Ref. Ares(2023)4412902 - 26/06/2023

Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

WP8 Project Management and Exploitation

D8.4 External Advisory Board report

This project has received funding from the European Union's Horizon 2020 research and innovation programme, EuroHPC JU, grant agreement No 956831

TEXTAROSSA

Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

Grant Agreement No.: 956831

Deliverable: D8.4 External Advisory Board reports

Project Start Date: 01/04/2021 **Duration**: 36 months
Coordinator: *AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE, L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE - ENEA, Italy.*

Deliverable No	D8.4
WP No:	WP8
WP Leader:	ENEA
Due date:	M24 (March 31, 2023)
Delivery date:	21/05/2023

Dissemination Level:

PU	Public	
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	X

This project has received funding from the European Union's Horizon 2020 research and innovation programme, EuroHPC JU, grant agreement No 956831

DOCUMENT SUMMARY INFORMATION

Project title:	Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale
Short project name:	TEXTAROSSA
Project No:	956831
Call Identifier:	H2020-JTI-EuroHPC-2019-1
Unit:	EuroHPC
Type of Action:	EuroHPC - Research and Innovation Action (RIA)
Start date of the project:	01/04/2021
Duration of the project:	36 months
Project website:	textarossa.eu

WP8 Project Management and Exploitation

Deliverable number:	D8.4					
Deliverable title:	External Advisory Board reports					
Due date:	M24					
Actual submission date:	M27					
Editor:	M. Celino					
Authors:	W. Fornaciari					
Work package:	8					
Dissemination Level:	Confidential					
No. pages:	25					
Authorized (date):	14/04/2023					
Responsible person:	William Fornaciari					
Status:	Plan	Draft	Working	Final	Submitted	Approved

Revision history:

Version	Date	Author	Comment
0.1	30.03.2023	Massimo Celino	Draft structure
1.0	21.05.2023	Massimo Celino	Final check and submission on ECAS

Quality Control:

Checking process	Who	Date
Checked by internal reviewer	Carlos Alvarez (BSC), W Fornaciari (CINI)	28.04.2023
Checked by Task Leader	Massimo Celino (ENEA)	30.04.2023
Checked by WP Leader	Massimo Celino (ENEA)	30.04.2023
Checked by Project Coordinator	Massimo Celino (ENEA)	20.05.2023

COPYRIGHT

© Copyright by the **TEXTAROSSA** consortium, 2021-2024
This document contains material, which is the copyright of TEXTAROSSA consortium members and the European Commission, and may not be reproduced or copied without permission, except as mandated by the European Commission Grant Agreement No. 956831 for reviewing and dissemination purposes.

ACKNOWLEDGEMENTS

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement no 956831. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Italy, Germany, France, Spain, Poland.
Please see http://textarossa.eu for more information on the TEXTAROSSA project.
The partners in the project are AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE, L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE (ENEA), FRAUNHOFER GESELLSCHAFT ZUR FOERDERUNG DER ANGEWANDTEN FORSCHUNG E.V. (FHG), CONSORZIO INTERUNIVERSITARIO NAZIONALE PER L'INFORMATICA (CINI), INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), BULL SAS (BULL), E4 COMPUTER ENGINEERING SPA (E4), BARCELONA SUPERCOMPUTING CENTER-CENTRO NACIONAL DE SUPERCOMPUTACION (BSC), INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK (PSNC), ISTITUTO NAZIONALE DI FISICA NUCLEARE (INFN), CONSIGLIO NAZIONALE DELLE RICERCHE (CNR), IN QUATTRO SRL (in4). Linked third parties of CINI are POLITECNICO DI MILANO (CINI-POLIMI), Università di Torino (CINI-UNITO) and Università di Pisa (CINI-UNIPI); linked third party of INRIA is Université de Bordeaux; in-kind third party of ENEA is Consorzio CINECA (CINECA); in-kind third party of BSC is Universitat Politècnica de Catalunya (UPC).
The content of this document is the result of extensive discussions within the TEXTAROSSA © Consortium as a whole.

DISCLAIMER

The content of the publication herein is the sole responsibility of the publishers and it does not necessarily represent the views expressed by the European Commission or its services. The information contained in this document is provided by the copyright holders "as is" and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the members of the TEXTAROSSA collaboration, including the copyright holders, or the European Commission be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of the information contained in this document, even if advised of the possibility of such damage.

Table of contents

List of Acronyms....5
Executive Summary....7
1 Introduction....8
2 TEXTAROSSA objectives....9
3 State of the art in HPC....10
3.1 HPC or edge computing....10
3.2 New IPs for heterogeneous computing....11
3.3 Cybersecurity....11
3.4 Heterogeneity....12
3.5 Energy consumption....13
4 External Advisory Board....14
5 EAB Member: Dr. R. Squarcini....16
5.1 Curriculum Vitae....16
5.2 Company....17
6 EAB Member: Dott. S. Cozzini....18
6.1 Curriculum Vitae....18
6.2 Company....19
7 EAB vision....20
8 Lesson learnt: IT knowledge best practices....23
9 Conclusions and future works....25

List of Acronyms

Acronym	**Definition**

EAB	External Advisory Board
PTC	Project Technical Committee
HBM	High-Bandwidth Memory
CXL	Compute Express Link
UCIe	Universal Chiplet Interconnect Express
OPS	Oracle Certified Professional
HEP	High Energy Physics
PC	Project Coordinator
PTM	Project Technical Manager
HPDA	High Performance Data Analytics
HPC-AI	High Performance Artificial Intelligence
FPGA	Field Programmable Gate Array

Executive Summary

In a rapidly changing HPC environment of increasing impact in society at all levels, it is important to have a keen and critical eye on the innovations and advances occurring in the technology sector. At the same pace, industry contexts are changing, and new professions are now available. For these reasons, it is crucial to have the support of experts from outside the project who can have a critical look toward the project. Of particular importance is having selected members of the External Advisory Board (EAB) who have strong experience in industry, both in energy management and new HPC architectures.

