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

WP7 Dissemination, Communication and Exploitation

D7.5 Communication and Dissemination Report 2

http://textarossa.eu

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: D7.5 Communication and Dissemination Report 2

Project Start Date: 2021-04-01 **Duration**: 36 months

Coordinator: *AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE, L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE - ENEA , Italy.*

Deliverable No	D7.5
WP No:	WP7
WP Leader:	CINI-POLIMI
Due date:	M36 (2024-03-31)
Delivery date:	2024-04-25

Dissemination Level:

PU	Public	X
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)	

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:	2021-04-01
Duration of the project:	36 months
Project website:	TEXTAROSSA.eu

WP7 Dissemination, Communication and Exploitation

Deliverable number:	D7.5					
Deliverable title:	Communication and Dissemination Report 2					
Due date:	M36					
Actual submission date:	2024-05-01					
Editor:	Federico Reghenzani					
Authors:	W. Fornaciari, G. Agosta, F. Reghenzani, M. Kulczewski, A. Oleksiak, S. Ciecielski, P. D'Ambra, B. Cantalupo, L. Eyraud-Dubois, B. Bramas, O. Beaumont, A. Lonardo, D. Jimenez, C. Alvarez, M.V. Fontanesi, A. Lonardo, M. Celino, K. Martin, C. Claire, P. Palazzari, M.A. Vignon, E. Boella					
Work package:	WP7					
Dissemination Level:	Public					
No. pages:	48					
Authorized (date):	2023-04-25					
Responsible person:	William Fornaciari					
Status:	Plan	Draft	Working	Final	Submitted	Approved

Revision history:

Version	Date	Author	Comment
0.1	2024-03-15	F.Reghenzani	Initial Draft
0.2	2024-03-22	F. Reghenzani	Update website section
0.3	2024-04-02	P. D'Ambra	Update CNR section
0.4	2024-04-02	W. Fornaciari, G. Agosta	Working on several sections
0.5	2024-04-04	B. Cantalupo, M. Kulczewski, A. Oleksiak, L. Eyraud-Dubois	Update of respective partners sections
0.6	2024-04-05	B. Bramas, W. Fornaciari	Update of respective partners sections
0.7	2024-04-07	D. Jimenez, C. Alvarez, M.V. Fontanesi, A. Lonardo, B. Cantalupo	Update of respective partners sections
0.8	2024-04-09	M. Celino, K. Martin	Various updates
0.9	2024-04-16	S. Saponara	Modified CINI part
0.10	2024-04-19	C. Claire, P. Palazzari, M.A. Vignon, G. Zummo	Update of respective partners sections
0.11	2024-04-20	E. Boella, A. Lonardo	Update of respective partners sections
0.12	2024-04-24	F. Reghenzani, W. Fornaciari	Final revisions

Quality Control:

Checking process	Who	Date
Checked by internal reviewer	Barbara Cantalupo	2024-04-23
Checked by Task Leader	Federico Reghenzani	2024-04-23
Checked by WP Leader	William Fornaciari	2024-04-24
Checked by Project Coordinator	Massimo Celino	2024-04-25

COPYRIGHT

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

Table of contents

List of Figures

List of Tables

No tables in the document.

Executive Summary

This deliverable reports the dissemination and exploitation activities of the TEXTAROSSA project and the achieved results. We report the overall dissemination and exploitation results achieved by the project, individual dissemination and exploitation plans undertaken by each partner. Dissemination activities are reported as an update with respect to the previous deliverable D7.4. In particular, the deliverable highlights the scientific output of the project, which consists in 10 journal articles, 26 conference papers, and 6 other publications. This result puts TEXTAROSSA over the initial plans, which foresaw a total of 5 journal articles and 15 conference paper during the entire span of the project. Regarding the exploitation, 17 out of 18 innovation products reached the target TRL or better, including both open source and commercial innovation products. All these innovation products are listed and described in Section 4.

The description of dissemination activities of Section 3 has been updated with respect to the previous deliverable D7.4 by integrating the activities of the second part of the project. It is worth noticing that CINI-UNITO purchased the first commercial E4 prototype of a GPU-enabled server with evaporative cooling, based on the TEXTAROSSA technology. It was installed in the HPC4AI green datacenter of CINI-UNITO for wider testing with the existing commercial users and to foster further scientific cooperation with the project partners.

1 Report and Plans on Communication and Dissemination Activities

The dissemination activities of the TEXTAROSSA project can be split into scientific publications, other dissemination activities (talks, presentations, general public dissemination, etc.), networking activities, and web-related dissemination (such as the website and social media). In the following sections, we will report the result of dissemination activities of the project for each of the mentioned categories.

1.1 Scientific publications

At the date of writing of this deliverable, the consortium published a total of 10 international journal articles and 26 conference papers related to the TEXTAROSSA project. In addition, the consortium presented 6 posters and talks at scientific venues. The list of publishers of the published articles is as follows:

- IEEE: 4 journal articles, 10 conference papers
- ACM: 1 journal, article, 2 conference papers
- Elsevier: 4 journal articles
- IOP Science: 1 journal article
- Springer: 6 conference papers
- Other publishers: 8 conference papers

The goal for publications reported in the grant agreement and D7.3 is 5 journal publications, 15 conference papers, and 1 book plan. At the end of the project, **the consortium achieved 200% of the goal for journal articles and 166% of the goal for conference papers** categories. In addition to this number that includes only already published documents, several articles are still under review. The full list of the published papers is available on the website (https://textarossa.eu/dissemination/publications/) and will be updated in the coming months after the end of the project adding the accepted papers that are currently under submission and not yet published. The current list of publications is also available in Appendix 1 of this document.

A special mention is for the conference paper at Euromicro Conference on Digital System Design (DSD) conference [P34] and the derived journal article in Microprocessors and Microsystems [P8] that present the project and its goals to the scientific community. The conference paper [P36] has been presented at the conference Digital System Design (DSD) in Palermo, Italy in 2021 (moved to virtual due to COVID'19).

Currently, under the guidance of BSC, we are working on a new joint paper to be submitted to DSD 2024 (August 28-30, 2024) with the goal to present the main achievements of the TEXTAROSSA project.

In accordance with the open access requirements of Article 29.2, all the articles either have been published with an open access policy (gold open access) or an authors' version has been published in institutional repositories (green open access). As soon as they were made available, the open access links have been published in the TEXTAROSSA project website.

Task T7.2 (M1-M36) prepared the prospective content of a book and contacted an editor with a proposal, thus achieving the task goals. We are currently in contact with the Springer Editorial Director Dr. Charles Glaser who accepted to publish the book and the contract has been already signed in April 2024. The plan

is to have the proofs of the book ready for review before October 1st, 2024. The Editors of the books will be prof. William Fornaciari (CINI-POLIMI), prof. Carlos Alvarez (BSC, UPC) and Dr Massimo Celino (ENEA, Project Coordinator). Prof. William Fornaciari has experience in these activities, having already published two books at the end of EU projects (Multicube and Harpa).

1.2 Other dissemination activities in scientific venues

In addition to the scientific publications above, the following presentations have been made in the context of the TEXTAROSSA project in scientific venues:

1. Olivier Beaumont presented the TEXTAROSSA project at the Teratec forum (https://teratec.eu/gb/forum_2021/index.html) in Paris, June 2021.
2. Dr. Francesco Simula (INFN) presented "Distributed and Plastic Spiking Neural Network model of the brain cortex behavior" at PSNC Internal Seminars Series, February 2022.
3. Dr. Alessandro Lonardo (INFN Roma), presented an overview of the project with a focus on INFN activities in a talk with title "Il progetto TEXTAROSSA - Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale" at the 2022 Workshop of INFN Computing and Network Committee, Paestum, Italy, May 2022. https://agenda.infn.it/event/30202/contributions/168462/attachments/91434/124134/Lonardo-TEXTAROSSA-WS%20Calcolo%20INFN.pdf
4. Dr. Paolo Cretaro (INFN Roma) gave a talk with title "APEIRON: a heterogeneous computing platform for real-time inference" at the AI@INFN – Artificial Intelligence at INFN Workshop, Bologna, Italy, May 2022. https://agenda.infn.it/event/29907/contributions/163460/attachments/90355/121698/APEIRON_WS_AI_Bologna.pdf
5. Dr. Giuseppe Zummo (InQuattro) presented "Innovative Two-Phase Cooling Solutions for the Exascale Computing Systems" at ISC High Performance Conference, May 2022.
6. Prof. William Fornaciari (CINI-POLIMI) presented the project with a talk entitled "Design of secure power monitors for accelerators, by exploiting ML techniques, in the Euro-HPC TEXTAROSSA project" at the SCADL workshop co-located with IPDPS conference, June 2022.
7. Prof. Marco Aldinucci (CINI-UNITO) presented "From small files to no files" at 6th Workshop on Performance and Scalability of Storage Systems. Saclay, France, Jun. 2022.
8. Prof. William Fornaciari (CINI-POLIMI) presented an overview of the project and a description of a specific technology, in the keynote "Design of secure power monitors for hardware accelerators", given at the Conference SAMOS 2022, Samos, Greece, July 2022.
9. Dr. Iacopo Colonnelli (CINI-UNITO) presented "Hybrid Workflows For Large-Scale Scientific Applications" at 6th EAGE High Performance Computing Workshop, Sep. 2022.
10. Dr. Iacopo Colonnelli (CINI-UNITO) presented "Hybrid workflows for heterogeneous distributed computing" at 3rd Italian Workshop on HPC (ITWSHPC), Torino, Italy, Sep. 2022.
11. Dr. Alessandro Lonardo (INFN Roma) gave a talk with title "APEIRON: composing smart TDAQ systems for high energy physics experiments" at the 21st International Workshop on Advanced Computing and Analysis Techniques in Physics Research - ACAT 22, Bari, Italy, Oct. 2022. https://indico.cern.ch/event/1106990/
12. Prof. Marco Danelutto (CINI-UNIPI) presented "HPC@CINI: the HPC Key technology and tools lab experience" at the ACM Computing Frontier, Torino, Italy, May 2022.
13. Dr. Iacopo Colonnelli (CINI-UNITO) presented at the NVidia HPC round table, Bologna, September 2022.

