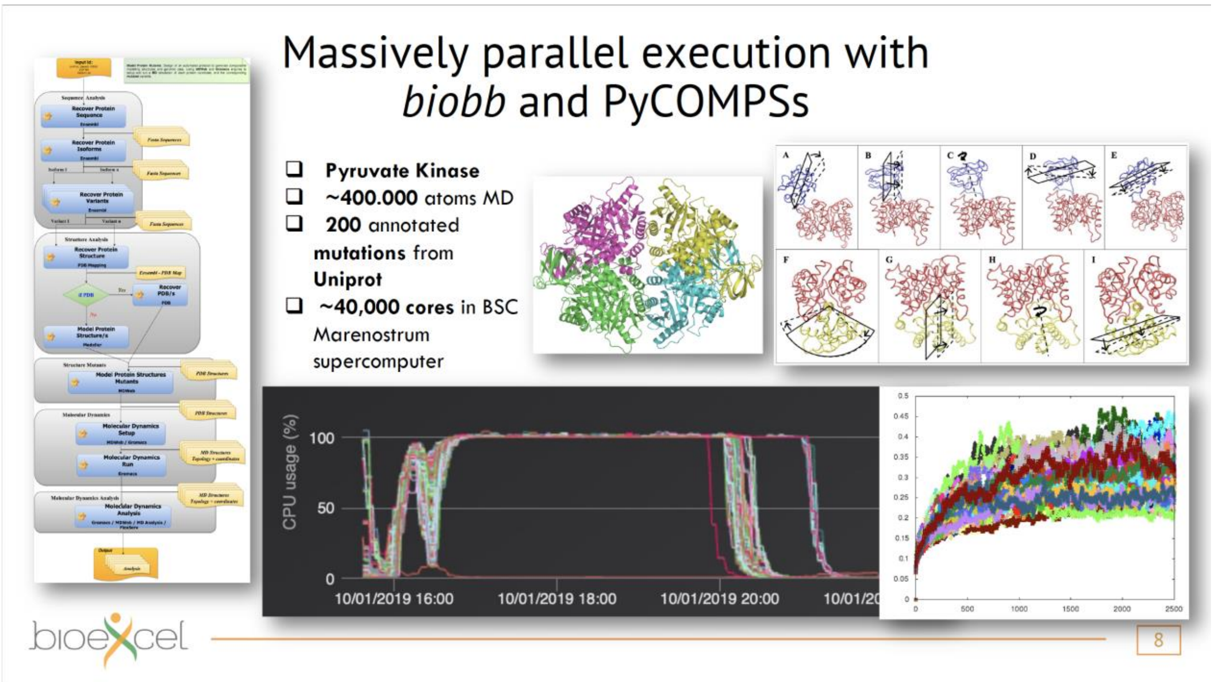

Figure 13 A use case (Pyruvate Kinase) from the BioExcel CoE implementing a massively parallel execution with biobb and PyCOMPSs (Python binding to COMPSs). PyCOMPSs exploits parallelism at runtime and can execute HPDA apps in distributed computing platforms (HPC clusters, cloud, edge computing).

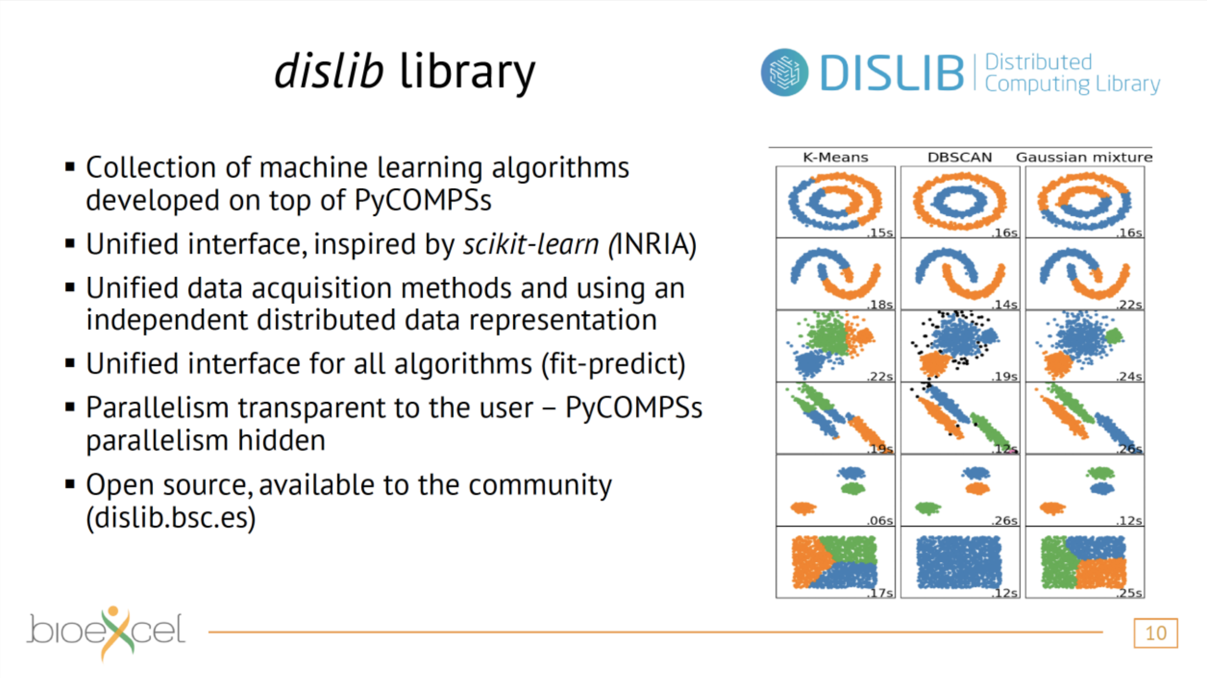

Figure 14 The dislib library in the BioExcel CoE provides a collection of machine learning algorithms built on top of PyCOMPSs. It provides a unified interface for all algorithms and it transparently enables parallelism, thus making easier for end users the development of HPDA applications/workflows.