To support the work of the partners of consortium and in particular the activities of the Project Technical Committee (PTC), an EAB was established. The task of the External Advisory Board is to monitor the adherence of project results to industry needs and to provide indications to facilitate rapid market adoption. The TEXTAROSSA EAB consists of internationally recognized experts in the field of energy management and cooling systems as well as experts in industrial processes and HPC and Big Data technologies.

The EAB will support the project by exchanging knowledge on best practices and experiences in implementing and using technological developments. The EAB will advise the consortium on project progress and scientific direction, advise on the policy relevance of project results, and facilitate industry and end-user participation to identify and pursue exploitation strategies.

It is worth to underline that the availability of the TEXTAROSSA EAB helped to better focus project deficiencies during the analysis of the project issues after the review report at M18. Moreover, the EAB provided a clearer vision of the TEXTAROSSA project indicating the more impactful results and achievements. This caused a delay in the submission of this deliverable, but it was the way to move forward and synthonize the project activities towards the current trends in future HPC.

1 Introduction

The TEXTAROSSA project has the ambition to bring significant progress in the HPC sector. The project addresses the HPC sector at multiple levels, from hardware to use cases via software systems for the management and efficient use of computing nodes. Furthermore, the project proposes to develop new cooling technologies at both chip and node level. The main target is to provide innovative technologies to support the development of future HPC infrastructures. However, several hardware and software technologies developed in TEXTAROSSA can be implemented in the edge computing infrastructures.

The problem is that the HPC world is evolving very rapidly also thanks to important investments, increasingly targeted and ambitious in recent years. In this context it is necessary to develop a detailed strategy to rapidly develop new solutions for HPC but consider the developments that will soon occur in these years. To this end it is essential that the project has a constant and continuous dialogue with experts in the sector to monitor the progress of the project with respect to what is evolving externally, both at the European and international level.

Various scenarios are opening that could greatly influence the development of HPC. Among the others, TEXTAROSSA is focusing edge computing developments because are going to influence the HPC roadmap. Another extremely interesting trend is the design of new computing nodes by inserting multiple units, each specialized in specific tasks. While this increases flexibility, it requires the development of new technologies and IPs.

Finally, it has become important to consider the problems related to the heterogeneity of systems, performance, data processing and the energy issue which seems to be increasingly prevalent. The project takes some of these into account and is striving to harmonize them to create state-of-the-art prototypes in the HPC sector. In this context it is important to have a continuous and punctual dialogue with experts external to the project. Since the project is focused on the HPC sector and its technologies taking into strong consideration the consumptions and management, a precise choice has been made regarding the components of the advisory board.

Compatible with the resources available to the project and the need to have an advisory board highly responsive in analyzing project issues, it was decided to limit the advisory board to two high-level members who are highly focused on the issues of interest to the project.

The report is organized as follows.

- The project objectives are summarized.
- An analysis of the new trends in the HPC sector is made to report open discussions at the international level that are of interest to TEXTAROSSA.
- Definition of EAB and its role within the project.
- Presentation of the the advisory board members.
- EAB vision about the future HPC trends
- Lessons learnt about project knowledge management.
- Conclusions.

2 TEXTAROSSA objectives

TEXTAROSSA project has a major vision of the future trends in HPC. To achieve high performance and high energy efficiency on near-future exascale computing systems, a technology gap needs to be bridged: increase efficiency of computation with extreme efficiency in HW and new arithmetic, providing methods and tools for seamless integration of reconfigurable accelerators in heterogeneous HPC multi-node platforms, and designing and implementing new algorithms and application software to exploit at the best the heterogeneous platforms. TEXTAROSSA aims at tackling this gap through applying a co-design approach to heterogeneous HPC solutions, supported by the integration and extension of IPs, programming models and tools derived from European research projects, led by TEXTAROSSA partners. The main directions for innovation are towards: enabling mixed-precision computing, through the definition of IPs, libraries, and compilers supporting novel data types (including Posits), used also to boost the performance of AI accelerators; implementing new multilevel thermal management and two-phase liquid cooling; developing improved data movement and storage tools through compression; providing IPs for fast task scheduling and low-latency intra/inter-node communication.

These technologies will be tested on the Integrated Development Vehicles mirroring and extending the European Processor Initiative ARM64-based architecture, and on an OpenSequana testbed. To drive the technology development and assess the impact of the proposed innovations TEXTAROSSA will re-design and implement a selected but representative number of HPC, HPDA and AI demonstrators covering challenging HPC domains such as general-purpose numerical kernels, High Energy Physics (HEP), Oil & Gas, climate modelling, and emerging domains such as High Performance Data Analytics (HPDA) and High Performance Artificial Intelligence (HPC-AI).