14. Dr. Iacopo Colonnelli (CINI-UNITO) and Prof. Marco Aldinucci (CINI-UNITO) presented “Hybrid Workflows for Large-Scale Scientific Applications” at 6th EAGE High Performance Computing Workshop, Milano, Italy, September 2022.
15. Prof. Marco Aldinucci (CINI-UNITO) presented “From HPC4AI to ICSC living lab: Where systems are the research” at Dell Advanced Computing Workshop 2023: HPC and Beyond in Bologna, Itay, February 2023.
16. Prof. Marco Aldinucci (CINI-UNITO) presented “HPC4AI: The Research on AI beyond the public cloud” at CENTAI kick-off meeting in Torino, Italy, March 2023.
17. Prof. Marco Aldinucci (CINI-UNITO) presented “Experimenting with Systems for Decentralized Machine Learning” at NVidia GTC 2023, 2023, virtual, March 2023.
18. Prof. Carlos Álvarez. Presented the invited talk “Improving resource usage in large FPGA Accelerators” at 15th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop in INRIA Bordeaux, France, March 2023.
19. Cristian Rossi (INFN Roma) gave a talk at the 26th International Conference on Computing in High Energy and Nuclear Physics - CHEP 2023, Norfolk, Virginia (USA), with title “APEIRON: a Framework for High Level Programming of Dataflow Applications on Multi-FPGA Systems”, May 2023. https://indico.jlab.org/event/459/contributions/11827/attachments/9668/14279/CHEP2023_APEIRON_INFN.pdf
20. Cristian Rossi (INFN Roma) presented a talk with title “APEIRON: a Framework for High Level Programming of Dataflow Applications on Multi-FPGA Systems” at the 2023 Workshop of the INFN Computing and Network Committee, Loano (SV), Italy, May 2023. https://agenda.infn.it/event/34683/contributions/197361/attachments/105626/148510/CCR2023_APEIRON.pdf
21. Prof. Daniele Gregori (CSO-E4) and Dr. Giuseppe Santomauro (ENEA) presented TEXTAROSSA project at ISC High Performance 2023 (Hamburg, Germany). The presentation was hosted at E4’s booth, May 2023.
22. ISC High Performance conference 2023. Research poster “An FPGA-based platform to evaluate Posit arithmetic in next-generation processors” (N. Neves, L. Crespo, F. Rossi, M. Cococcioni, S. Saponara, M. Kuehn, J. Krueger, P. Tomas, N. Roma), Hamburg, Germany, May 2023.
23. Prof. Daniele Gregori (CSO-E4) presented TEXTAROSSA project at EuroHPC Meeting in Turin, June 2023.
24. Prof. Carlos Alvarez. Presented the talk “Extracting performance out of task-based FPGA programs” at HEART 2023: 13th International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies in Lake Biwa, Japan, June 2023.
25. Prof. William Fornaciari (CINI-POLIMI) gave a half-day tutorial “Energy and power management in the computing continuum”, during the SAMOS Conference, presenting some of the achievements of TEXTAROSSA regarding the thermal management, July 2023.
26. Prof. Marco Aldinucci (CINI-UNITO) delivered the invited talk “Federated Learning: A Distributed System Viewpoint” in the context of the Bicocca University seminars, Milan, Italy, December 2023.
27. Prof. Robert Birke (CINI-UNITO) gave a keynote speech “The impact of the advances in generative models on applications and systems” at 8th GDR RSD / ASF Winter School on Distributed Systems & Networks 2024} at Le Pleynet, France, January 2024.
28. Dr. Paolo Palazzari (ENEA) and Cristian Rossi (INFN Roma) gave a joint talk with title “Multi-FPGA performance scaling of High-Level Synthesis applications through the APEIRON Framework” at the EuroHPC Projects Shaping Europe's HPC Landscape HiPEAC 2024 associated workshop, Munich, Germany, January 2024.

29. Joint talk with the other companion project at EuroHPC conference. Contribution to the presentation “European System Architecture Advancements” (H.-C. Hoppe, FZJ) at the “Co-Designing the Future of European HPC: Eco-friendlyTechnologies, Systems and Software” parallel session, March 2024. https://www.hipeac.net/media/private/56/8085/hipeac-2024-enea-infn-compresso.pdf
30. Prof. Carlos Alvarez presented the talk “New scheduling strategies for task-based (hardware) runtimes” at the 6th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop in Kobe, Japan, April 2024.

1.3 Dissemination activities toward the general public

In addition to the dissemination activities delivered by the consortium members in scientific venues, prof. Marco Aldinucci (CINI-UNITO) presented the TEXTAROSSA project in the following events for general public:

- UNIGHT 2022: EU Researchers’ night, Torino, Italy, Sep. 2022.
  - > 1000 contacts, mostly with students (secondary school, BSc, MSc).
- Lectio Magistralis at the finals of Italian Olympic Games of Informatics, Biella, Italy, Sep. 2022.
  - > 100 secondary school finalists to the Olympic games 2022 and their families.
- “La convergenza HPC-cloud è l'anello mancante tra il calcolo scientifico e l'IA applicate”, Virtual, Sep. 2022.
  - > 100 participants from Italian SMEs.
- “Da HPC4AI al living lab dello spoke FutureHPC del Centro Nazionale HPC, Condivisioni – Conferenza GARR”, Palermo, Italy, Jun. 2022
  - > 100 participants from the community of Italian universities in the area of networking and cloud.

Olivier Beaumont (INRIA) participated in "Maths en Jeans" in 2023 and 2024 with groups of students (4 groups of 4-5 high school students working on a given topic throughout the school year) on data distributions for linear algebra problems.

1.4 Networking activities

The consortium, in particular Prof. Marco Aldinucci (CINI-UNITO) as leader of the task T7.3 (Networking with EU HPC landscape and Centres of Excellence) and E4, participated in the following international networking events among researchers:

- EuroHPC EoCoE final summit, Napoli, Italy, Jun 2022
- ELIXIR Cloud, Data & AAI Bi-weekly Technical Calls, Virtual, 2022
- HiPEAC Vision meeting, Brussels, Belgium, May 2022
- EuroHPC summit week, Paris, France, Mar 2022
- Teratec meeeting, Saclay, France, Jun 2022
- HPC Day within "Critical Infrastructure Protection Forum - CIP FORUM V 2022 - Critical Infrastructure Protection & Resilience Europe 2022”, Bucharest, Jun 2022
- First EuroHPC19 Workshop to Seed and Foster Collaborations Across Europe, Madrid, Spain, Sep. 2022

- HiPEAC 2023, Toulouse, France, 2023
- Supercomputing 2023, Denver, United States, 2023
- HPCAI Advisory Council, Lugano, Switzerland, 2023
- 2nd Italian Conference on Big Data and Data Science (ITADATA), Napoli, Italy, 2023
- HiPEAC 2024, Munich, Germany, 2024
- EuroHPC 2024 Summit, Antwerp, March 18-21, 2024

A more detailed description of the networking activities can be found in the deliverable D7.8 concerning the final report of the collaboration plan.

Figure 1: Participants at the EuroHPC19 Collaboration Workshop – Madrid – Sep 2022 (participants in TEXTAROSSA: M. Celino, M. Aldinucci, P. Palazzari)

1.5 Website and Social Media reports

All the journal articles, conference papers, and presentations published and presented by the consortium have been published on the website in the "Publications" page. This page will be updated even after the end of the project for papers that are currently under submission to be disseminated. All the papers have been published either in gold open access (so freely accessible via the editor website) or with green open access (editor version is closed and requires a fee, while an author's version is freely available through institutional repositories). In the "Publications" page, to each article is indicated the Digital Object Identifier and the relative link, in addition to a link to the pre-print version of the article if published in green open access.

The consortium has published several blog posts on the website, containing information on the status of the research project. These have been authored by ATOS, INRIA, ENEA, INFN, CNR, and CINI. These blog posts increased the dissemination to the general public, also with reposting on social media. In addition to the blog posts, news about project events and general updates are also published on the website. For instance, we reported the experiments, with the relative photos, of the new two-phase cooling system by InQuattro and CINI-POLIMI (link).

All the public deliverables have been uploaded to the project website (link) and are indexable by search engines.

1.5.1 Website analytics

As explained in D7.4, we had to change the analytics engine used for the website due to privacy reasons. For this reason, the analytics of the website are reported separately for periods M1-M15 and M16-M36.

1.5.2 Website analytics – M1-M15

During the first 16 months of the project, the website was accessed by 1750 unique visitors, with about ~20 unique users/day at the beginning of the project, with an expected decrease to ~10 unique users/day after the first 6 months. Figure 2 shows the daily trend. The dates of blog posts, especially from January 2022, are clearly visible and demonstrate the effectiveness and outreach of the blog posts, with an average of ~40 extra unique users/day with a peak of 52 extra unique users/day.

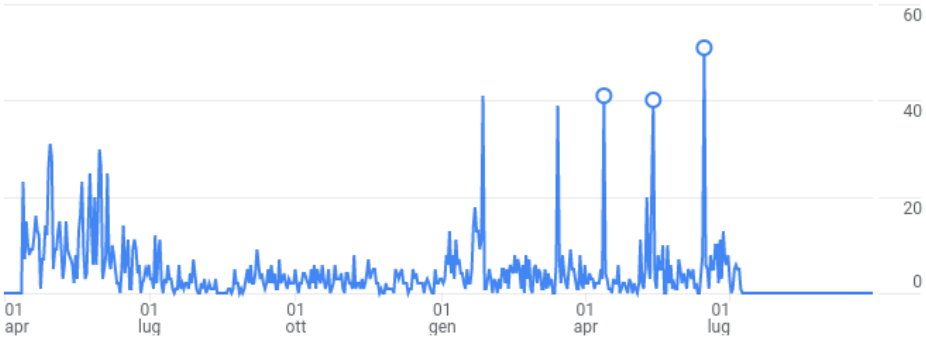

Figure 2: Daily unique users as reported by GA

Geographically, the users are mostly distributed among European Union and United States (US). The highest number of unique users is from Italy: this is a direct consequence of the fact that many partners of the consortium are based in Italy (> 500 unique visitors). It is worth mentioning that also from US, even if no partners are based in US, there is a considerable number of unique visitors (> 400 unique visitors). The geographical distribution and evolution are shown in Figure 3.

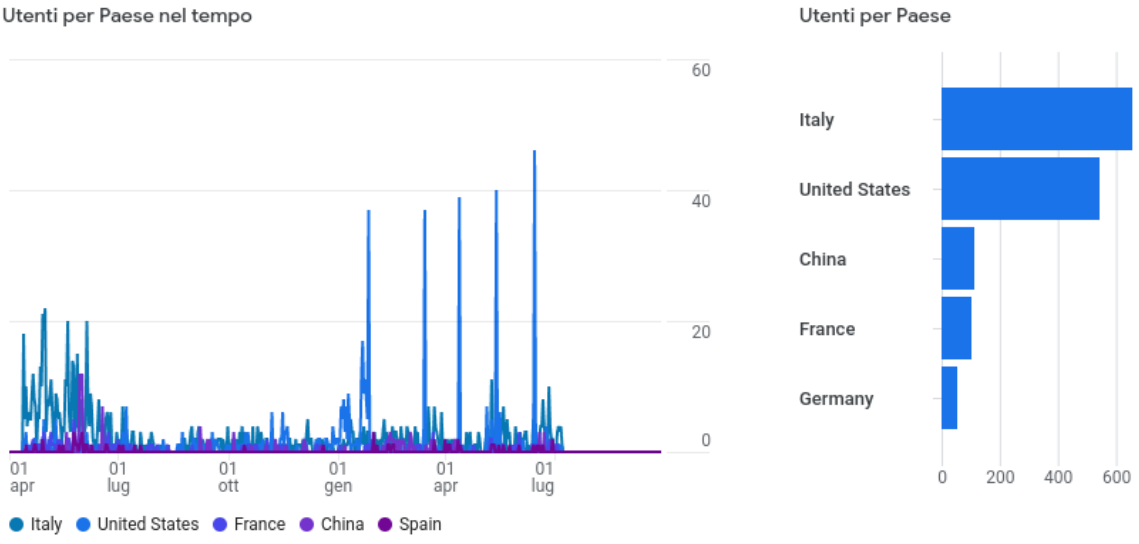

Figure 3: Geographic distribution of the users

Regarding the referral information of the users, depicted in Figure 4, most of the users come directly typing in the website address. This is the usual effect of the "direct" dissemination activities -- such as presentations, flyers, etc. -- where the website address or qr-code is visible. The second source is from search engines, and then referral from other websites, and social media.