5 Summary of FETHPC Presentations

In the next sub-sections, the HPDA contributions from the following FETHPC projects are presented: MAESTRO, ASPIDE, VESTEC, ExaNeSt, and Sage2.

5.1 Maestro: Towards a Memory- and Data-aware Middleware

Speaker: Dirk Pleiter (Jülich Supercomputing Centre) - MAESTRO Project

Abstract

In this talk, we introduce the Maestro project, started in September 2018, and involving the following partners: Appentra, CEA, Cray, CSCS, ECMWF, Forschungszentrum Juelich, Seagate. The goal is to develop and demonstrate a new middleware that aims at addressing two key shortcomings observed for today's supercomputing solutions. First, available solutions typically lack data awareness as they are optimised for filling processing pipelines, not for managing data. Furthermore, available solutions lack memory awareness, i.e. available middleware do not have information about available memory (or storage) hardware and thus are not able to make data transport decisions. Maestro applies a co-design approach for designing a solution addressing these challenges. The co-design applications, in particular, include data-intensive applications like the weather prediction workflow from ECMWF. In this case, simulations of the IFS model produce huge amounts of data that are consumed by other applications, which generate products for ECMWF's products member states and customers. This matches well Maestro's concept of providing a pool of data objects, where data producers give objects to the pool and data consumers take objects from the pool. Using suitable metadata, the Maestro middleware can manage storing and transport of data objects taking the hardware

architecture into account, on which the applications are being executed. We will conclude this talk with an outlook on the work of the Maestro project.

Talk highlights

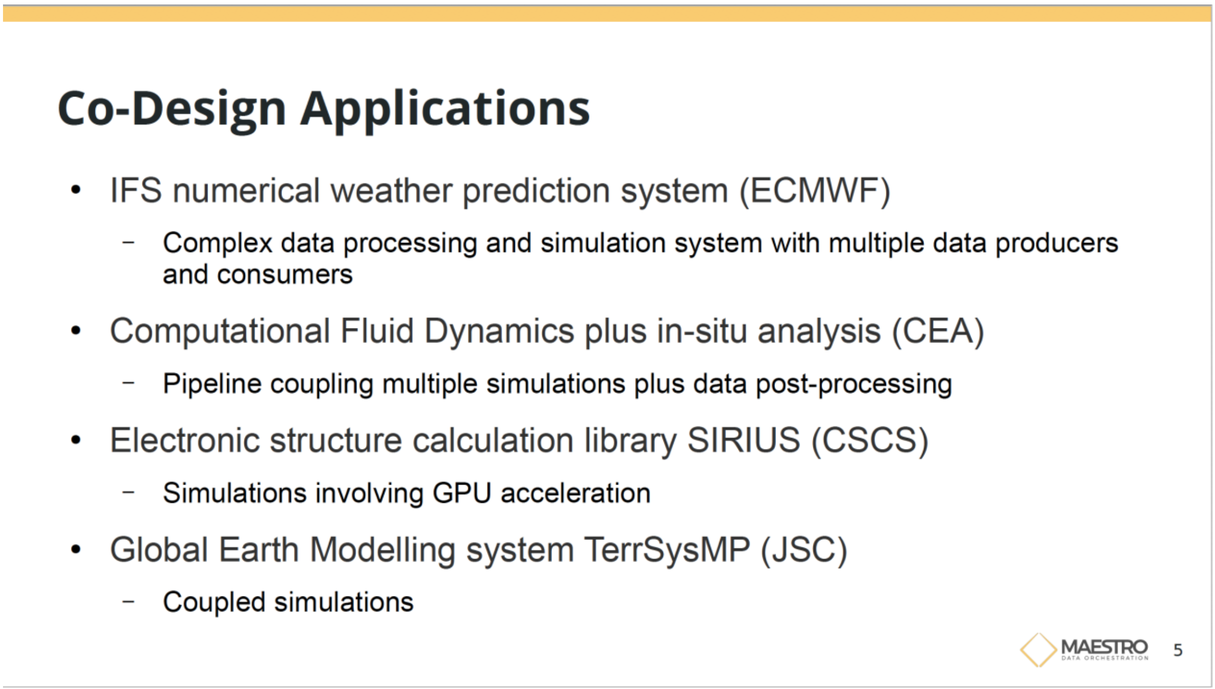

Figure 15 The Co-design applications in the MAESTRO FETHPC project. Apps span from numerical weather prediction system (IFS model at ECMWF), the computational fluid dynamics joining in-situ analysis at CEA, the electronic structure calculation library SIRIUS at CSCS and the global Earth modelling system, TerrSysMP at JSC. They all include HPDA challenges and requirements.