Given the objectives of TEXTAROSSA and the new technologies that are under development it is very important to have a clear vision of the current development in the European HPC ecosystem:

- HPC or edge computing
- New Ips for heterogeneous computing
- Cybersecurity that needs fast secure communications and accelerators for data encryption
- Heterogeneity of systems
- Energy consumption

All these topics are analyzed in the following sections.

3 State of the art in HPC

3.1 HPC or edge computing

In recent years, the COVID-19 pandemic has had a significant impact on several sectors of the economy and society, including the high-performance computing (HPC) industry and data centers. Despite these challenges, the demand for more computing power and reduced latency has driven the growth of the HPC market beyond the initial forecasts made in the early 2020s.

The contemporary landscape is characterized by new work methodologies, educational organization and social interaction patterns. This change requires greater flexibility and advanced data processing and sharing capabilities, resulting in increased productivity, information accessibility, and collaboration efficiency. As a result, HPC is no longer limited to big data centers, research labs, and supercomputers. It is now being used in a variety of industries for such tasks as product design, financial modeling, and weather forecasting. HPC's ability to provide increasingly powerful computing capabilities has made it an integral part of industries and applications closely related to everyday life. The HPC market is expected to continue to evolve and expand, with an increasing number of industries requiring interconnected computing systems and high-speed networks.

The need for HPC systems to predict catastrophic weather events, accelerate energy transitions, and address post-COVID medical emergencies has become increasingly apparent. Furthermore, the emergence of HPC for consumer-oriented applications is gaining momentum due to the availability of HPC in the cloud. Historically, high-performance data centers were isolated and accessible only to research organizations, governments, and businesses with substantial budgets. It is anticipated that HPC at the edge will become more prevalent than the exception. The industrial sector will utilize HPC for applications in robotics, vision systems, preventive maintenance, and monitoring, such as predetermined or predictive failures on assembly lines, essentially encompassing all industrial areas requiring computing power precisely where time-saving devices are employed during downtime.

The main driver for the growth and expansion of HPC is the growing importance of data. Data must be processed, analyzed and transferred faster than ever before. As a result, the HPC market is expanding its scope to include new types of work, integrating artificial intelligence (AI) and data analysis with traditional simulation and modeling. This, coupled with the growing demand for faster data processing and higher levels of accuracy in various vertical industries (life sciences, automotive, financial, gaming, manufacturing, aerospace, etc.), has led to a growing demand for more data processing. Technologies such as artificial intelligence, edge computing, 5G and Wi-Fi 6 will enhance the capabilities of HPC, leading to the development of new chip/system architectures that can provide industries with advanced processing and analysis capabilities.

Without a doubt, edge computing is a key trend that is transforming the overall computing landscape. The main challenge lies in its apparent opposition to HPC, which typically relies on large, centralized processing and storage resources. In contrast, edge computing focuses on processing data at or near the edge of the network, rather than transmitting it to a central location. This approach can offer reduced latency and, in many cases, more secure operational characteristics. However, these domains are converging: edge computing increasingly incorporates HPC outside traditional data centers. The volume and complexity of data generated in the edge is growing exponentially, encompassing a wide

range of Internet of Things (IoT) devices driven by demand for intelligent systems. Edge computing is particularly important for latency-sensitive and high-content applications, where round trips to a cloud/data center do not meet the required response time. In some cases, files are too large to be transmitted to the cloud for processing or storage. Examples of such applications are urban traffic management, automated driving systems, precision medicine, fraud detection, business intelligence, and smart city development. As a result, edge computing will have a significant impact on HPC system providers, cloud, network and storage service providers as organizations strive to integrate remote HPC capabilities with locally generated and processed data strategies. In this context, the physical footprint of HPC is expected to expand from a centralized delivery model to a more distributed approach, including locations near data-intensive edge points.

3.2 New IPs for heterogeneous computing

From a chip design perspective, edge computing requires optimal power, performance, and area (PPA), with a key priority: reducing latency in processing and data transfer. Design strategies must prioritize the speed and efficiency of data transfer in integrated circuits, such as those using chiplet architectures. A chip design solution must consider all aspects of PPA tradeoff scenarios and provide advanced capabilities for designing and analyzing an integrated circuit optimized for any application requirement. This includes powerful simulation and verification tools, thermal and power analysis capabilities, intelligent implementation of design layouts, and a range of certified IP blocks for key functions and interfaces. In the future, there will be a growing demand for energy-efficient design solutions, from data centers to battery-powered IoT devices. One of the latest trends in HPC is the use of multi-die systems. Traditional monolithic systems-on-chip (SoCs) are becoming too large and expensive to produce for advanced designs, and performance risk is growing along with design size. The multi-die approach is attractive in that it is a viable way to extend the benefits of Moore's Law in terms of PPA, offering increased processing capacity without requiring an increase in chip area or power. It also supports a heterogeneous mix-and-match approach to maximize application-optimized process technologies. Decomposing SoC components, manufacturing them separately, and then bundling these distinct functions into a single package reduces waste and enables the rapid creation of new product variants with optimized power and system performance. However, the multi-die approach presents heterogeneous integration, interconnection, and packaging problems, requiring advanced IP, methodologies, and tools to address these issues.