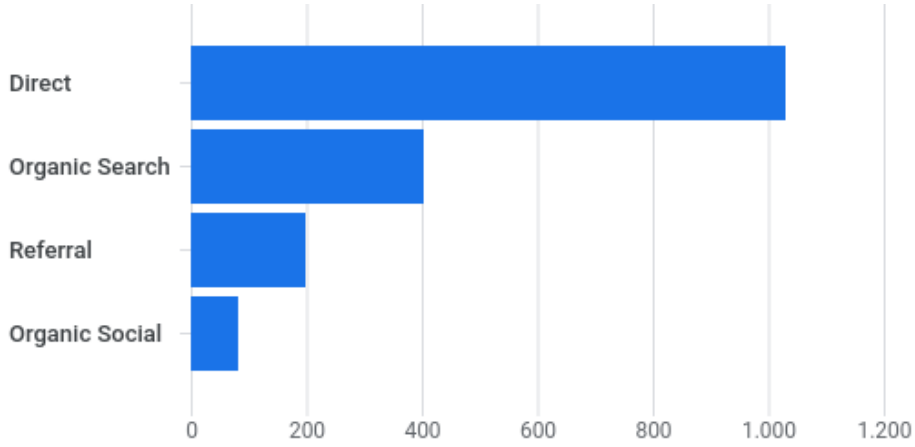

Figure 4: How the users find the TEXTAROSSA website

Finally, the information on specific pages is shown in Figure 5. Apart from the homepage which is clearly the most visited, the consortium page and the page containing project description are the most accessed. The list of blog posts (News page) is also in the top-5 visited list, including the page describing the TEXTAROSSA co-design approach.

	Titolo pagina e...asse schermata +	↓ Visualizzazioni	Utenti
		5.961 100% del totale	1.750 100% del totale
1	TEXTAROSSA – Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale	2.167	1.107
2	Consortium – TEXTAROSSA	769	385
3	Project Overview – TEXTAROSSA	484	307
4	News – TEXTAROSSA	266	168
5	A Co-Design Approach – TEXTAROSSA	223	149
6	Work Packages – TEXTAROSSA	223	148
7	Page not found – TEXTAROSSA	197	197
8	Publications – TEXTAROSSA	197	131
9	Press Release and Media Coverage – TEXTAROSSA	167	100
10	Applications – TEXTAROSSA	161	108

Figure 5: Most visited pages (first column shows visits, second column shows unique users)

1.5.3 Website analytics – M16-M36

During the last 20 months of the project, the website was accessed by 2825 unique visitors. Figure 6 shows the daily trend which is consistent with the previous period. The peak visit was during the week April 3-9, 2023, when the website received 113 visits.

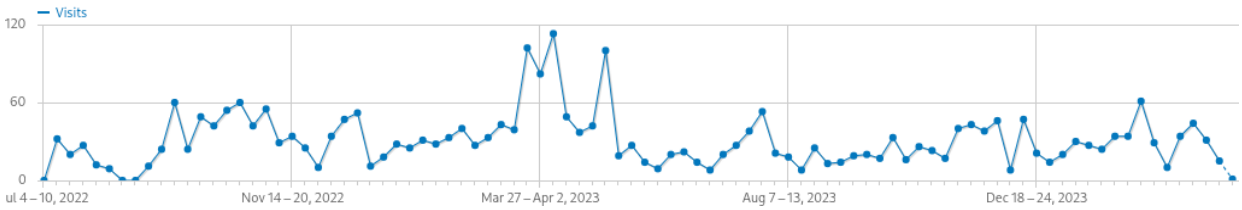

Figure 6: Daily visits as reported by Matomo

Geographically, the users are mostly distributed among European Union and United States (US) as visible in the graph of Figure 7. The highest number of unique users (28% of the total) is from Italy: this is a direct consequence of the fact that many partners of the consortium are based in Italy. It is worth mentioning that also from US, even if no partners are based in US, there is a considerable number of unique visitors (15% of the total).

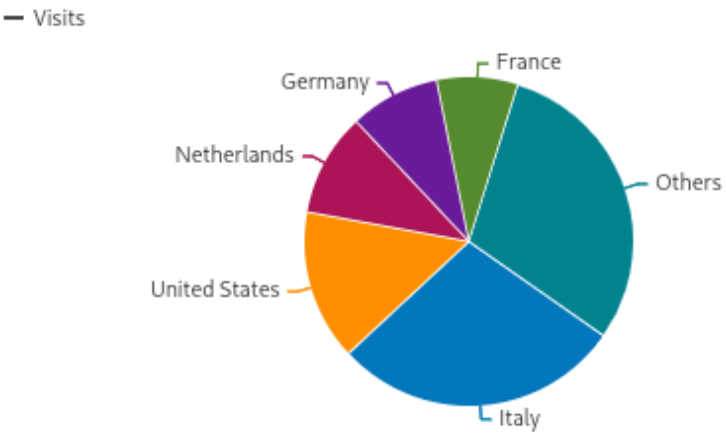

Figure 7: Geographic distribution of the users

Differently from the previous period, the most accessed pages (excluding the homepage) are the pages of the dissemination material and, in particular, the publication list. The statistics are visible in Figure 8.

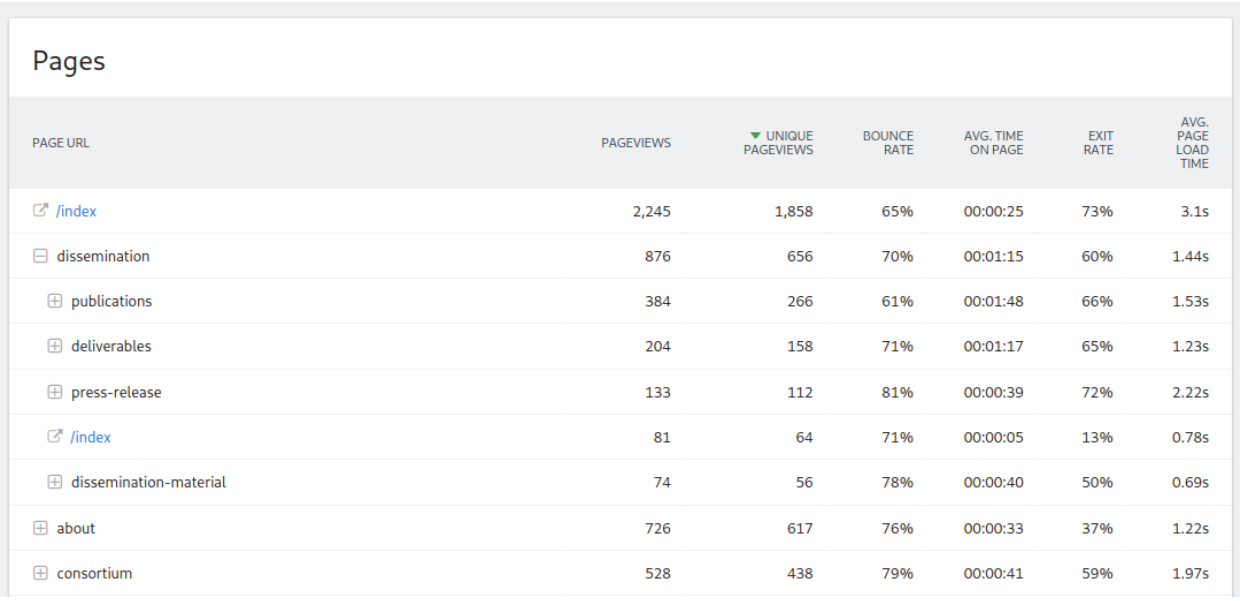

Pages

PAGE URL	PAGEVIEWS	▼ UNIQUE PAGEVIEWS	BOUNCE RATE	AVG. TIME ON PAGE	EXIT RATE	AVG. PAGE LOAD TIME
/index	2,245	1,858	65%	00:00:25	73%	3.1s
dissemination	876	656	70%	00:01:15	60%	1.44s
publications	384	266	61%	00:01:48	66%	1.53s
deliverables	204	158	71%	00:01:17	65%	1.23s
press-release	133	112	81%	00:00:39	72%	2.22s
/index	81	64	71%	00:00:05	13%	0.78s
dissemination-material	74	56	78%	00:00:40	50%	0.69s
about	726	617	76%	00:00:33	37%	1.22s
consortium	528	438	79%	00:00:41	59%	1.97s

Figure 8: Page statistics

1.5.4 Social Media

The TEXTAROSSA project has a landing page on three social media: Facebook, Twitter, and Linkedin. The latter is the most accessed by the users. The social networks have been primarily used to promote the website news, event announcements and blog posts. Additionally, the LinkedIn social network has been used to publicize the job offers from the consortium partners.

2 Individual Dissemination Plans and Reports

In addition to the dissemination activities performed by the consortium as a whole presented in the previous sections, some partners report the individual dissemination activities and plans for future beyond the project timeframe in the following sections.

ATOS and InQuattro have only limited dissemination effort and employed it to participate joint dissemination efforts including project-wise papers and participation to poster sessions at relevant conferences. For the other partners, individual dissemination plan and reports will follow.

2.1 ENEA

ENEA participated in several dissemination actions. Among the others, ENEA contributed to the project papers and posters. ENEA presented the TEXTAROSSA project at the SC21 Conference in USA in November 2021, 2022 and 2023, at ISC22 Conference in June 2022, at the EuroHPC19 Coordination Workshop in Madrid (September 2022) and HiPEAC 2024. The TEXTAROSSA information is delivered on the ENEA websites, among the other the ICT website (www.ict.enea.it). Finally, an ENEA press release was published in 2021.

2.2 FHG

Reduced precision and mixed precision floating point operations become more widely accessible on various hardware platforms. The usability and value of mixed precision for numerical algorithms has been demonstrated, creating the opportunity to improve the current state of the art and to make the achievement accessible through submission to scientific conferences or journals.

2.3 CINI

CINI leads the dissemination and communication activities. As such, it has set up the website, the social media accounts, and the project communication. Furthermore, CINI presented the paper [P34] to the Euromicro Conference on Digital System Design (DSD) 2021 edition and led the writing of the journal extension [P8] published by MICPRO. In addition to these joint publications with other partners, CINI co-authored 6 journal articles [P1] [P3] [P6-P8] [P10] and 14 conference papers [P15-P18] [P22-P26] [P28-P31] [P34]. CINI also presented the project at the SCADL workshop co-located with IPDPS conference, Lyon, France. CINI further disseminated information about TEXTAROSSA through the flyer at the SAMOS conference 2022 and during a tutorial held during SAMOS conference in 2023. CINI also presented the project in several talks at scientific venues (see Section 1.2) and events for general public (see Section 1.3). In the context of networking activities, reported in Section 1.4 and detailed in D7.8, CINI had the opportunity to further disseminate the project technology and results in the European projects landscape.

CINI plans to further disseminate project awareness through several channels, including the submission of an article focusing on the project achievements to the 2024 fall issue of HiPEAC Newsletter. CINI has also articles under review, expected to be published in the coming months and, in addition CINI is preparing various future papers submissions (e.g., a submission to DSD'24 in cooperation with BSC) and leads the efforts for the joint book proposal with Springer. Furthermore, CINI contributes to the overall dissemination plan by maintaining the website and coordinating the activities for the project blog posts. The TEXTAROSSA web site will be maintained beyond the duration of the project, for at least 3 years and after such timeframe it will be made static and perpetual.