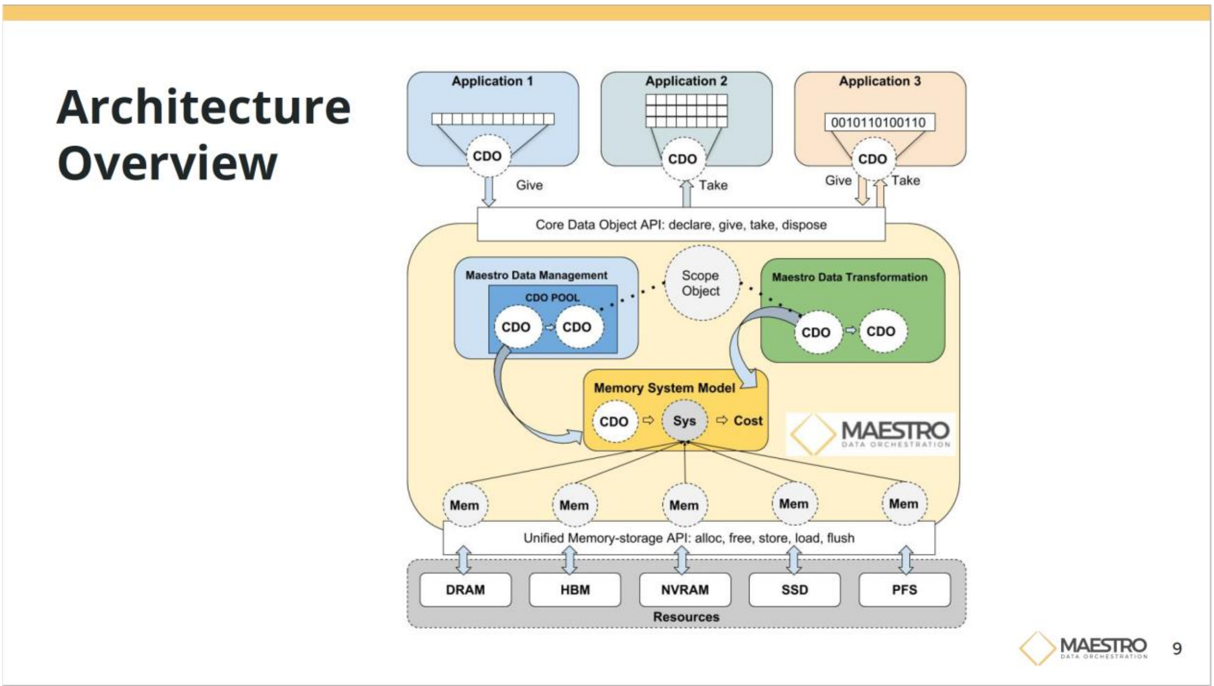

Figure 16 The architectural overview in the MAESTRO FETHPC project. A core element in the architecture is represented by the CDO, which are the Core Data Objects (object-based approach to encapsulate data with application and Maestro related metadata). The access to different types of memory is performed via a unified memory storage API.

5.2 Towards data intensive aware programming models for Exascale systems

Speaker: Francisco Javier Garcia-Blas (Universidad Carlos III de Madrid) ASPIDE Project

Abstract

Extreme Data is an incarnation of Big Data concept distinguished by the massive amounts of data that must be queried, communicated and analysed in (near) real-time by using a very large number of memory/storage elements and Exascale computing systems. Immediate examples are the scientific data produced at a rate of hundreds of gigabits-per-second that must be stored, filtered and analysed, the millions of images per day that must be mined (analysed) in parallel, the one billion of social data posts queried in real-time on an in-memory components database. Traditional disks or commercial storage cannot handle nowadays the extreme scale of such application data. Following the need of improvement of current concepts and technologies, ASPIDE's activities focus on data-intensive applications running on systems composed of up to millions of computing elements (Exascale systems). Practical results will include the methodology and software prototypes that will be designed and used to implement Exascale applications. The ASPIDE project will contribute with the definition of a new programming paradigms, APIs, runtime tools and methodologies for expressing data-intensive tasks on Exascale systems, which can pave the way for the exploitation of massive parallelism over a simplified model of the system architecture, promoting high performance and efficiency, and offering powerful operations and mechanisms for processing extreme data sources at high speed and/or real-time.

Talk highlights

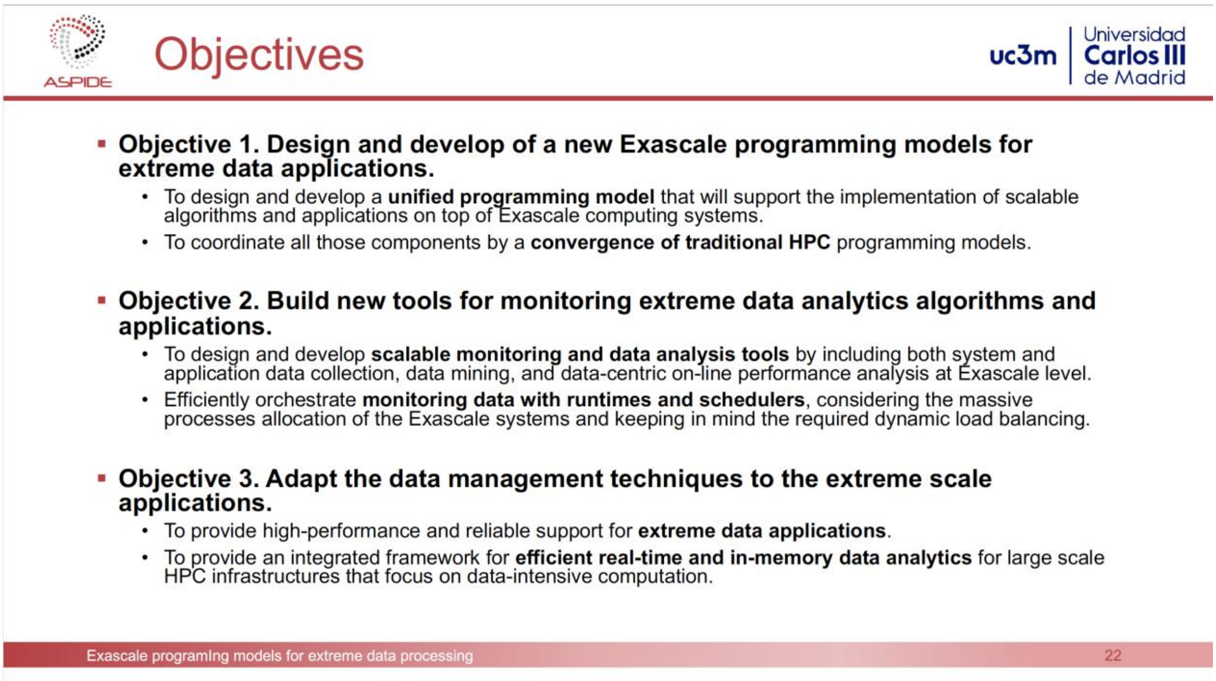

Figure 17 The three main objectives of the ASPIDE FETHPC project. HPDA is at the core of the project, design, and use cases.