3.3 Cybersecurity

The importance of protecting information, data confidentiality and integrity, and providing data access controls has increased significantly. As the value of data in the infrastructure increases, the number of cyber attacks is expected to increase, making security from hardware to all levels of the stack increasingly important to protect information. Fast and secure communications will be needed, as well as accelerators, of which FPGAs are the most promising for data encryption.

To access data, people must validate their identity and prove that they are authorized to access it. This requirement is expected to intensify further in the coming years. Hardware is already being built to support this, and we will see trust roots embedded in every element of an infrastructure. This will allow them to authenticate each other and ensure that before data is shared with another device, that device is authorized to use and process it.

As digitization increases in many markets, so do the opportunities for security risks. As processing performance moves away from the data center, the number of opportunities for attacks that cannot be fully mitigated with software patches will increase. This will put pressure on design teams to rush hardware to solve these problems, resulting in accelerated hardware design cycles. Increasing the productivity of designers to keep up with time-to-market demands will be a critical need.

As the amount of data increases, it is not only security that needs to be considered. Storage infrastructure will have to increase, as will data processing capacity. This will require a comprehensive approach to hardware design, including the use of accelerators and other specialized hardware to handle the increased workload. In addition, the hardware must be designed to be scalable and flexible to accommodate future growth and changing needs.

3.4 Heterogeneity

The architectural landscape of HPC is undergoing radical change, driven by evolving artificial intelligence workloads, processing flexibility (CPUs, GPUs, FPGAs, DPUs, etc.), cost, memory, and IO throughput. Advances in microarchitecture include faster interconnects, higher computing densities, scalable storage, more efficient infrastructure, environmental friendliness, space management, and improved security.

From a systems perspective, next-generation HPC architectures will see an explosion of disaggregated architectures, in which memory is decoupled from processors and accelerators, and heterogeneous systems, in which several specialized computing architectures (FPGAs, GPUs, CPUs, etc.) are integrated into a single node, allowing flexible switching between modules at scale. A key recipe for achieving this type of integrated system is the use of "chiplets." Chiplets are small modular components that can be combined to create a larger, more complex system. This approach allows for greater flexibility in system design, as well as better scalability and cost-effectiveness.

However, such complex systems pose a great verification challenge, especially with respect to IP/node-level verification in the context of the system, dynamic hardware-software orchestration, performance under workload and power, etc. This will require a push toward a new hardware-software verification approach. Verification of these systems will require a comprehensive approach that includes both hardware and software verification and system-level verification. This will require the development of new tools and methodologies that can handle the complexity of these systems, as well as the ability to verify system performance and power consumption under different workloads.

In addition, the use of chiplets and disaggregated architectures will require a new approach to system design and integration. This will require a shift toward a more modular approach to system design, where components can be easily replaced as needed. This will require the development of new standards and interfaces that can facilitate the integration of different components into a single system.

In general, the evolution of HPC architectures is driven by a combination of factors, including evolving artificial intelligence workloads, computational flexibility, cost, memory, and IO throughput. The use of chiplets and disaggregated architectures is a key trend in this evolution, but it also poses significant challenges in terms of verification and system design. Addressing these challenges will require a new approach to hardware-software verification and system integration.

3.5 Energy consumption

One of the challenges facing system operators today is the high energy and time cost of moving data, both of which are limited resources. To address this challenge, 2022 will see an accelerating trend of moving processing closer to the data to reduce the amount of data. In addition, there is a need to continue to scale resources to meet the growing demand for computing power.

Despite the rapid improvement in HPC capabilities and the extension of the benefits to many aspects of our lives, we are paying a price in terms of the environmental impact of these power-hungry systems. Data centers alone are expected to consume between 3 percent and 7 percent of the world's energy by 2030. Many data centers face obstacles and even new building permits because of the amount of electricity and water they need. Powering and cooling these huge IT platforms has become a hot topic for sustainability, and metrics such as PUE (Power Usage Effectiveness) and carbon dioxide emissions are taking center stage, starting with the boards of directors.

Fundamental changes in powering data centers through sustainable energy sources, such as hydropower, solar, and wind, are becoming more prevalent. New approaches, such as immersive or liquid cooling techniques (including underwater data centers), the redistribution and recycling of energy and water consumed by data centers for other uses (e.g., for heating buildings), and the use of greener components, materials, and manufacturing approaches by the supply chain ecosystem all have significant potential.

Another possibility is to improve the efficiency of power consumption and heat dissipation at the chip level. For example, high-performance integrated circuit designs can be better optimized for power with the use of advanced low-power design methodologies and the use of power-optimized IP cores to reduce overall chip and system power consumption.

The chiplet trend offers another significant way to reduce energy, not to mention reducing physical waste in manufacturing through higher yields. More power-sensitive data transport methods, such as high-bandwidth memory (HBM), can also make chips and the systems they power more energy efficient. These methods are favored by open-source standards and initiatives such as CXL, UCIe, and OCP.

In summary, the HPC industry is constantly evolving and expanding, bringing new improvements to our lives every day. However, its proliferation is a double-edged sword, as it creates a relentless and costly increase in performance in data creation and consumption, which can cause harmful environmental effects. To address this challenge, industry is exploring new approaches to power and cooling, as well as optimizing chip and system design to achieve energy efficiency.

4 External Advisory Board

The External Advisory Board (EAB) is a voluntary body to consult and support project management. Its opinions are not binding because they are not formal. Unlike the Project Technical Committee, the EAB does not act on behalf of the organization and therefore has no responsibilities.