2.4 INRIA

INRIA participated in several dissemination actions. INRIA presented the TEXTAROSSA project at the Teratec forum (https://teratec.eu/gb/forum_2021/index.html) in June 2021. In a very different context, we also presented the project, in relation to the energy minimization issue, to high school students in Libourne (France). The high school students are to take up the elements of our discussion during an event organized (mid-November 2024) by the scientific culture center of Bordeaux CapScience (https://www.cap-sciences.net/en/homepage/) in the framework of the COP 27 conference. Finally, we have published several scientific papers related to:

- Scheduling algorithms tailored for heterogeneous computing systems [P27], [P11].
- Optimization of linear algebra kernels, in particular related to data distributions for symmetric kernels [P20], [R4], [R5].
- Memory efficient techniques for both inference (using StarPU and OmpSs as runtime systems in addition to ONNX) and training (based on re-materialization and offloading techniques) [R2], [R3].

2.5 E4

During the duration of TEXTAROSSA, E4 performed a series of dissemination and communication activities to showcase the project and its results.

Figure 6 9: Example of LinkedIn post by E4

In 2022, E4 presented TEXTAROSSA at the ISC High Performance event in Hamburg running from May 29th to June 2nd. In particular, E4 hosted at its booth talks by TEXTAROSSA partners Paolo Palazzari from ENEA and Giuseppe Zummo from InQuattro. The talks were well attended, as shown in Figure 6. The talks were also advertised on the company social media channels (see for instance Linkedin post in Figure 6). E4 was one of the sponsors of the International Conference in High Energy Physics (ICHEP) 2022 held in Bologna

from July 6th to 13th. A poster about TEXTAROSSA was shown by Daniele Gregori (E4 Chief Scientific Officer) during the duration of the conference. In 2022, E4 participated also in the Big Science Business Forum (Granada, October 4th – 6th). In this occasion, the company had a small booth where the TEXTAROSSA poster was shown, and Daniele Gregori gave a talk where the project TEXTAROSSA was explained.

In 2023, E4 presented a poster about TEXTAROSSA at the High Performance, Edge And Cloud computing (HiPEAC) conference in Toulouse (January 16th – 18th). The poster was showcased also in the booth of the company. E4 had a booth also at the ISC High Performance event in Hamburg in May 2023. This time, E4 could show IDV-E. The node was displayed on the wall of the booth (see Figure 7) and attracted many visitors. E4 discussed the TEXTAROSSA project at the 31st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (Naples, March 1st – 3rd, 2023) with a poster presentation and at the 2023 Swiss Conference & HPCXXL User Group (Lugano, April 3rd – 6th) with an oral presentation delivered by Daniele Gregori. A poster about TEXTAROSSA was presented at the event Computing Frontiers held in Bologna from May 9th to 11th, 2023. On June 7th, 2023, Daniele Gregori gave a talk focused on the project at the EuroHPC Project Coordination Meeting in Turin. He also described TEXTAROSSA in an oral presentation at the 2nd Italian Conference on Big Data and Data Science (Naples, September 11th – 13th).

In 2024 E4's effort in Communication and Dissemination to promote TEXTAROSSA at a national and international level continued. TEXTAROSSA poster was shown at the E4 booth at HiPEAC 2024 (Munich, January 17th – 19th). Finally, an oral presentation about TEXTAROSSA was given by Elisabetta Boella (E4 HPC product specialist) at the Italian national workshop Future HPC and Big Data organised by the Italian Research Center on High Performance Computing, Big Data and Quantum Computing (ICSC) on February 14th, 2024, in Turin.

Figure 7:10 E4 booth at ISC High Performance 2023, where IDV-E was displayed.

2.6 BSC

BSC has participated in the dissemination of TEXTAROSSA by participating on the papers published in DSD 2021 and MICPRO 2022, with the description of the plan for implementing the fast task scheduler to be used by the TEXTAROSSA programming models, and the benefits that we have been observing during our developments. BSC published and presented two TEXTAROSSA related papers in the 9th BSC Doctoral Symposium [P31], a paper at the 2023 FCCM [P12] and a paper at the 2023 FPT conferences [P13]. Also, an IEEE Transactions on Computers journal paper [P2] has been published with the work regarding the fast task scheduler integration in a RISC-V core.

In addition, BSC organized three editions of the OmpSs@FPGA PATC course around the use of the OmpSs programming model for FPGA devices including the new developments achieved in the TEXTAROSSA project. BSC also presented OmpSs@FPGA TEXTAROSSA improvements to researchers at EuroCC, at the 2023 HEART conference in Japan and in the 15th JLESC conference in Bordeaux. Another presentation was done at the 16th JLESC conference in Japan in 2024.

BSC will lead the planning of the paper and the collection of the contributions by all the project partners to submit a comprehensive and final paper to the DSD 2024 conference that, possibly, will be published as an extended version on the Micpro Journal, Elsevier. Two more papers regarding TEXTAROSSA developments are currently under review, one for the HEART conference in 2024 and another one for the FGCS journal.

2.7 PSNC

PSNC: UrbanAir application is able to provide weather forecast, assess air quality in urban environments or support renewable energy sources modelling. Therefore, the audiences to target are:

- Photovoltaic operators and users
- Distribution system operators
- Companies involved in building and/or running wind farms
- Government institutions and environmental bodies

With the recently developed UrbanAir service during the project execution, reaching target audience will be easier. PSNC reached stakeholders by participating to HPC SuperComputer Conference in 2023 and DestinationEarth User Exchange Forum. At national level, PSNC reached Distributions System Operators and photovoltaic operators by participating and giving talks during exhibitions. Direct contact demonstrating results with perspectives for future cooperation was made with one of the larges polish DSO and government bodies in two polish cities.

2.8 INFN

During the first phase of the project, we had the opportunity to introduce TEXTAROSSA and specific INFN achievements in several workshops and international conferences. In particular INFN developments in TEXTAROSSA was presented at *TWEPP international conferences* with a paper [P9] included in the conference proceedings. Furthermore, we gave talks at the *Workshop sul calcolo nell'INFN* (Title "Il progetto TEXTAROSSA"; Paestum (SA) May 23-27, 2022) and at workshop *AI@INFN- Artificial Intelligence at INFN* (Title: "APEIRON: a heterogeneous computing platform for real-time inference", Bologna May 2-3, 2022), introducing the application of TEXTAROSSA findings in INFN internal research activities. The TEXTAROSSA ideas and preliminary results were also introduced at the *NVIDIA roundtable meeting* (Cineca Bologna, September 12, 2022) having the opportunity to disseminate our findings in a mixed academic-

industrial environment. INFN also contributed to the first project-wise joint papers, the proceedings of DSD2021 conference [P34] and to the derived MICPRO Journal paper [P8]. More recently, INFN presented its activities in the project in several international workshops.
Dr. Alessandro Lonardo (INFN Roma) gave a talk with title “APEIRON: composing smart TDAQ systems for high energy physics experiments” at the 21st International Workshop on Advanced Computing and Analysis Techniques in Physics Research - ACAT 22, Bari, Italy, on October 2022; proceedings are about to be published. Cristian Rossi (INFN Roma) gave a talk at the 26th International Conference on Computing in High Energy and Nuclear Physics - CHEP 2023, Norfolk, Virginia (USA), with title “APEIRON: a Framework for High Level Programming of Dataflow Applications on Multi-FPGA Systems” on May 2023; proceedings are about to be published.
Cristian Rossi (INFN Roma) gave a joint talk together with Dr. Paolo Palazzari (ENEA) with title “Multi-FPGA performance scaling of High-Level Synthesis applications through the APEIRON Framework” at the EuroHPC Projects Shaping Europe's HPC Landscape HiPEAC 2024 associated workshop, Munich, Germany, in January 2024.

2.9 CNR

CNR participates to animate the website by news. Furthermore, CNR presented activities developed within TEXTAROSSA at an invited seminar to the dissemination activities of the Computation-based Science and Technology Research Center (CaSToRC) of The Cyprus Institute (2021 Summer seminar series), a plenary invited talk at the Italian National Conference on Scientific Computing and Mathematical Models (SCMM 2022), an invited presentation at the special session on “Large-Scale Models: Numerical Methods, Parallel Computations and Applications” within the 13th International Conference on Large-Scale Scientific Computations (LSSC 2021), a plenary invited talk at the 31st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing PDP 2023, an invited presentation at the special session “Preconditioning and Multilevel Methods, In memory of O. Axelsson ” within the 14th International Conference on Large-Scale Scientific Computations (LSSC 2023). Furthermore, CNR contributed to the proceedings paper for Euromicro Conference on Digital System Design (DSD) 2021 and to the journal extension published by MICPRO. In addition to these joint publications with other partners, CNR wrote a journal paper published on IEEE on Parallel and Distributed System, a journal paper on Software Impacts, and a Proceedings paper on IEEE Proc. of the 31st Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP 2023). Finally, CNR organized a Special Session on “Mathematical software for Computational and Data Science at Extreme scales", within the First Thematic Conference on Emerging Technologies in Computational Science for Industry, Sustainability and Innovation (M2P2023), and a Special Session on “Sparse Linear Solvers for Computational Science at Extreme Scales", within the 10th International Congress on Industrial and Applied Mathematics (ICIAM 2023).

3 Individual Exploitation Plans and Reports

The use of results of the TEXTAROSSA project for both commercial purposes and scientific area was a key objective of the project. Exploitation was embedded in the vision of the project. Hence all partners within the project were aware of and committed to the exploitation of the project results. TEXTAROSSA management structure coordinated the exploitation activities among the partners, in close cooperation with the PTC, from the innovative idea to the market, assuring support to overcome issues and providing the fastest and less critical path.

The exploitation results for the whole project are proposed by the partners in this section, while the complete list of innovation products is available in the next Section 4.

3.1 ENEA

In the second half of the project, ENEA focus shifted towards further enhancing the utilization of high-performance computing (HPC) infrastructures and technologies. Specifically, the efforts were directed to drive innovation, optimize performance, and enhance the accessibility of HPC resources, ultimately advancing the project's objectives and fostering collaborative advancements in high-performance computing. Specifically, the action included the incorporation of emerging technologies, including novel hardware accelerators and specialized processors, into the HPC ecosystem. This entailed evaluating the feasibility and potential benefits of integrating technologies such as field-programmable gate arrays (FPGAs) and other accelerators to augment computational capabilities and energy efficiency. The enhanced HPC capabilities allowed ENEA to exploit its infrastructure in new European and Italian projects devoted to: Italian exascale HPC infrastructure (ICSC), the realization of a HPC in cloud infrastructure (IPCEI Cloud), the European energy oriented center of excellence in HPC (EoCoE-III), HPC for SMEs in a new European Digital Innovation Hub (CETMA-EDITH).

3.2 FHG

The activity in TEXTAROSSA allowed FHG to further develop its Oli&GAS RTM application plus, in collaboration with other partners, to test new technologies like Posit. As outcomes of this activity there will be new scientific publications and inputs for further developed to increase the TRL in new EU proposals. The activities on the development of a 16-bit PPU posit processing unit, with CINI-UNIPI, are planned to be exploited in the EPI project.

3.3 CINI

CINI, as an academic partner, focuses its exploitation strategy on three main aspects:

- Consolidating its positioning as a leading expert in resource management to secure funding for further research activities in future European calls.
- Technology transfer to either established industrial partners or spin-off companies.
- Exploiting the research carried out to improve its offer of tertiary education targeting Master and Doctoral students.