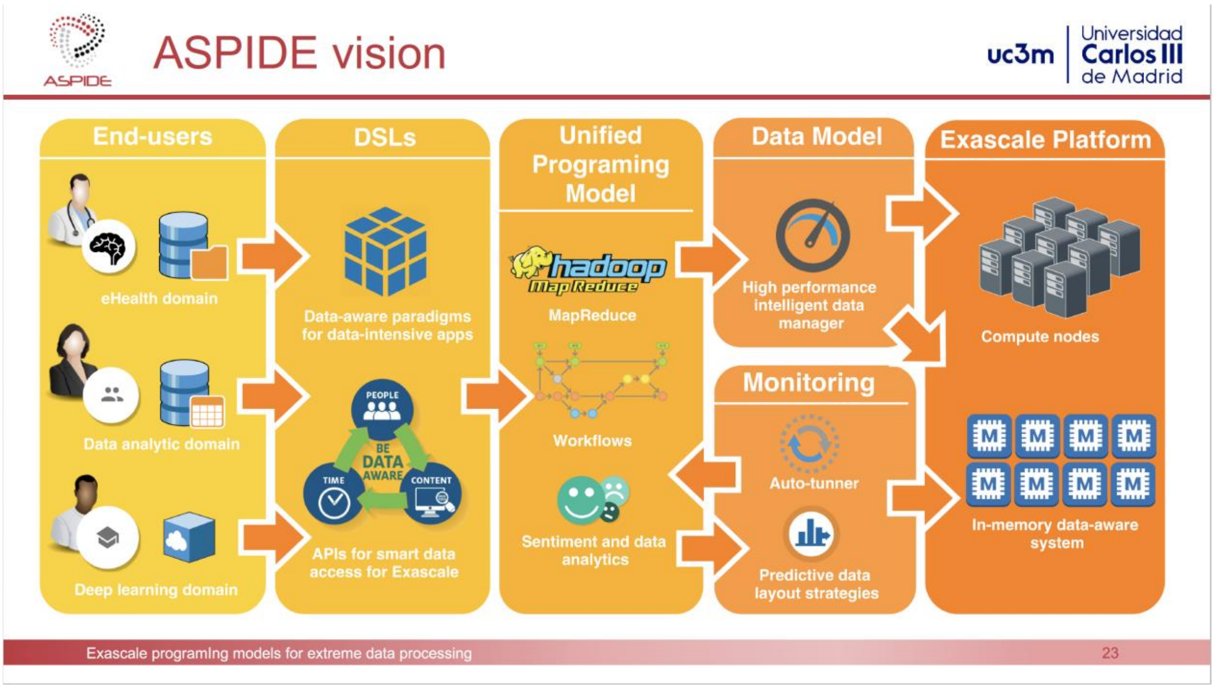

Figure 18 The ASPIDE vision. From end-users to exascale platforms through Domain Specific Languages (DSL), Unified programming models, data models and monitoring.

5.3 The VESTEC project: Fusing HPC and real-time sensor data for urgent decision making

Speaker: Gordon Gibb (EPCC, Univ. of Edinburgh) - VESTEC Project

Abstract

The VESTEC project, an acronym for Visual Exploration and Sampling Toolkit for Extreme Computing, aims to investigate how HPC resources can be used for disaster management/mitigation. Ongoing, rapidly evolving disasters (e.g. earthquakes, forest fires, disease outbreaks, etc.) can benefit from HPC simulations to help forecast their evolution to aid Crisis Managers in producing an appropriate response to them. As such events are rapidly evolving, real-time sensor data must be incorporated into any forecasts to ensure they are as up-to-date as possible. A typical workflow would be initiated when a disaster occurs. Sensor data would be collected, very possibly form a variety of different sources, and a series/ensemble of simulations would be run to simulate the disaster and determine uncertainty quantification. It is possible that initial simulations may be required to run initially to gain initial conditions or boundary conditions for the main simulations (e.g. a weather forecast to determine weather conditions). Results can then be forwarded to the Crisis Management Centre to help guide the relief effort. As the crisis continues, new sensor data will be incorporated into running simulations, or new simulations will be run to ensure the simulations remain up to date. Crucially, we envisage the capability for the Disaster Management Centre to provide feedback to steer running simulations, for example to explore the effects of possible mitigation actions on the evolution of the disaster. We envisage a central server that is responsible for monitoring and preparing new sensor data, managing simulations, and communicating with the Crisis Management Centre. This server will monitor a pool of HPC machines (ideally most HPC facilities in Europe) and try to submit jobs to the machine with the highest probability of running jobs quickly, to ensure as close to real-time execution as possible. As simulations can potentially produce GBs or even TBs of data, we plan to have in-situ data reduction/processing

so that only vital information from each simulation needs to be transferred to the Disaster Management Centre. The VESTEC project hence has to deal with all four Vs of Big Data: volume, variety, veracity and velocity. The volume comes from the potential size of output data from simulations that needs to be reduced, whilst the variety comes from the different sources of sensor data, as well as results from different simulations. The veracity of the data is included in uncertainties from sensor data, and from ensembles being used for uncertainty quantification. Finally, the velocity of the data comes from the rates at which new data may be available from sensors. Although still in the early stages, we think that this is a very exciting research project and we relish the challenges involved in getting this vision to completion. The VESTEC project has received funding from the European Union's Horizon 2020 Programme for research, technological development and demonstration under grant agreement n° 800904.