The EAB task is mainly to help, advise and support the organization according to everyone's possibilities, skills, and opportunities: from marketing to relationship facilitation; from management to influence to third-party publics, from technical innovations to exploitation. Specifically, the tasks of the EAB could be simplified as follows:

- Gives a third-party perspective to issues of interest and on the organization's own processes.
- Gives advice and/or opinions on actions identified by the PTC.
- Suggests new actions or proposes development ideas.
- Identifies and reports opportunities.
- Facilitates the organization's relationships.

The establishment of an EAB is considered strategic, particularly for TEXTAROSSA, that has ambitious goals and wants to position itself as a leader in the European HPC scenario.

In terms of management, it is preferable for an EAB to consist of a small number of people: usually 3 to 5 elements. AB members should be identified for real opportunities with long-term impact. Thus, political choices that, while comfortable at first, may prove unfortunate in the long run are to be ruled out. In TEXTAROSSA the EAB is composed by only two members as it will be discussed later.

The EAB assignment is defined upstream until the end of the project and possibly renewed to follow the adoption of project innovations beyond the end of the project. In the start-up phase after the end of the project, a one-year assignment will be considered and proposed to EAB members.

For what concern specifically the TEXTAROSSA project, the External Advisory board is intended to provide input and recommendations to the project, help in disseminating results to the wider HPC, AI and HPDA community as well as represent an antenna for the project to gather industry trends and technology adoption.

Within the first three months of the project, the TEXTAROSSA Project Technical Committee (PTC) agreed on the criteria to select members of the Advisory Board. They should be:

- independent and external experts
- renowned experts in the field of HPC, AI and HPDA
- renowned and prominent scientists in their respective scientific communities
- renowned experts in national and international policy making
- experienced in research and project management
- linked to relevant TEXTAROSSA stakeholders
- experienced in industrial processes and environments.

Moreover, the PTC agreed to select Advisory Board members to get:

- expert advice on the project as a whole and on individual tasks
- comments and recommendations regarding requirements, objectives, and development, as well as exploitation and dissemination activities. This will maximize the impact of the TEXTAROSSA project and future exploitation of project results.

- recommendations to be discussed and processed in GA meetings.

In TEXTAROSSA the EAB is expected to be responsible for providing advice and strategic inputs considering risks and issues too. The EAB members should follow and consider the assessment of risks by the Consortium and the resolution actions.

A call for external advisory board members was then issued. Each TEXTAROSSA partner proposed one member for the Advisory Board, and two of them were selected to be actively involved. The PTC chose to have a two-person Advisory Board for two main reasons: the reduced resources devoted by the project to this activity and the great interest in focused discussions on the project tasks.

Members of the advisory committee were also selected based on their experience in industry, an important factor in the subsequent improvement of the project products. Thus, the members of the advisory committee were chosen among others because they have extensive experience in the HPC field in both the public and industrial research worlds. They have expertise in power management, thermal load, and a broad view of possible HPC applications.

The plan for the EAB management is the following:

- share with EAB the project documents
- give them access to project document repository
- give them access to project software repository
- invite them to GA meeting
- collect from them interviews and reports on specific project advancements and results.

The EAB has a tight connection with both the project coordinator (PC) and the project technical manager (PTM). Online and in-person meeting are organized to quickly update EAB members and plan future activities.

5 EAB Member: Dr. R. Squarcini

Dr. Raffaele Squarcini (PhD)
R&D – Modeling & Simulation
Pierburg Pump Technology Italy S.p.A.
Via Salvatore Orlando, 12
57123 Livorno, Italy

raffaele.squarcini@it.rheinmetall.com

5.1 Curriculum Vitae

Raffaele Squarcini was born on the 14th May 1973 in Livorno, he has 2 kids (3 and 6 years old).

Education:

- Secondary School: Liceo Scientifico Enriques Livorno, mark 58/60
- Master Degree in Nuclear Engineering (1999), mark 110/110
- PhD in Mechanical Engineering (2004), excellent

About Career path:

- 2002 –2004 Siemens VDO Automotive, he worked as Advanced\Innovation Engineer
- from 2005 Pierburg Pump Technology Italy SpA
    1) Pre-Development Oil Pumps
    2) Calcuation Department Manager
    3) Calculation Testing Department Manager
    4) **R&D Senior Manager**

R&D and Product Advanced has been his main environment up to now.
He always had continuous relations with universities, in particular:

- Pisa University Dept. Mechanical Eng., Electric Systems, Electronics
- Polito Dept. Of Oleodynamic
- Bologna University Dept. Of Mechanic (Fluid Machines)
- INSA Lyon Acoustic Department
- Firenze University Mechanical Department

He is Author\co-Author of more than 20 Patents, he is Author\co-Author of more than 30 papers about dynamic of mechanisms, fluid machines, acoustics, tribology, thermofluid dynamic of systems of system including electric motors and electronic control.