CINI has identified 3 main innovation products, as reported in the Table 2.2.b of the Grant Agreement. For the first product (Mixed-precision compiler) we defined an open-source exploitation strategy, and it has been already licensed under the permissive MIT license: this will allow to further cooperate with both academic and industrial partners for the development of the toolchain. Commercial exploitation can be

still pursued through cooperation agreements with companies to develop support for specific architectural targets. For the remaining two innovation products (Posits Processing Unit and Crypto accelerator for secure HPC services), discussions are ongoing regarding the licensing model as there is opportunity to pursue licensing agreements for hardware components or to release such components as open hardware.

The activity in TEXTAROSSA allowed CINI-UNIPI to further develop its posit technology (C++ Posit library and IPs) as main exploitable result thanks to the development of light posit processing unit (useful for data compression of CNNs) and full posit processing unit (for compression and AI computation) and its integration with RISC-V cores. The developed IPs lead to rise of collaborations with other partners (being used by CINI-POLIMI, INFN, FHG) with publications in journals and top HPC conferences (e.g. ISC2023, HIPEC2023, Conga2023 part of AsiaSiupercomputing2023). This also contributed to enrich PhD and MS thesis offered at CINI-UNIPI, to strength the position of CINI-UNIPI in the European scenario since CINI-UNIPI is entering as a key partner for digital IPs in the most important HPC new proposals, such as RISC-V FPA, EPI2, Italian HPC national center. The activities on the development of a 16-bit PPU posit processing unit, with FHG, are planned to be exploited in the EPI project.

Concerning part of the activities carried out by CINI-POLIMI, based on the final project achievements, a detailed analysis is in progress regarding the possibility of a partnership for the improvement and licensing of a thermal modelling and control strategy tailored for the evaporative cooling. Specific agreements regarding the restriction and confidentiality in the use of the thermal models developed by CINI-POLIMI have been signed between CINI-POLIMI, InQuattro and PSNC.

As a notable early exploitation results of the project, CINI-UNITO purchased the first commercial prototype of a GPU-enabled server with evaporative cooling. The server, based on the TEXTAROSSA technology, is a system provided by E4 with Intel dual-socket plus 4 NVidia H100SXM24 NVIDA GPUs and exploiting the two-phase cooling provided by InQuattro. The prototype has a primary evaporative circuit and a secondary liquid-cooled circuit (then to internal air or external cooling). It is fully instrumented with hundreds of sensors (temperature, pressure, power, etc.) and implements a LabVIEW programmable cooling strategy. The prototype also includes a programmable heater of the secondary to test the machine on different operational settings. This platform was installed at the HPC4AI green datacenter (https://hpc4ai.unito.it/) of CINI-UNITO for deep testing with the existing commercial users and to foster further scientific cooperation with the project partners.

3.4 INRIA

During the first half of the project, INRIA ported StarPU to Xilinx FPGA using the OpenCL interface. However, the performance results for Chameleon and Scalfmm did not meet our expectations, such that we focus to GPU for the remainder of the project. Additionally, we developed a new scheduler in StarPU to enhance execution on heterogeneous architectures. This work was conducted during Hayfa Tayeb's PhD. The outcomes were promising and surpassed the state-of-the-art for irregular applications. We are currently refining this scheduler to incorporate energy considerations into its decision-making process. Although we only have preliminary results at this stage, we anticipate developing a robust scheduler by the end of 2024. We also proposed a new software package, STARONNX https://gitlab.inria.fr/topal/staronnx, which builds on StarPU and ONNX to create an efficient inference server for DNN models. The software is a product of work carried out as part of Jean-Francois David's PhD thesis in TEXTAROSSA. It is still at the prototype stage, but initial results are very encouraging.

3.5 ATOS

During the first period of the project, Atos firstly carried out the pre-conception phase of two blade servers with four GPUs and two CPUs, either equipped with the Intel Ponte Vecchio or the Nvidia Redstone-Next GPU-baseboards. At the time of the beginning of the project, these GPUs were the latest and most powerful available from each manufacturer. After deliberation with the partners of the project, the Nvidia server was selected for the prototype whose conception was made during the second half of the first reporting period.

The second reporting period was dedicated to designing a custom-made frame and supporting InQuattro for the integration of their two-phase cooling technology inside the OpenSequana architecture.
A thermal analysis of the IDV-A prototype with single-phase cooling was done to produce a benchmark to compare to the two-phase cooling technology. During a joint experiment carried out by CINI-POLIMI, InQuattro and Atos, the two-phase system was installed, and measures were made, including the introduction of a hierarchical thermal control encompassing both the flow of the coolant and the DVFS actuators.

The results obtained during this project showed that the two-phase cooling is capable to maintain the temperature set point with the possibility to remove the target of 3.5Kw. Moreover, this innovative technology is for sure promising and several findings on how to optimize the setup to enhance the integration withing the Atos HPC rack have been identified. For this reason, Atos remains open to any collaboration which could lead to the improvement of the cooling capacities. More details can be found in D5.1 and D3.5.

3.6 E4

In the first part of the project, E4 explored a server solution that could meet the constraints defined in the proposal and resulting from the analysis of WP3 and WP5. At the end of this phase, the most suitable design was identified among different commercial proposals. Accelerators based on FPGA technology were selected by working jointly with partners. Many online meetings with them helped define the right model. The selected server based on ARM technology and Xilinx FPGA accelerators was then installed on E4 premises and made available for remote access to the partners. After a first phase of tests, the server was turned off for a short period during the 2023 summer to allow InQuattro to install the two-phase cooling system developed within the project. The server, now equipped with the proper cooling system, returned to be available for the partners during 2023 fall.

The main Key Exploitable Result indicated in the Exploitation Plan, i.e. the development of a two-phase cooling system - designed in collaboration with InQuattro, was reached. The platform is now part of E4 commercial offer and is present among its solutions. E4 is currently preparing a datasheet describing its main features and a brochure to use for marketing purposes. The first commercial system equipped with a two-phase cooling system was sold to the University of Turin and installed at the end of March 2024.

3.7 BSC

BSC has a two-pronged approach to exploit the results of the TEXTAROSSA project. The first big development, the OmpSs@FPGA framework is the main beneficiary of the developments. The framework has been completely revamped to take advantage of the new technologies that have appeared along the project development like OmpSs-2, new CLang compiler versions, VITIS HLS, etc. This has improved the

framework scope and capabilities (including performance and programmability). New features have been added like power monitoring, HBM memory management, etc. Key developments that are actively being used in research and may even make their way to industrial exploitation are the support for IDV-E like platforms (PCIe FPGAs attached to non-x86 host CPUs) and support for clusters of FPGAs (like MEEP or ESSPER FPGA clusters). The latter development, unforeseen at the time of the project proposal introduces support for clusters of FPGAs both with MPI like directives or the new and promising IMP approach developed along the project. Although in early stages of development this TEXTAROSSA contribution is generating several research initiatives regarding its applicability to different distributed memory systems.

The second direction that we are following exploiting the results of the TEXTAROSSA project regards the use of the Fast Task Scheduling IP (and other OmpSs@FPGA ideas like the IMP programming model) in many-core RISC-V nodes and even clusters. These ideas, already used for research and published in high-quality journals are being fed to projects like EPI SGA2, EUPilot to evaluate the possibility of incorporating the IPs into the final designs to strengthen the competitive advantage of the European developed systems.

3.8 PSNC

PSNC exploitation activities focused on UrbanAir application. An online service with weather prediction and air quality forecast was prepared to support different stakeholders. PSNC is now finalising preparation phase for an environmental project, where UrbanAir will be used by Polish city government to adapt to air-quality policy.

UrbanAir application is able predict air quality in urban environments but can also be used in renewable energy sources domain to predict amount of electrical energy produced by wind farms and photovoltaic system, or to assess the probability of damages which may happen to the overhead electrical network infrastructure. The solution is being tested by one of the largest Polish Distribution System Operator to assess the probability of damages. In another EU project, UrbanAir is used to predict amount of energy produced by wind farms and photovoltaic systems. Is it further developed in cooperation with DestinationEarth initiative in a form of digital twin.

3.9 INFN

During the first period of project, the exploitation activities of INFN focused on the integration of our inter-FPGA low-latency communication IP with the HLS toolchains in HEP experimental initiatives. Contacts and dissemination activities towards our reference science community (HEP Experiments) led to the demonstration of a first prototype of a ML architecture FPGA-based for Cherenkov ring identification of the CERN NA62 experiment RICH detector [P9].

Leveraging the developments done in TEXTAROSSA a large part of the INFN team have joined the Electron-Ion Collider experiment, contributing to the design of a data reduction stage for the dual RICH detector based on the APEIRON framework.

In addition, we are exploring the opportunity to licence the communication IPs for industrial exploitation, supported by the INFN Transfer Technology office.

3.10 CNR

CNR hired two young collaborators to implement the Math-lib for parallel sparse matrix computations. Participation to the TEXTAROSSA project was an added value to be invited to collaborate to some

dissemination activities and to lead a WP in a new project proposal for the EuroHPC call for Center of Excellence on Supercomputing Applications.

3.11 InQuattro

During the initial half of the project, InQuattro started the exploration and development of two-phase cooling technology for HPC servers, in collaboration with Atos and E4. This phase primarily focused on both technical solutions—encompassing design and testing as detailed in Deliverable 3.2—and gathering insights into the needs of users and stakeholders, which is crucial for achieving a market-ready solution.

The project engages a diverse group of users and stakeholders, including E4 and Atos as potential customers of the cooling systems, and several universities and research institutions as end-users of the hardware. Through direct dialogues with Atos and E4, we have outlined the cooling system requirements based on the specified hardware. These discussions are important not only in capturing the technical specifications but also in understanding the underlying reasons for these requirements.

A similar approach was adopted to grasp the perspectives of the final users of the HPC systems. This was achieved through direct interactions with some partners and at the ISC event in Hamburg in June 2022. In addition, our outreach has expanded through additional contacts with various operators, effectively broadening our network of potential customers and stakeholders. These interactions have provided a clearer understanding of the market dynamics specific to HPC data centers. The insights gained from these engagements have led to the development of the initial prototypes of the two-phase cooling system for HPC servers.

As we moved in the second half of the project, with the cooling systems operational on the servers, we have continued this engagement process to gather feedback, suggestions, and constructive criticism. The aim has been to continuously refine and enhance the design of the two-phase cooling system prototypes, ensuring they effectively meet user and market needs, that allowed InQuattro to achieve the goal of a setup ready for commercialization.

At the end of the project, the development of our prototypes culminated in a significant achievement: the University of Turin placed an order to install our innovative two-phase flow cooling system in new GPU servers with 4 Nvidia H100 GPUs. This system was successfully delivered and installed in the University's data center by the end of March 2024. This installation marks an important milestone in the market development of our two-phase cooling solution.

4 Products and IPs

The next sections, one for each partner, shows a characterization table for main products/IPs developed or improved in the TEXTAROSSA project. These tables identify the innovations and exploitable results of the project.