Talk highlights

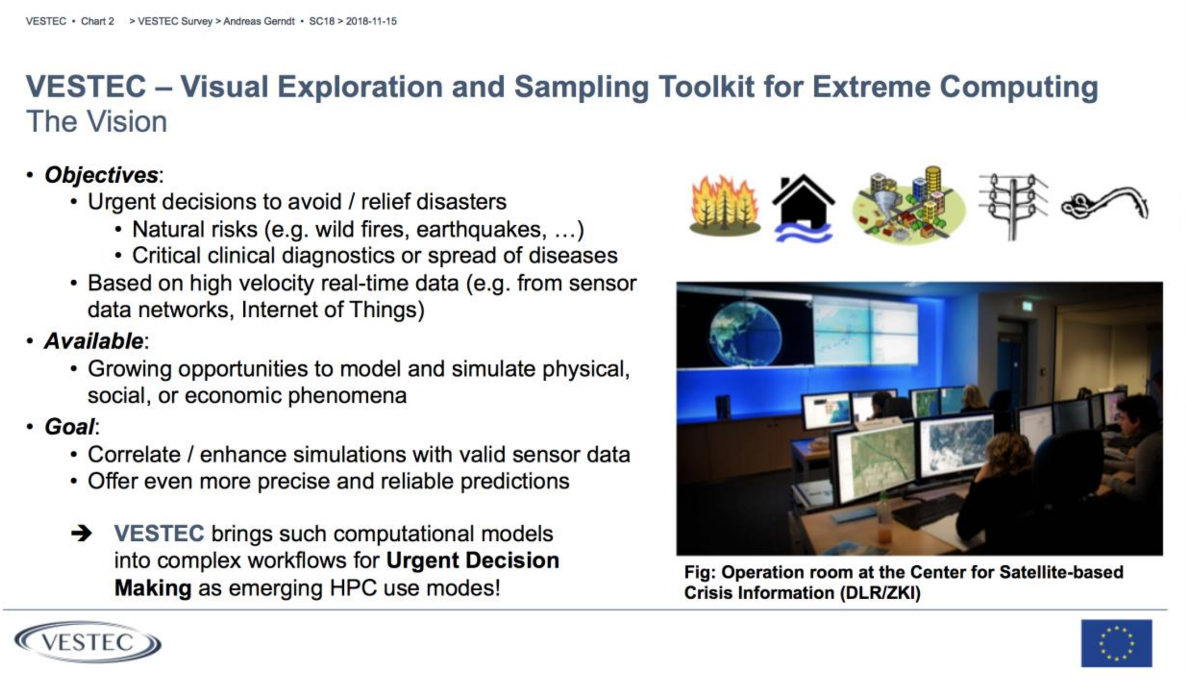

Figure 19 The VESTEC FETHPC project vision bringing together HPC and real-time data to enable urgent decisions to be made for disaster response. Whilst some other approaches have looked at HPC for urgent decision making, they did not leverage real-time data, which adds extra complexity to the problem as ensemble simulations been to be updated/run on the fly. There is a strong need to present correct information in a way that the urgent responder can make the correct decision first time, every time.

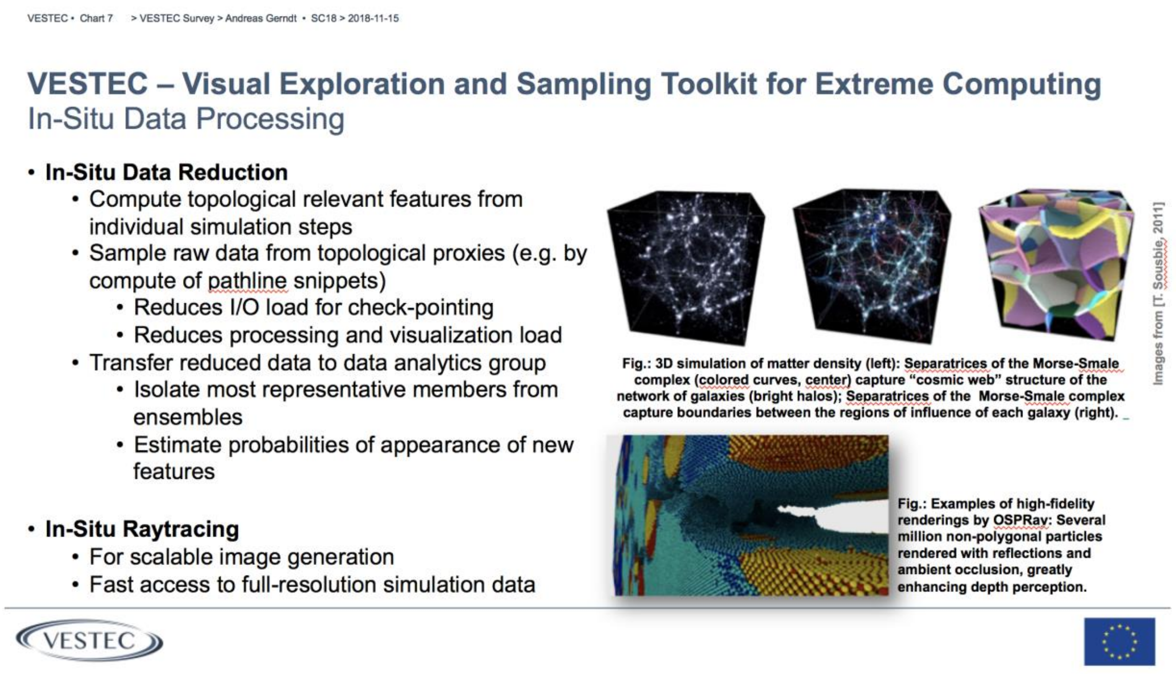

Figure 20 Another critical aspect faced by the VESTEC FETHPC project is the data analysis and visualization at extreme-scale. Several techniques for in-situ data reduction and in-situ raytracing are investigated and developed in the project.