Main Passions: Cinema, Reading, Windsurfing, Kayak

5.2 Company

Dr. Raffaele Squarcini is now working for Rheinmetall. Rheinmetall operates as an internationally active group with technologically leading products and services in various markets. Sales are focused on the security technology and mobility segments. The Sensors and Actuators division is synonymous with outstanding expertise in thermal and fluid management applications in the mobility domain and industry. Its five business units – Actuators, Automotive Emission Systems, Commercial Diesel Systems, Pump Technology, and Solenoid Valves – develop solutions for major challenges facing the automotive sector, including downsizing, reducing emissions, and boosting efficiency, while simultaneously optimizing performance. The Pump Technology business unit combines top-notch system expertise with a wide spectrum of pumps for oil, vacuum and cooling systems destined for a full range of engine and cooling applications. The two main souls under the big Rheinmetall group are: Defense and Automotive.

6 EAB Member: Dott. S. Cozzini

Dott. Stefano Cozzini

Direttore Istituto Ricerca e Innovazione Tecnologica (RIT)
(Head of Research and Technology Innovation Institute)
AREA SCIENCE PARK
Padriciano, 34149 Trieste, Italy

stefano.cozzini@areasciencepark.it
www.areasciencepark.it

6.1 Curriculum Vitae

Stefano Cozzini born on 24th April 1966 in Trento (Italy) married, two daughters,

Education
PhD in Physics, University of Padova, Italy (1995)
"Laurea" in Physics, University of Trento, Italy (1991)

Short professional Profile

Dr. Stefano Cozzini has over 20 years' experience in scientific computing and HPC and data management e-infrastructures. His main professional activity is to plan and install and then efficiently maintain IT infrastructures. He is presently director of the Research and Technology Institute of Area Science Park, where he coordinates team of more than 20 people in maintaining cutting-edge research Genomic and Epigenomic infrastructures delivering high level services for scientific institutions.

Before this current position SC was founded and CTO of a start-up company (www.exact-lab.it) providing HPC services. In this role he was active in Strategic planning /Business developing activities and in the coordination of the whole R&D activities and research project.

S.C. has considerable experience in leading and coordinating research infrastructure projects at national and international levels. He also serves regularly as a referee for international conferences and scientific review and served as scientific referee and scientific advisor of several European projects.

S.C. has a large teaching experience in the above areas. Over the last 20 years he has been organizer and director of more than 20 International training activities on High Performance and GRID CLOUD and Data Infrastructures all over the world. From 2014 on he contributed actively to the International Master in High Performance Computing promoted by Sissa and ICTP (http://www.mhpc.it) where he

was its first coordination till 2016. He is also in the scientific council of the recently launched DSSC Laurea Magistrale (https://dssc.units.it/)) where he coordinates and teaches the course “Foundation of HPC”.

His main research interests are in the field of High-Performance Computing applied to computational science and industrial problems and in scientific data management Infrastructure and associated tools. In the last few years, he is actively working in the area of Machine Learning to provide advanced data services on top of computational infrastructure for heterogeneous scientific communities.

6.2 Company

Area Science Park is a public national research organization that promotes the development of innovation processes. For 40 years, its mission has been to boost connections between research and enterprise, public administration, and the private sector, supporting national and international initiatives and fostering territorial development.

The current lines of activity are:

- Science and Technology Park: attracting the best excellence and supporting growth.
- Business generation: support for business ideas.
- Innovation and complex systems: creation of innovation systems to support businesses, organizations, public administration.
- Technology platforms: supporting research activities.

The main AREA’s mission is to build a structured ecosystem of private companies and public operators capable of generating economic development in the different regional, national and international levels, by enhancing the assets of digitization, advanced technology and higher education present in the reference territories on a regional scale, in order to optimize resources and at the same time attract national and international investment in new strategic development projects.

AREA Science Park, the national entity that places the interaction between research and industry as the driving factor in this model to ensure development and growth. AREA Science Park, which for almost 40 years has been operating with concrete and measurable results under the coordination of the Ministry of Research, stands as a major player in innovation, using tools adapted to the economic, social, and technological context of reference. Therefore, not only an active player in the enhancement of research, but a player in industrial development.

AREA SCIENCE PARK

7 EAB vision

The EAB members received project documents and they were informed about objectives and current activities. Moreover, they were informed about their role and about what is expected from them to increase the competitiveness of TEXTAROSSA advancements.

Then their impressions about the project and their general vision of future trends in HPC were collected and reported here.

Question: Your first impression of the project is:

Squarcini: I had a very good impression of this project. I think it is necessary to deal with these current and future issues in the field of high-performance computing for both communications and technology evolutions.

Cozzini: The first impression is quite positive, and the project seems well structured and with a set of interesting and well posed challenges within the HPC Arena. The consortium is of high quality and with the right blend of competencies.

Question: What do you think are the most promising directions in hpc and big data?

Squarcini: I can answer regarding HPC and the possibility of increasing computing power. This makes it possible to reduce the time to market of product development as much as possible, reducing time and allowing more and more tests to be replaced by complex computations. It makes it possible to analyze complex problems from a high-level overview down to a very low level of detail.

Just-in-time visualization of modeled physical phenomena is and will increasingly be the true benchmark for product development.

Regarding Big Data, I cannot add specific professional comments (it is not my field of expertise), but it seems absolutely clear to me what the trend is. We have had a strong acceleration in this direction also "thanks to Pandemic," now it seems difficult to turn back. It is just a matter of managing and increasing capacity.