4.1 ENEA

Innovation	FPGA ImgLib		
Problem	How to perform high-performance image processing on FPGA.		
Alternative solutions	Xf::OpenCV		
Innovation of your solution	Kernels are optimized for variable bithwidths, so can be easily adapted to the available I/O bandwidth.		
Description	Implemented several operators for image and stream processing (color space conversion, splitting and merging of image components, FIR filtering, median filtering, histogram equalization, various image generation and mixing, single pixel transformation, rotation, resizing, zooming, stream splitting, stream copying, stream merging, …).		
IPR & License	Open Source		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	0	4	4
Availability	Will be released as open source through github after the end of the project		
Reference Person	Paolo Palazzari <paolo.palazzari@enea.it>		
Possible market	Image processing		

4.2 CINI

Innovation	TAFFO for CUDA and OpenCL
Problem	Support Mixed-Precision computing on heterogeneous platforms leveraging GPGPUs, programmed using either OpenCL or CUDA.
Alternative solutions	Currently, none of the frameworks for precision tuning can directly target GPGPU programming models. See [R1] for a recent survey on such tools.

Innovation of your solution	The proposed solution enables automated management of computation precision, thus reducing the burden on the application developer.		
Description	TAFFO is a set of plugins for the LLVM compiler framework to perform precision tuning (both in terms of short floating point data types and of integer/fixed point data types). The innovation provided by TEXTAROSSA enables the plugins on the main commercial GPGPUs by supporting CUDA and OpenCL. TAFFO has been also extended to support VItisHLS, enabling the support of reconfigurable fabrics, and the Posit number system.		
IPR & License	CINI-POLIMI, permissive open source (MIT)		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	5	5
Availability	- Publicly available: https://github.com/TAFFO-org - Deliverable D4.3		
Reference Person	Daniele Cattaneo <daniele.cattaneo@polimi.it>		
Possible market	Developers of HPC and edge applications for heterogeneous systems		

Innovation	IP security acceleration		
Problem	Advanced secure solutions exploiting new schemes such homomorphic encryption beyond classical cryptography (AES, SHA, ECC, RSA, ...)		
Alternative solutions	Mainly SW solutions are available, but HW acceleration needed for computing intensive part.		
Innovation of your solution	HW acceleration of computing intensive secure kernels like SHAKE and NTT		
Description	See deliverable D2.4, D2.5		
IPR & License	CINI-UNIPI, open source licence for research use		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**

	3	5	5
Availability	- SEAL embedded IP repository - https://drive.google.com/drive/folders/1UxmG3t8YLPV6oRAi0B9zwMZK13cAjAET?usp=sharing - SHAKE IP repository - https://drive.google.com/drive/folders/1V07qDMeCX8IdoJrVOXxuQuj82D_gl_a-?usp=sharing. - Deliverable D2.4, D2.5		
Reference Person	Sergio Saponara <sergio.saponara@unipi.it>		
Possible market	Developers of secure HPC and edge applications for heterogeneous systems		

Innovation	IPs for Light and Full PPU units		
Problem	Implements Posit arithmetic and use Posits for computation and storage to provide same accuracy of floats but with reduced data-size		
Alternative solutions	Market is looking at several available alternatives of floats but there is no a clear winner. Posits can be a solution with companies worldwide already interested in them, such as VividSparks in Asia		
Innovation of your solution	Efficient implementation of Posit operations (add, sub, mul, divide, fused multiply-add) and of transform Posits-to/from-Floats, with IPs used and tested by many partners, PLIMI; INNFN; Fraunhofer		
Description	See Deliverable D2.2, D2.3, D2.6, D2.7		
IPR & License	CINI-UNIPI, open source licence for research use		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	5	5
Availability	- The IP called Full PPU IP has a repository (where it is designed in SystemVerilog) in https://github.com/federicorossifr/ppu_public/tree/ppu2 - The IP is also available in a data-based ready to be synthetised by VITIS High Level Synthesis tool https://github.com/federicorossifr/ppu_public/tree/pipeline/vitis - Deliverable D2.2, D2.3, D2.6, D2.7		
Reference Person	Sergio Saponara <sergio.saponara@unipi.it>		
Possible market	Developers of HPC and edge applications for heterogeneous systems		

Innovation	First production deployment of a CPU+GPU server with a two-phase cooling system		
Problem	Power saving in existing HPC datacentres (retrofit).		
Alternative solutions	Today the market for two-phase cooling is very small, the only other player selling servers with two-phase cooling (zutacore) does not support CPU+GPU servers.		
Innovation of your solution	The solution is suitable for cooling servers with GPUs; to our knowledge, there are no other solutions on the market. We plan to demonstrate that it is possible to run servers with full-power Nvidia H100 GPUs (700W per GPUs) on air-cooled data centres, which can be crucial for small-medium-sized data centres.		
Description	Supermicro server with 2 Intel sockets and 4 Nvidia H100 SXM modified (by E4) with a 2-phase cooling system from InQuattro. The system has been integrated within the CINI-UNITO's cloud-HPC system (https://hpc4ai.unito.it/documentation/) and made available via CINI-UNITO's integrated cloud-HPC management system based on the StreamFlow cloud-HPC management system (https://streamflow.di.unito.it) and Jupyter-workflow (https://jupyter-workflow.di.unito.it). StreamFlow and Jupyter-workflow development has been funded by EuroHPC project. Streamflow was selected in April 2023 by EU Innovation Radar program (https://www.innoradar.eu/innovation/49626).		
IPR & License	Solution commissioned and funded by CINI-UNITO, integrated by E4 with cooling system made by InQuattro.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	N/A	8/9	9
Availability	http://hpc4ai.unito.it		
Reference Person	Marco Aldinucci <aldinuc@di.unito.it>		
Possible market	Existing cloud-HPC datacenters (120-250KW)		

4.3 INRIA

Innovation	Use of FPGA in StarPU
Problem	How FPGA can be efficiently used in a dynamic task-based runtime system
Alternative solutions	OmpSs@FPGA (also exploited in TEXTAROSSA)

Innovation of your solution	Technical upgrade of StarPU.		
Description	FPGA are managed as any other PU and thus can compute tasks provided by the users.		
IPR & License	LGPL-2.1 license (same of StarPU)		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	5	5
Availability	Publicly available online: https://starpu.gitlabpages.inria.fr/		
Reference Persons	Hayfa Tayeb <hayfa.taybe@inria.fr>, Bérenger Bramas <berenger.bramas@inria.fr>		
Possible market	No market directly, but through the use of StarPU.		

Innovation	Multreeprio scheduler in StarPU		
Problem	Scheduling of tasks over heterogeneous processing units.		
Alternative solutions	Heft and dmda are the references.		
Innovation of your solution	Provide speedup for several test cases, do not need input from the user. Evaluation in progress.		
Description	Our scheduler has an efficient multi-binary tree that allows to set several priorities to the tasks.		
IPR & License	LGPL-2.1 license (same of StarPU)		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	5	5
Availability	Publicly available online: https://starpu.gitlabpages.inria.fr/		
Reference Persons	Hayfa Tayeb <hayfa.taybe@inria.fr>, Bérenger Bramas <berenger.bramas@inria.fr>		
Possible market	No market directly, but through the use of StarPU.		

Innovation	StarONNX inference framework		
Problem	Performing inference requests over heterogeneous architectures		
Alternative solutions	NVIDIA TritonServer has similar capabilities		
Innovation of your solution	More efficient resource usage, asynchronous data management, dynamic scheduling		
Description	Our framework combines the efficient kernel implementations of ONNX Runtime with the dynamic runtime system StarPU.		
IPR & License	GPL-2.1 license		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	1	3	3
Availability	Publicly available online: https://gitlab.inria.fr/topal/staronnx		
Reference Persons	Jean-Francois David <jean-francois.david@inria.fr>, Lionel Eyraud-Dubois <lionel.eyraud-dubois@inria.fr>, Olivier Beaumont <olivier.beaumont@inria.fr>		
Possible market	AI users with need of efficient inference on HPC servers or embedded architectures.		

4.4 ATOS

Innovation	Two-phase cooling in Atos HPC Rack
Problem	Blade server cooling
Alternative solutions	Possibility to cool down the HPC blade server with different other solutions: - Air-cooling with forced airflow through a heat sink mounted on the processor. - Single-phase liquid inside a micro channel cold plate which is in contact with the processor through a thermal interface. - Single-phase liquid with micro jets directly spread on the surface of the processor. - Immersion cooling with single-phase liquid that directly cools down the processor.
Innovation of your solution	Better thermal performances to be able to improve the thermal resistance of the cooling solution to support higher TDP (Thermal Design Power) or lower limited temperature processors with still warm water at the inlet (optimization of the datacenter PUE).

Description	The solution consists of using a two-phase liquid inside a micro channel cold plates in contact with the processor through a thermal interface. These cold plates are integrated into our blade server module with the required associated components for this two-phase loop (tank, exchanger, monitoring and control devices) and interfaced with the current single-phase loop of the SequanaXH3000 HPC cabinet through a heat exchanger.		
IPR & License	- InQuattro and Atos IP - Commercial license - Eventual hardware patent at blade level		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	0	4	6
Availability	Contact the Atos sales team to get more information (https://atos.net/en/contact-high-performance-computing)		
Reference	Fabien Demange <fabien.demange@atos.net>		
Possible market	HPC market		

4.5 E4

Innovation	Efficient, ecofriendly compute platform
Problem	Current computing platforms require a significant amount of power and imply a related significant amount of emission of GHG.
Alternative solutions	One solution is reducing the frequency of the processor, but performance decreases more than linearly with power consumption while the heat released is not significantly affected.
Innovation of your solution	Phase-change, pumped cooling decreases the temperature of processors and accelerators, enabling it to run at target frequency – and possibly at turbo mode. As of today, TEXTAROSSA is the only project developing the phase-change technology on ARM platforms.
Description	One of the Integrated Development Vehicle (IDV-E) is co-designed encompassing the technology regarding the two-phase cooling and the thermal monitoring and control, beside the presence of heterogeneous computing nodes (CPUs plus top-class FPGA) to accommodate several programming paradigms. The outcome of such a co-design activity will be a node/blade well suited for the HPC, AI and HPDA markets, featuring a match between high level of performance (provided by the compound CPU+accelerator(s)) and a significant reduction in power consumption and related thermal requirements (because of the two-phase thermal management solution), that will become part of the E4 product portfolio.
IPR & License	E4 is the owner of the integration technologies; licensing policy will be defined at a later time.

TRL	At the beginning of the project	Target goal at the end of the project	Achieved at the end of the project
	3	6/7	7
Availability	The product is proposed as a customized solution to E4's customers. A webpage dedicated to the product is currently under preparation for E4 webpage.		
Reference	Cosimo Gianfreda <cosimo.gianfreda@e4company.com>		
Possible market	Any markets using computing platforms and requiring a reduction of their TCO and GHG impact.		