5.4 ExaNeSt: Low-Latency Communication and Acceleration in a liquid-cooled energy-efficient Prototype Rack

Speaker: Manolis G.H. Katevenis (Univ. of Crete) - ExaNeSt Project

Abstract

ExaNeSt, ExaNoDe, and EcoScale, three "sister" projects, as well as EuroEXA, their follow-up project, work for advancing the state-of-the-art in specific technologies that are needed for exascale computing over an energy-efficient platform based on 64-bit ARM processors and Reconfigurable Accelerators (FPGAs). ExaNeSt focuses on dense packaging, the interconnection network, storage and data access, and also works with a rich set of real, full HPC Applications that have been ported, optimized, and evaluated on the ARM platform. ExaNeSt has built an HPC Testbed consisting of very dense "QFDB" daughter cards plugged onto Mezzanine boards and immersed into liquid-cooled blades; this now contains 96 FPGAs, 1.5 TBytes of DRAM, and 6 TBytes of non-volatile solid-state storage, and runs a full stack of systems software and libraries, HPC job manager, and our rich set of Applications, while its size will soon be doubled. The nodes are interconnected via a 3-Dimensional Torus network with 10 Gbps links and 185 ns per-hop latency. The virtualized network interfaces, integrated on-chip with the processing cores, feature a zero-copy, protected, user-level RDMA engine with 1024 channels, packetizers, mailbox queues, and offer resilient communication at half-microsecond one-way, one-hop, user-to-user software latency. Global storage, with per-job SSD/NVM on-demand temporary parallel file system, is provided by BeeGFS with replication extensions; low-latency memory-mapped storage access path, and Virtual Machine (VM) support have been added. Selected real, full Applications from materials science, climate forecasting, computational fluid dynamics, astrophysics, neuroscience, and an open database suitable for large-scale data analytics have been ported and optimized for this platform. Evaluation is on-going and shows a significant advantage over Intel-based platforms, ranging from 30% to half, and sometimes one full order of magnitude reduction in Energy-to-Solution.

The ARM processors that have been used out of necessity (A53, the only ones available in FPGAs in 2016) only offer a very modest floating-point performance, so the time-to-solution is quite longer than on Intel platforms, but the full performance advantages, while retaining energy-efficiency, will appear with next generations of ARM processors. Additional evaluations of FPGA-accelerated applications, in collaboration with the EcoScale and EuroEXA projects, have demonstrated significant energy advantages at performances that are competitive to Intel and Nvidia GPU platforms.

Talk highlights

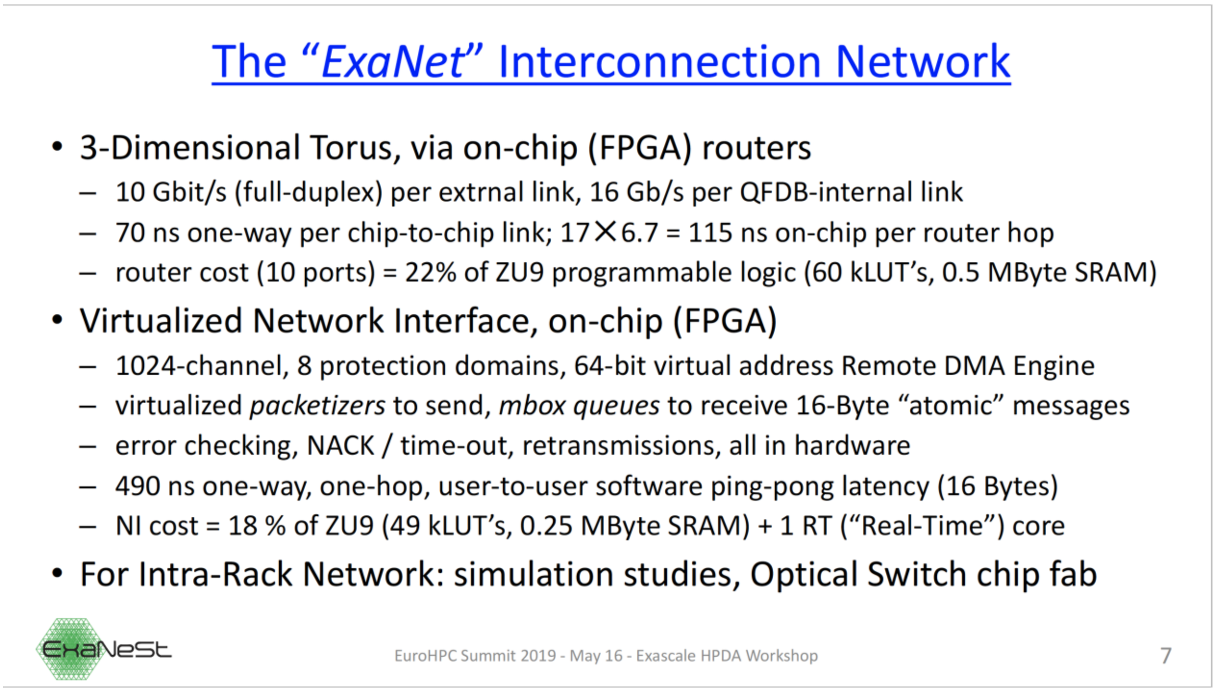

Figure 21 The physical details about the ExaNeSt interconnection network.