Cozzini: My feeling here is that HPC Arena is becoming increasingly heterogenous in term of computational workloads and therefore one of the most promising directions to explore is to develop and provide tools and methodologies to cope with this extreme diversity of the new types of computational workload. It is no longer a matter of computational power in rigid monolithic supercomputers, but the emphasis is now on the need for flexible and scalable HPC solutions that, sometimes, can be also made available on the cloud infrastructure.

Big Data and associated data analysis by means of Artificial intelligence and other techniques are just the most prominent workloads that modern HPC infrastructure should take care of. This makes quite evident that the real challenge here is to evolve the traditional HPC infrastructure toward a more integrated HPC/AI and HPDCI: High performance Data Intensive Computing Infrastructure.

The converge of HPC, Big Data, Deep Learning is becoming the next technical challenges and a concrete business opportunity.

Question: why do you think a company should use HPC?

Squarcini: I answer from the point of view of an R&D Expert in Mechanic\Mechatronic product. Because the product Developments involves day by day more and more the modulization of physical phenomena via computer, the time to market is reduced a lot in automotive for example and the Validation Testing phase is more and more reduced.

In addition, the integration of the pure mechanical + Electric Motors + Electronic controls requires HPC as a must for evalutions.

Finally, HPC for Data Server is absolutely a must.

Cozzini: As founder of the small company providing HPC services I strongly believe that "outcompute is outcompete". So, my short answer here is that HPC is a key-enabler technology that allows a company to stay ahead in terms of innovation of services and products. There are many excellent examples of successful stories with the traditional HPC approach.

Nowadays when AI and Big Data comes into play this aspect is even more important.

Question: what you think are the biggest obstacles for a company to use HPC?

Squarcini: Now I do not see any problems. Costs, storage space and a strategical balance between buy or use by third party paying a lease can be the question driving the HPC usage.

Cozzini: The lack of flexibility in the usage of the standard HPC approach and the lack of the human capital within the companies to manage and use an HPC infrastructure at best, even if resources are provided on the Cloud.

Question: what do you think are the most important requirements for a product developed in the HPC field? reliability, performance, power consumption, ease of use, etc.

Squarcini: I think in the question we can also find the answers: Reliability for sure, Performance for sure, and especially in this period the Power Consumption can play a fundamental role in the implementation of HPC in a Company and they can also influence the strategy buy or use via Third parts.

Cozzini: The most important ones are from one side the easiness of use and on the other the flexibility: the product should be able to be easily adapted/used under different scenarios. So generally speaking, a successful HPC product for the industrial market should be developed keeping in mind these two aspects and then developed to be reliable and well performant in terms of energy consumption.

Final discussion on the vision of the TEXTAROSSA project, as requested by the reviewers in the M18 review report. On this point, both EAB members agreed on the overall vision that TEXTAROSSA should have. We summarize their position below.

The activities of the TEXTAROSSA project focus on future trends in HPC, particularly on bridging the technology gap to achieve high performance and high energy efficiency on exascale computing systems of the near future. The project, called TEXTAROSSA, aims to apply a co-design approach to heterogeneous HPC solutions by integrating and extending IP, programming models and tools derived from European research projects. The main directions of innovation include enabling mixed-precision computing, implementing new multilevel thermal management and two-phase liquid cooling, developing improved tools for data movement and storage through compression, and providing IP for rapid task scheduling and low-latency intra/inter-node communication.

The project will test these technologies on prototypes that mirror and extend the ARM64-based architecture of the European Processor Initiative and on an OpenSequana testbed. To guide technology development and evaluate the impact of the proposed innovations, TEXTAROSSA will redesign and implement a selected but representative number of HPC, HPDA, and AI demonstrators covering challenging HPC domains.

The project focuses on HPC, typically used for large-scale scientific simulations and data analysis. Edge computing, on the other hand, focuses on bringing computing power closer to the data source, enabling real-time data processing and analysis. Although both HPC and edge computing are important in the world of computing, the TEXTAROSSA project focuses specifically on advancing HPC technology.

Although the TEXTAROSSA project is focused on advancing HPC technology, it could indirectly impact edge computing technology in a few ways.

First, the project's focus on improving energy efficiency and thermal management could lead to the development of more efficient and sustainable computing systems. These advances could be applied to edge computing systems, which often have limited energy and cooling resources.

Second, the project's focus on developing IP, libraries, and compilers that support new data types, including Posit, could lead to the development of more efficient and accurate artificial intelligence accelerators. These accelerators could be used in edge computing systems to enable real-time data processing and analysis.

Finally, the project's focus on developing better tools for moving and storing data through compression and providing IP for rapid task scheduling and low-latency intra/inter-node communication could lead to the development of more efficient and reliable edge computing systems.

Overall, although the TEXTAROSSA project focuses primarily on advancing HPC technology, its advances could have an indirect impact on the development of more efficient and sustainable edge computing systems.

8 Lesson learnt: IT knowledge best practices

TEXTAROSSA will produce an enormous amount of new knowledge with high potential for impact toward both the research and industrial worlds. To better interact with the EAB and all the external stakeholders, the project proposed to develop fundamental components of a knowledge management system that can enable value to be placed on the new products both during and after the end of the project.

To set up an effective line of development with the EAB, best practices have been defined within the first year of the project that should be considered guidelines in discussions with the EAB regarding project innovations. Such guidelines have been discussed in the TEXTAROSSA PTC and agreed by everyone to have a common feeling regarding the new ideas developed in the project. If everyone is aware of the value of the new knowledge developed in TEXTAROSSA, it means that everyone considers high value the process of knowledge management. In this respect we shared the idea that knowledge management refers to the process of creating, curating, sharing, using, and managing knowledge within the consortium.