4.6 BSC

Innovation	Fast Task Scheduler (FTS)		
Problem	When sending tasks to the FPGAs, fast task scheduling is necessary to tackle large amounts of small to medium size tasks. Task size is limited by the necessity of obtaining parallelism and FPGA capacity so in a task-based model, task scheduling directly impacts full-system performance. In available solutions the slow scheduling of tasks results in a performance loss that hinders the usefulness of the model. The model has been extended to also address fast task scheduling in many-core SMPs where the large number of cores needs fast scheduling to keep the computing units fed.		
Alternative solutions	All other systems that integrate FPGAs and CPUs (like Vitis) have their own task scheduler (like the Xilinx Runtime Library, XRT) but it usually resides in the host CPU instead of the FPGA hardware. Even the parts that are in hardware usually are not optimized for speed as the FTS is. If we talk about the many-core SMPs scenario, the alternative solutions are typically software ones that are not competitive against our proposed hardware one.		
Innovation of your solution	Allows to deploy a large set of tasks in the computing resources (FPGA, accelerators or SMPs) very fast to keep pace with a growing number of computing units. From our research, this is critical for performance in task-based environments.		
Description	Our solution is a hardware IP that on one side is connected to the software running on the host and on the other to the hardware runtime (it can be connected directly to the accelerators or used in conjunction with task managers like PSC proprietary IP PICOS). The IP ensures that tasks are deployed as fast as possible to the accelerators moving several synchronization mechanisms to inside the hardware.		
IPR & License	The FTS is open-sourced (LGPSv3) as part of the OmpSs@FPGA runtime.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	6	6
Availability	FTS is publicly available: https://github.com/bsc-pm-ompss-at-fpga		

Reference	Carlos Álvarez <carlos.alvarez@upc.edu>
Possible market	FPGA companies. As ASIC inside a multiprocessor, Processor Design Companies and processor design projects as EPI.

Innovation	OmpSs@FPGA over IDV-E		
Problem	Execute applications using the heterogenous resources of the IDV-E platform (FPGA).		
Alternative solutions	Currently there are no alternative solutions to the problem. We are working to develop (unofficial) support to Xilinx Vitis Flow in the same platform using the same approach as OmpSs@FPGA.		
Innovation of your solution	There is no alternative solution currently on the market that supports standalone FPGAs (connected through PCIe) with an ARM CPU host.		
Description	Our solution supports work offloading from a CPU host computer to a standalone FPGA (PCIe) in a way that is transparent to the programmer. This improves portability (as the same program targeting a x86 host can be used transparently to the programmer), programmability (by using a High-Level Programming Model like OmpSs) and even performance (as performance results from the model are competitive with the state-of-the-art results for the same platform).		
IPR & License	The OmpSs@FPGA over IDV-E is going to be open-sourced (LGPSv3) as part of the OmpSs@FPGA runtime.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	2	6	6
Availability	Publicly available: https://github.com/bsc-pm-ompss-at-fpga		
Reference	Carlos Álvarez <carlos.alvarez@upc.edu>		
Possible market	FPGA companies & FPGA users/developers.		

Innovation	IMP for multi-node task/stream programming
Problem	Execute task-based applications in multi-node platforms leveraging task-based and stream-based solutions. The initial proposal was to only explore the possibilities of this approach, but the good results obtained encouraged us to continue development of the system.
Alternative solutions	Currently there are only independent solutions to different parts of the problem. Multi-node applications use OpenMP+MPI. Some initial tests have tried to apply this solution to heterogeneous systems, but they rely on a host-accelerator organization that limits scalability. Also, some models (like OmpSs@cluster) try to apply task-

	based programming models to distributed memory systems, but their scalability is limited (usually up to 16 nodes) and only apply to homogeneous SMP systems.
Innovation of your solution	Currently, there is no alternative solution on the market that supports programming heterogeneous distributed systems with a task-based programming model. In addition, achieving good scalability and performance is a bonus only shown by our early tests in TEXTAROSSA.
Description	Our solution supports programming a cluster of distributed heterogeneous resources with a task-based programming model in a way that is transparent to the programmer. This improves portability (as the same program targeting a x86 host can be used transparently to the programmer), programmability (by using a High-Level Programming Model like OmpSs) and even performance (as performance results from the model are competitive with the state-of-the-art results for other platforms).
IPR & License	The IMP programming model is going to be open-sourced (LGPSv3) as part of the OmpSs@FPGA runtime.

TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	0	2 (tentative)	2

Availability	Publicly available: https://github.com/bsc-pm-ompss-at-fpga
Reference	Carlos Álvarez <carlos.alvarez@upc.edu>
Possible market	FPGA companies, FPGA users/developers, programmer community in general.

4.7 PSNC

Innovation	Liquid-based cooling models for DCworms
Problem	Simulations aimed at comparison of thermal and energy efficiency of the PSNC server room after changing to IDV-A nodes with two-phase cooling. Experiments aimed towards Exascale simulations evaluating the performance of servers and applications developed in the TEXTAROSSA project.
Alternative solutions	N/A
Innovation of your solution	Adaptation of thermal models of liquid-based cooling in DCworms.
Description	DCworms is designed as an object-oriented, plugin-based, event-driven simulator. Thus, it provides easy extension capabilities that allow us to plug in workload and resource management policies, as well as to integrate the corresponding energy, thermal and performance models.

IPR & License	PSNC, open source		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	7	6
Availability	Publicly available: https://git.man.poznan.pl/stash/projects/WORMS/repos/dcworms		
Reference	Sebastian Ciesielski <sciesielski@man.poznan.pl>		
Possible market	Developers of HPC and edge applications for heterogeneous systems.		

Innovation	Accelerating UrbanAir with GCRK routine		
Problem	How to improve GCRK/UrbanAir in terms of performance and energy efficiency		
Alternative solutions	N/A		
Innovation of your solution	Technical upgrade to GCRK/UrbanAir.		
Description	Improvements by adapting to the air quality case, exploiting multi-gpus and studying mixed-precision.		
IPR & License	Bilateral agreement (contact kulka@man.poznan.pl for further details)		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	2	4	4
Availability	Contact kulka@man.poznan.pl		
Reference	Michał Kulczewski <kulka@man.poznan.pl>		
Possible market	Environmental protection, Earth sciences, renewable energy sources		

4.8 INFN

Innovation	Communication IP
Problem	Enabling low-latency communication between HLS kernels on the same FPGA (Intra-node communication) and on different FPGAs (Inter-Node communication).

Alternative solutions	For inter-node communication A. Use the host to perform communication (HLS kernel communicate data to host). B. Directly interface the HLS kernel to a network channel implemented using FPGA resources.		
Innovation of your solution	- With respect to A: much lower latency and energy consumption. - With respect to B: concurrent access to the network channels by HLS kernels, flexible partitioning of available ports between I/O and communication network channels, no need for external network switch, lower latency and energy consumption.		
Description	Our solution allows communication between HLS kernels deployed on FPGAs without involving the host and system bus resources, ensuring low latency and reducing energy consumption. It is also able to manage parallel data flow, solving contentions for shared resources.		
IPR & License	INFN, non-exclusive licensing for source code version of the IP.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	3	6	6
Availability	- The IP is available in form of compiled object file (.xo) in the APEIRON framework git repository (https://github.com/APE-group/Textarossa). - INFN TT service will address requests from potential customers for dedicated developments.		
Reference	Alessandro Lonardo <alessandro.lonardo@roma1.infn.it>		
Possible market	HPDA (real-time high throughput data analysis, also based on NN inference), Government Research, Security, BFSI, Energy & Utilities		

4.9 CNR

Innovation	BootCMatchGX
Problem	Solution of Large and Sparse Algebraic Linear Systems
Alternative solutions	Main alternative solutions is Nvidia AmgX library
Innovation of your solution	The library leverages on an innovative Algebraic MultiGrid Method and parallel design patterns which demonstrated their benefits in terms of robustness, efficiency and scalability when compared with the alternative solutions.
Description	BootCMatchGX is a library for Nvidia multi-GPU systems. Sparse solvers are one of the building blocks of any software technology for reliable and high-performance

	scientific and engineering computing. In BootCMatchGX we make available an Algebraic MultiGrid method for preconditioning algebraic linear systems Ax = b, where A is a symmetric positive definite (s.p.d.), large and sparse matrix. All the computational kernels for setup and application of the adaptive AMG method, as preconditioner of an efficient version of the Conjugate Gradient Krylov solver, were designed and tuned for hybrid MPI-CUDA programming environments when multiple distributed nodes hosting Nvidia GPUs are available.		
IPR & License	CNR, open source MIT license.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	2	4	4
Availability	Publicly available: https://github.com/bootcmatch/BootCMatchGX		
Reference	Massimo Bernaschi <massimo.bernaschi@cnr.it>, Pasqua D'Ambra <pasqua.dambra@cnr.it>		
Possible market	Computational Scientists and Software Developers.		

4.10 InQuattro

Innovation	Evaporative Liquid Cooling		
Problem	Thermal management of HPC data center with high power processors (GPU, CPU) using free cooling in hot summer season (with external temperature of 40°C).		
Alternative solutions	Air Cooling, liquid cooling, immersion cooling.		
Innovation of your solution	The new cooling technology, helps data centers to increase processing power while using free cooling in hot climates (up to 40°C), and significantly less energy and space than conventional cooling systems. Besides, the new cooling technology is able to cool high power processors up to the boundary of 1000 W and beyond.		
Description	The innovative feature of this system is the use of flow boiling heat transfer for cooling electronic devices. Compared to traditional cooling systems (liquid cooling, heat pipes), significantly higher heat transfer coefficients can be achieved at significantly low flow rates and pumping power.		
IPR & License	Patent owners: Francesco Romanello, Antonio Scotini, Luca Saraceno, Giuseppe Zummo. In Quattro has the patent licence.		
TRL	**At the beginning of the project**	**Target goal at the end of the project**	**Achieved at the end of the project**
	4	7	7
Availability	www.in-quattro.com		

Reference	Giuseppe Zummo <g.zummo@in-quattro.com>
Possible market	Data center Cooling

5 Conclusions

This deliverable described the dissemination activities and exploitation results obtained by the consortium during the 36 months of the project. The consortium fully achieved the goals for scientific publications expected at the end of the project. Partners presented the project in several scientific and general public events, as highlighted from Section 1.1 to Section 1.4. Section 1.5 described the performance of the website and social media, highlighting the actions undertaken to maintain a good level of dissemination in such platforms. Section 2 reported the individual dissemination activities and the updated dissemination plans, while Section 3 contains individual plans and reports for the exploitation activities.

Notably, a set of HPC nodes has been purchased by CINI-UNITO to provide computing services, as well as to keep on experimenting on the TEXTAROSSA technologies, especially the 2-phase cooling, also beyond the project timeframe.

6 References

[R1] 2020, Cherubin, Stefano, and Giovanni Agosta. "Tools for reduced precision computation: a survey." *ACM Computing Surveys (CSUR)* 53.2 (2020): 1-35.

[R2] Xunyi Zhao, Théotime Le Hellard, Lionel Eyraud–Dubois, Julia Gusak, Olivier Beaumont. Rockmate: an Efficient, Fast, Automatic and Generic Tool for Re–materialization in PyTorch. ICML 2023 – 40 th International Conference on Machine Learning, Jul 2023, Honolulu (HI), United States. ⟨hal–04095305v5⟩

[R3] Olivier Beaumont, Lionel Eyraud–Dubois, Alena Shilova. Pipelined Model Parallelism: Complexity Results and Memory Considerations. Europar 2021, Aug 2021, Lisbon, Portugal.

[R4] Olivier Beaumont, Philippe Duchon, Lionel Eyraud–Dubois, Julien Langou, Mathieu Vérité. Symmetric Block–Cyclic Distribution: Fewer Communications Leads to Faster Dense Cholesky Factorization. SC 2022 – Supercomputing, Nov 2022, Dallas, Texas, United States.

[R5] Olivier Beaumont, Lionel Eyraud–Dubois, Mathieu Vérité, Julien Langou. I/O–Optimal Algorithms for Symmetric Linear Algebra Kernels. ACM Symposium on Parallelism in Algorithms and Architectures, Association for Computing Machinery : SIGACT, SIGARCH, Jul 2022, Philadelphie, United States.