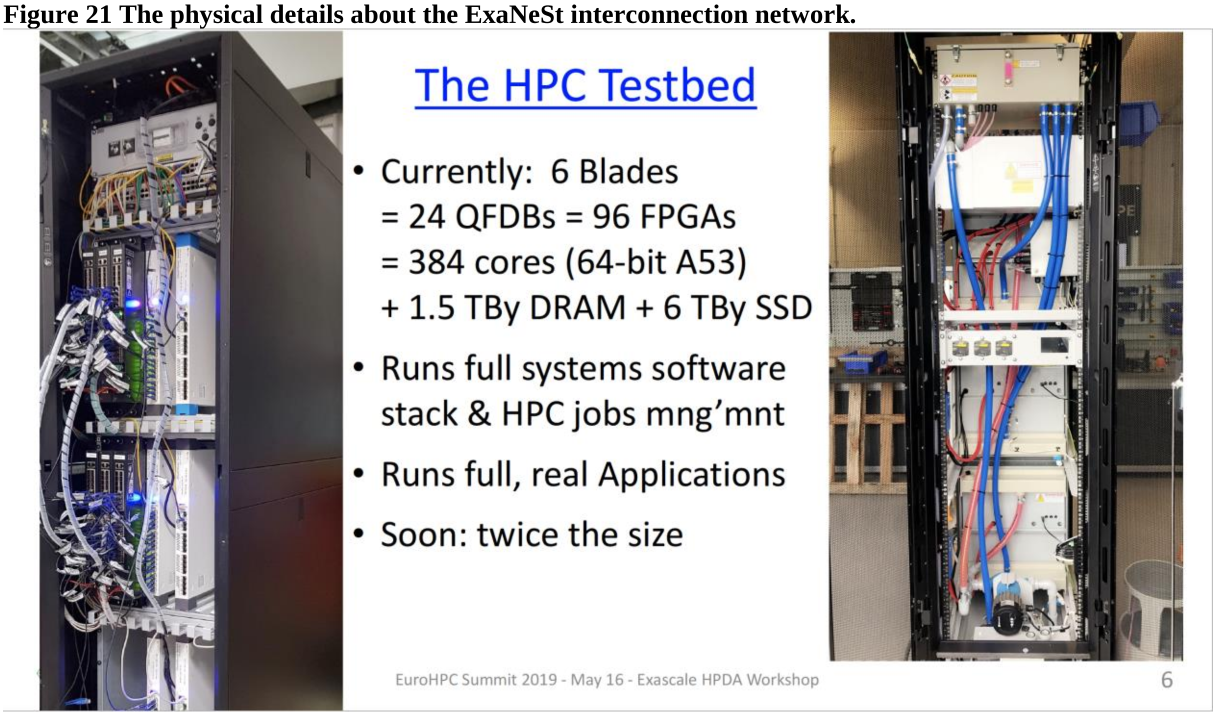

Figure 22 The HPC testbed in the ExaNeSt FETHPC project. The hardware has been tested on real applications and it is able to run full systems software stack as well as HPC jobs. The project is going to deliver soon twice the current size of this testbed. Data-intensive use cases have been considered too.

5.5 Sage2: Architecting a storage platform for the extreme data era

Speaker: Sai Narasimhamurthy (Seagate Technology, LLC) Sage2 Project

Abstract

The landscape for High Performance Computing is changing with the proliferation of enormous volumes of data created by scientific instruments and sensors, in addition to data from simulations. This data needs to be stored, processed and analysed, and existing storage system technologies in the realm of extreme computing need to be adapted to achieve reasonable efficiency in achieving higher scientific throughput. We started on the journey to address this problem with the SAGE project. The HPC use cases and the technology ecosystem are now further evolving and there are new requirements and innovations that are brought to the forefront. It is extremely critical to address them today without “reinventing the wheel” leveraging existing initiatives and know-how to build the pieces of the Exascale puzzle as quickly and efficiently as we can. The SAGE paradigm already provides a basic framework to address the extreme scale data aspects of High Performance Computing on the path to Exascale. Sage2 (Percipient StorAGe for Exascale Data Centric Computing 2) intends to validate a next generation storage system building on top of the already existing SAGE platform to address new use case requirements in the areas of extreme scale computing scientific workflows and AI/deep learning leveraging the latest developments in storage infrastructure software and storage technology ecosystem. Sage2 aims to provide significantly enhanced scientific throughput, improved scalability, and, time & energy to solution for the use cases at scale. Sage2 will also dramatically increase the productivity of developers and users of these systems.

Talk highlights

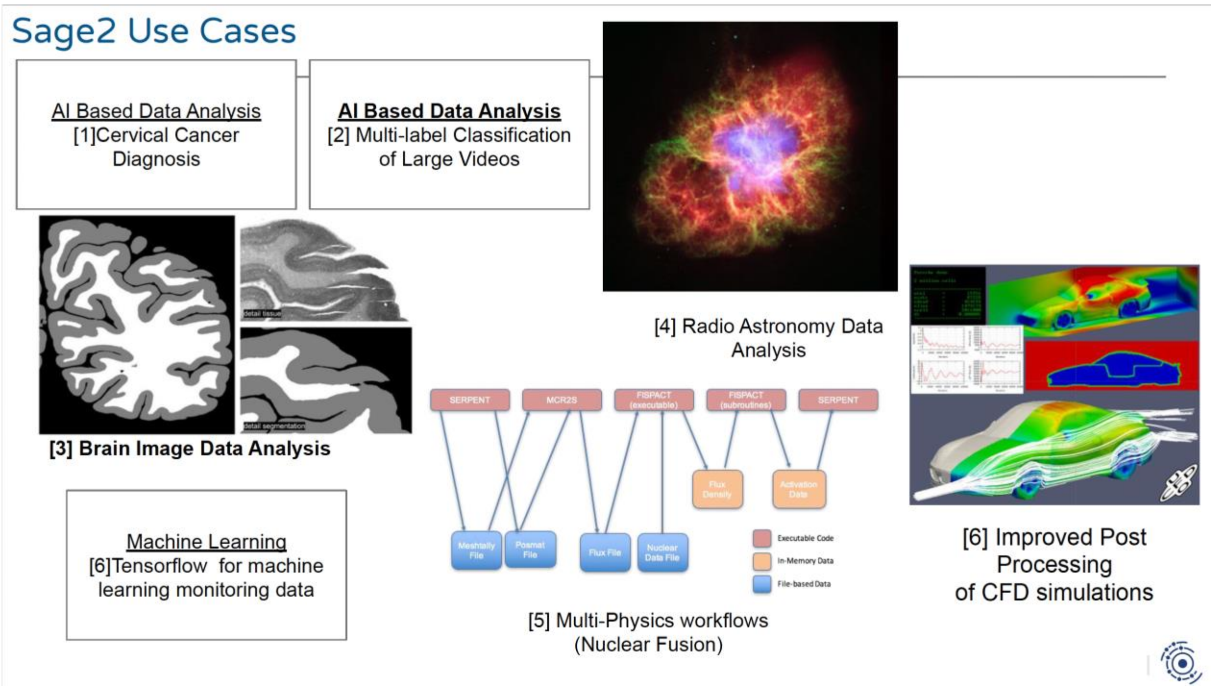

Figure 23 The complete set of use cases in the Sage2 FETHPC project. HPDA, ML and AI are core parts of the proposed applications. The use cases span from cervical cancer diagnosis to radio astronomy data analysis to improved post processing of Computational Fluid Dynamics (CFD) simulations.