Knowledge management is crucial because it aims to ensure that stakeholders receive the right information, in the right format, at the right level and at the right time, based on their level of access and other relevant policies. This requires a procedure for knowledge acquisition, including the development, acquisition, and collection of unstructured knowledge, regardless of whether the knowledge is formal and documented or informal and tacit.

More precisely we consider knowledge base as an online, self-service collection of information about a product, service, department, or topic. The data in the knowledge base can come from anywhere, but it is usually provided by several contributors who are experts on the relevant topic. The knowledge base can include frequently asked questions, discussions, troubleshooting guidelines, and any other information that is of interest to the community.

As IT services become more complex, teams must keep up with the wide range of technologies and procedures needed to provide effective support. Knowledge management, therefore, becomes more important than ever. It allows to:

- Creating value, it provides the right information to the right people at the right time.
- Promote innovation because it uses shared knowledge to foster brainstorming, collaboration, and big ideas.
- Achieve goals because it gives teams the tools, they need to set goals and achieve them.

For technical IT projects, knowledge management puts content at the fingertips of those who develop and deliver your products and services. In addition to being a benefit, this also shortens development cycles for new initiatives, improves relationships and communication between internal and external staff, enables more effective management of business environments, and leverages the intellectual capital and resources of your workforce.

Knowledge is one of the most valuable assets and open knowledge sharing can help partners stay current, collaborate, and make better decisions faster. Open sharing amplifies the value of knowledge because information is no longer an asset owned by the individual, but by the community. To promote more open knowledge sharing, here are some recommended best practices:

- Aggregate project's knowledge into a single repository or system. As technology in the work environment evolves, knowledge resides in increasingly disparate locations-emails, tickets, and in the minds of individual team members. As important as it is, choosing the right technology is only one step in a broader knowledge management strategy.
- Increase transparency with open, shared information: team members are invited to collaboratively edit pages, provide feedback via online comments, or mention team members for peer review.

The TEXTAROSSA activities are visible with a project poster. For each major initiative, a project poster is created to share goals and progresses with the rest of the team and stakeholders. This is a dynamic and accessible document that can help you explore the problem space, define the scope, and receive feedback.

Focus on short articles or responses. Shared documentation does not always mean shared understanding. Instead of creating long, extensive documents, the content will be tailored to WPs team. So, the entire team can learn and absorb information more quickly if it is quick to reference, written in easy-to-understand language, and posted in a timely manner.

Promote a culture of knowledge sharing. Reward the best employees with an ongoing recognition program that values both quality and quantity. Your management team can do much to set a positive example by regularly providing information such as important organizational updates. It can also direct staff to use your chosen tool and use it to interact directly with teams.

What is said, it is not always straightforward to be implemented. To this end some best practices are identified within the project:

- A clear organization of the project and a clear identification of the objectives
- Involve the project partners in the implementation of a knowledge management system
- Identify people from the partners interested to be active in knowledge management
- GA assembly and PTC meeting are good opportunity to check progresses in knowledge sharing and reuse.
- Aggregates all the information in one place, in a way that everyone can easily access to it and work in it.

Knowledge management benefits all the partners and improves communication both the EAB and to internal and external stakeholders. The underline idea is that all partners in TEXTAROSSA have valuable knowledge that is worth sharing. This is true for all project activities, all of them have knowledge that needs to be constantly shared with a team or an entire organization. With a plan in place, everyone can access project knowledge base to solve problems and prevent future ones.

To support the knowledge sharing the Office365 tools is used:

- A TEXTAROSSA team is created
- Private channels are created for each WP
- A cronoprogram for both activities and deliverables has been created for all partners
- A sharepoint environment has been created
- In the sharepoint a news section is available to everyone to share easily and friendly the project advancements.
- A project github is available to share codes and documents.

9 Conclusions and future works

This paper provides a comprehensive overview of the actions taken by the consortium in the first half of the TEXTAROSSA project. The consortium has made significant progress in defining the project goals, identifying the most important trends in HPC, and defining the role of the EAB in the project.

One of the key outcomes of this deliverable is the creation of a knowledge base management system that will enable the consortium to present project activities to both the EAB and external stakeholders. This system will be critical to ensure that stakeholders receive the right information, in the right format, at the right level, and at the right time, based on their level of access and other relevant policies.

In addition, this paper reports EAB members' vision of the future of HPC and its impact on the industrial environment. The insights provided by EAB members will be invaluable in defining the vision of the project and guiding it toward achieving its goals. This activity will ensure that the project's outcomes are relevant and impactful to the HPC community and society at large.

Looking forward, the EAB will be critically important in the second half of the project. The consortium will continue to work closely with the EAB to ensure that the project stays on track and that the project outcomes are aligned with the needs of the industrial environment. The EAB will be asked to maintain continuous liaison with the project management team and the PTC during and beyond the end of the project, ensuring that project outcomes continue to have a lasting impact.

Overall, this document highlights the importance of collaboration and knowledge sharing in the TEXTAROSSA project. The consortium is committed to working closely with the EAB and other stakeholders to ensure that the project results are relevant, impactful and aligned with the needs of the industrial environment.