Appendix 1: List of publications

Journal articles

[P1] 2023, D. Zoni, et al. A Survey on Run-time Power Monitors at the Edge. ACM Comput. Surv. 55, 14s, Article 325 (December 2023), 33 pages. (doi:10.1145/3593044, open access).

[P2] 2023, L. Morais et al., Enabling HW-based Task Scheduling in Large Multicore Architectures, in IEEE Transactions on Computers, 2023. (doi:10.1109/TC.2023.3323781)

[P3] 2023, S. D. Matteo et al., VLSI Design and FPGA Implementation of an NTT Hardware Accelerator for Homomorphic SEAL-Embedded Library, IEEE Access, vol. 11, pp. 72498-72508, 2023 (doi:10.1109/aldACCESS.2023.3295245 – open access).

[P4] 2023, M. Bernaschi, et al. A Multi-GPU Aggregation-Based AMG Preconditioner for Iterative Linear Solvers. IEEE Transactions on Parallel and Distributed Systems, vol. 34, no. 8, August 2023. (doi:10.1109/TPDS.2023.3287238, open access).

[P5] 2022, P. D'Ambra, et al. Parallel Sparse Computation Toolkit. Software Impacts, Volume 15, March 2023, 100463. (doi:10.1016/j.simpa.2022.100463 (open access))

[P6] 2022, D. Cattaneo, et al. TAFFO: The compiler-based precision tuner. SoftwareX. Volume 20, December 2022, 101238. (doi:10.1016/j.softx.2022.101238 (open access))

[P7] 2022, D. Zoni, et al. Cost-effective fixed-point hardware support for RISC-V embedded systems. Journal of Systems Architecture, Volume 126, May 2022, 102476. (doi:10.1016/j.sysarc.2022.102476, open access)

[P8] 2022, G. Agosta, et al. Towards EXtreme scale technologies and accelerators for euROhpc hw/Sw supercomputing applications for exascale: The TEXTAROSSA approach. In Microprocessors and Microsystems, 104679. (doi:10.1016/j.micpro.2022.104679, open access)

[P9] 2022, R. Ammendola, et al. Progress report on the online processing upgrade at the NA62 experiment. Journal of Instrumentation, Volume 17, April 2022. (doi:10.1088/1748-0221/17/04/C04002, open access)

[P10] 2021, M. Cococcioni, et al. A Lightweight Posit Processing Unit for RISC-V Processors in Deep Neural Network Applications. IEEE Transactions on Emerging Topics in Computing, In press. (doi:10.1109/TETC.2021.3120538, open access)

Conference papers

[P11] 2024, H. Tayeb et al., Dynamic Tasks Scheduler with Multiple Priority-based Trees on Heterogeneous Computing Systems, IEEE Heterogeneity in Computing Workshop (HCW'24), IPDPS 2024. (open access).

[P12] 2023, A. Filgueras, et al. Improving performance of hpc kernels on fpgas using high-level resource management. IEEE 31st Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM). IEEE, 2023.

[P13] 2023, A. Filgueras, et al. FPGA Framework Improvements for HPC Applications. 2023 International Conference on Field Programmable Technology (ICFPT). IEEE, 2023.

[P14] 2023, A. Ottimo et al. "FSP: a framework for data stream processing applications targeting FPGAs." 2023 31st Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP). IEEE, 2023.

[P15] 2023, W. Fornaciari et al., Hardware and Software Support for Mixed Precision Computing: a Roadmap for Embedded and HPC Systems, 2023 Design, Automation & Test in Europe Conference & Exhibition (DATE), Antwerp, Belgium, 2023, pp. 1-6 (doi:10.23919/DATE56975.2023.10137092, open access).

[P16] 2023, A. Ghiassi et al., Robust Learning via Golden Symmetric Loss of (un)Trusted Labels, SDM

'23: SIAM International Conference on Data Mining, pp. 568–576 (doi:10.1137/1.9781611977653.ch64, open access).

[P17] 2023, A. Galimberti et al., An Evaluation of the State-Of-The-Art Software and Hardware Implementations of BIKE, 14th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 12th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms. Open Access Series in Informatics (OASIcs), Volume 107, pp. 4:1-4:12, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2023) (doi:10.4230/OASIcs.PARMA-DITAM.2023.4, open access).

[P18] 2023, M. Piccoli et al., Dynamic Power consumption of the Full Posit Processing Unit: Analysis and Experiments, PARMA-DITAM 2023. (doi:10.4230/oasics.parma-ditam.2023.6, open access).

[P19] 2023, N. Neves et al., An FPGA-based platform to evaluate Posit arithmetic in next-generation processors, ISC High Performance 2023 (link).

[P20] 2023, O. Beaumont et al., Data Distribution Schemes for Dense Linear Algebra Factorizations on Any Number of Nodes, IEEE International Parallel & Distributed Processing Symposium, In publishing (open access)

[P21] 2023, P. D'Ambra, F. Durastante, S. M. Ferdous, M. Halappanavar, S. Filippone, A.Pothen, AMG Preconditioners based on Parallel Hybrid Coarsening and Multi-objective Graph Matching, Proc. of the 31th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), 2023. (doi, open access)

[P22] 2023, M. Danelutto, et al., FastFlow targeting FPGAs, In Publishing, PDP 2023 – 31st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing. March 1-3, Naples, Italy (doi, open access)

[P23] 2022, G. Sansone, et al. Experiments on Speeding Up the Recursive Fast Fourier Transform by using AVX-512 SIMD instructions. In Publishing, ApplePies'22 – International Conference on Applications in Electronics Pervading Industry, Environment and Society. (open access)

[P24] 2022, G. Montanaro, et al. Hardware-Software Co-Design of BIKE with HLS-Generated Accelerators, 2022 29th IEEE International Conference on Electronics, Circuits and Systems (ICECS), Glasgow, United Kingdom, 2022, pp. 1-4. (doi:10.1109/ICECS202256217.2022.9970992, open access)

[P25] 2022, M. Cococcioni, et al. Small Reals Representations for Deep Learning at the Edge: A Comparison. Conference on Next Generation Arithmetic CoNGA 2022 pp 117-133. (doi:10.1007/978-3-031-09779-9_8, open access)

[P26] 2022, M. Cococcioni, et al. Experimental Results of Vectorized Posit-Based DNNs on a Real ARM SVE High Performance Computing Machine. ApplePies 2021: Applications in Electronics Pervading Industry, Environment and Society pp 61-68. (doi:10.1007/978-3-030-95498-7_9, open access)

[P27] 2022, H. Tayeb, et al. MulTreePrio: Scheduling taskbased applications for heterogeneous computing systems. COMPAS 2022 – Conférence francophone d'informatique en Parallélisme, Architecture et Système, Jul 2022, Amiens, France. (available online)

[P28] 2022, D. Galli, et al. On the Effectiveness of True Random Number Generators Implemented on FPGAs. In Lecture Notes in Computer Science book series (LNCS, volume 13511), Springer. (doi:10.1007/978-3-031-15074-6_20, open access)

[P29] 2022, A. Galimberti, et al. On the use of hardware accelerators in QC-MDPC code-based cryptography. In Proceedings of the 19th ACM International Conference on Computing Frontiers (CF '22). Association for Computing Machinery, New York, NY, USA, 193–194.

(doi:10.1145/3528416.3530243, open access)

[P30] 2022, D. Cattaneo, et al. Ahead-Of-Real-Time (ART): A Methodology for Static Reduction of Worst-Case Execution Time. In Third Workshop on Next Generation Real-Time Embedded Systems (NG-RES 2022), Schloss Dagstuhl–Leibniz-Zentrum für Informatik. (doi:10.4230/OASIcs.NG-RES.2022.4 open access)

[P31] 2022, Morais, Lucas; Jiménez-González, Daniel; Álvarez, Carlos. Exploring Task Scheduling Sensitivity to L1 Cache configuration on a 32-core RISC-V Processor. In BSC International Doctoral Symposium - 9th BSC Doctoral Symposium, 10th, 11th and 12th May 2022 (http://hdl.handle.net/2117/383675, open access).

[P32] 2022, Filgueras Izquierdo, Antonio; Alvarez, Carlos; Jiménez González, Daniel. Improving resource usage in large FPGA accelerators. In BSC International Doctoral Symposium - 9th BSC Doctoral Symposium, 10th, 11th and 12th May 2022 (http://hdl.handle.net/2117/383418, open access).

[P33] 2022, W. Fornaciari, et al. The TEXTAROSSA Approach to Thermal Control of Future HPC Systems. In 2022 SAMOS Conference. July 2022. Lecture Notes in Computer Science, vol 13511. Springer, Cham. (doi:10.1007/978-3-031-15074-6_27, open access)

[P34] 2022, P. D'Ambra, et al. AMG Preconditioners for Computational and Data Science at Extreme Scale, Calcolo Scientifico e Modelli Matematici, Rome, April 8, 2022 (open access).

[P35] 2021, M. Aldinucci, et al. The Italian research on HPC key technologies across EuroHPC. In ACM Computing Frontiers, Virtual Conference, May 2021. (doi:10.1145/3457388.3458508, open access)

[P36] 2021, G. Agosta, et al. TEXTAROSSA: Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale. In 2021 24th Euromicro Conference on Digital System Design (DSD) (pp. 286-294). IEEE. (doi:10.1109/DSD53832.2021.00051, open access)

[P37] 2021, P. D'Ambra, et al. Scalable preconditioners for Computational Science at extreme scale, 13th International Conference on Large-Scale Scientific Computations, Sozopol (BG), June 9, 2021. (open access)

Presentations

[P38] 2024, C. Alvarez. New scheduling strategies for task-based (hardware) runtimes. In 16th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop, Kobe, Japan, April 2023.

[P39] 2023, C. Alvarez. Extracting performance out of task-based FPGA programs. HEART 2023: 13th International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies, Lake Biwa, Japan, June 2023.

[P40] 2023, Carlos Álvarez. Improving resource usage in large FPGA Accelerators. In 15th Joint Laboratory for Extreme-Scale Computing (JLESC) Workshop, INRIA Bordeaux, France, March 2023.

[P41] 2022, M. Bernaschi et al. BootCMatchGX: a scalable iterative linear solver for multi-GPU systems. LN HPCKTT General Assembly and 3rd Italian Workshop on HPC, Turin, Italy. September 2022.

[P42] 2022, W. Fornaciari. Design of secure power monitors for accelerators, by exploiting ML techniques, in the Euro-HPC TEXTAROSSA project. SCADL WS, Co-located with IPDPS 2022. June 2022.

[P43] 2022, G. Zummo. Innovative Two-Phase Cooling Solutions for the Exascale Computing Systems. ISC High Performance 2022. May 2022. Slides

[P44] 2022, A. Lonardo. TEXTAROSSA presentation at the annual INFN Workshop on Computing. May

2022

[P45] 2022, F. Simula. Distributed and Plastic Spiking Neural Network model of the brain cortex behavior. PSNC TEXTAROSSA Internal Seminars Series. February 2022.

[P46] 2021, P. D'Ambra, et al. Algebraic MultiGrid Preconditioners for Sparse Linear Solvers at Extreme Scales on Hybrid Architectures, EuroCC seminar series, July 2021 – CaSToRC: HPC National Competence Center in Cyprus. (slides)