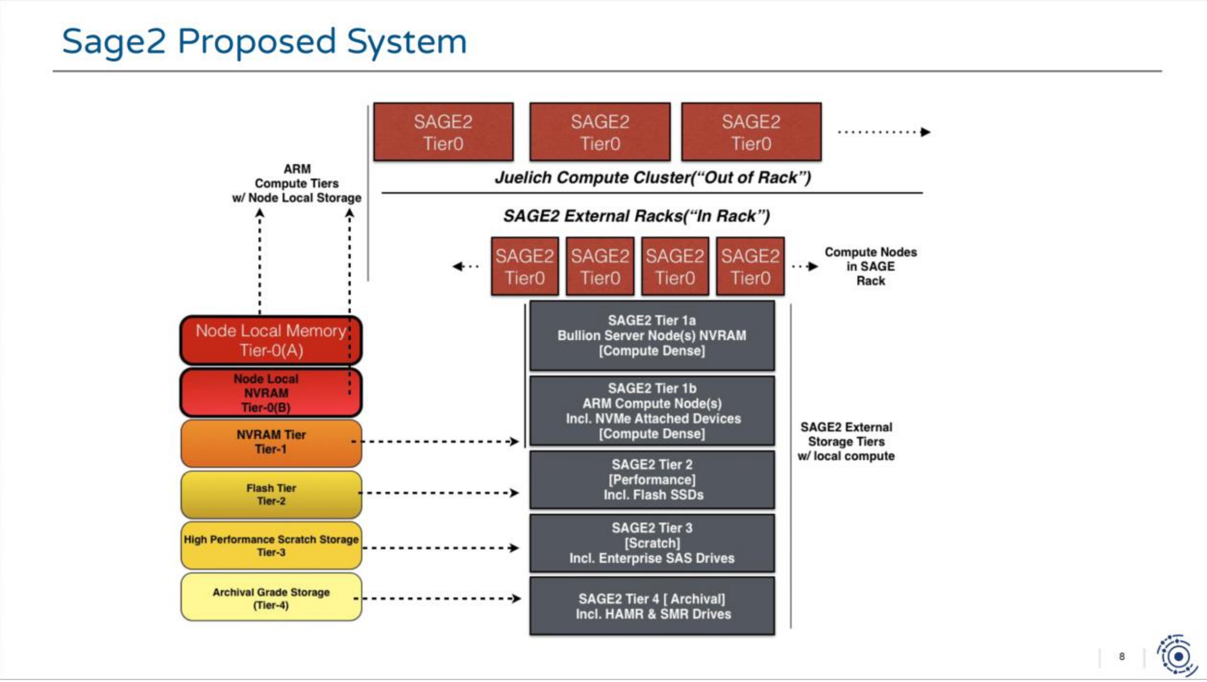

Figure 24 The Sage2 proposed system. A full stack of storage tiers provides the "SAGE2 External Rack" (In Rack). The compute tier hosts in-storage analytical capabilities which allow to perform analysis, reduce data and data movement with the compute cluster ("Out of Rack").

6 Conclusions and Next Steps

This deliverable summarizes the most relevant information about the EXDCI-2-WP3-organized "*1st European Communities Workshop on Exascale Computing*" with a focus this year on "*High Performance Data Analytics*" held during the EuroHPC Summit Week 2019.

The workshop was organized in two sessions, one dedicated to the CoEs and the other one to the FETHPC projects, for a total of 11 presentations. The presentations covered various aspects of HPDA from Big Data processing and parallel data analysis, heterogeneous infrastructure to deal with different aspects of HPDA, to provenance and reproducibility of simulation results.

The discussions and questions showed that many European communities are facing similar challenges and would be thankful for more joint efforts and combined solutions as well as networking activities like this workshop to ease experiences, results, solutions and best practices sharing.

In order to summarize main challenges with respect to HPDA for European communities, similarities within the talks were identified:

1) Hardware: Analyzing big data requires suitable hardware. For example, challenges in the area of I/O and storage need to be overcome and analytic tools must be integrated into the existing work flows.
2) Data transfer: Moving data is expensive. Thus, integrated work flows with an efficient use of data are important. Several in-situ solutions are currently under development or already in use.
3) Interface / integration: Many statistical or machine learning tools exist but the integration of these tools into existing work flows is challenging. Several CoEs work on interfaces making the integration of these tools simpler.

4) Reproducibility / quality of data: Big data analytics requires data with high quality. Ensuring this is a challenge if data origin, for example, from simulations, which cannot be reproduced. Systematic approaches, which will increase data reproducibility, are currently under development.
5) Real time information: The availability of real time information from sensors etc. has increased over the last years. The integration of this information, for example, as edge to cloud application raises new challenges, which need to be tackled. Examples are latency issues and data privacy.

These findings were shared and discussed with the workshop participants after the workshop in order to help them to synchronize developments and use synergies. Furthermore, discussions with the BDEC (Big Data and Extreme-scale Computing) Application Group were started to get an international view on this topic.

For WP3, it was important to get feedback with respect to this emerging topic, which helps to produce a more relevant HPC roadmap. Overall the workshop was successful, with about 50 people attending each of the two sessions. It is planned to organize a *"2nd European Communities Workshop on Exascale Computing"* during the next EuroHPC Summit Week in 2020, with a different focus, most likely (based on preliminary thoughts) on Artificial Intelligence (AI) and Machine Learning (ML) topics.