E-Infrastructures
H2020- INFRAEDI-2018-2020

INFRAEDI-01-2018: Pan-European High Performance Computing infrastructure and services (PRACE)

PRACE-6IP

PRACE Sixth Implementation Phase Project

Grant Agreement Number: INFRAEDI-823767

D7.1

Periodic Report on Results from Projects Supported by Applications Enabling and Porting Services
Final

Version: 1.0
Author(s): Gabriel Hautreux, CINES; Chris Johnson, EPCC; Sebastian Lührs, JUELICH
Date: 16.10.2020

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: INFRAEDI-823767**	
	Project Title: PRACE Sixth Implementation Phase Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D7.1	
	Deliverable Nature: Report	
	Dissemination Level: PU*	**Contractual Date of Delivery:** 30 / October / 2020
		Actual Date of Delivery: 30 / October / 2020
	EC Project Officer: Leonardo Flores Añover	

* - The dissemination level are indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2005/444/EC.

Document Control Sheet

Document	**Title: Periodic Report on Results from Projects Supported by Applications Enabling and Porting Services**	
	ID: D7.1	
	Version: <1.0>	**Status:** ***Final***
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2016	
	File(s): D7.1.docx	
Authorship	**Written by:**	Gabriel Hautreux, CINES; Chris Johnson, EPCC; Sebastian Lührs, JUELICH
	Contributors:	Adam Wlodarczyk, WCSS Ahmet Tuncer Durak, UHeM Anastasiia Shamakina, HLRS Andrew Emerson, CINECA Andrew Sunderland, STFC Claudio Domenico Arlandini, CINECA Cristian Morales, BSC Damian Podareanu, SURFSARA Engelbert Tijskens, UANTWERPEN Eric Pascolo, CINECA Eva Casoni, BSC Frédéric Pinel, UL Guillaume Houzeaux, BSC Isabelle Dupays, IDRIS Ivan Spisso, CINECA Jacob Finkenrath, CaSToRC Jing Gong, KTH Kübra Nari, UHeM

		Lilit Axner, KTH Karim Hasnaoui, IDRIS Kim Serradell, BSC Mariusz Uchronski, WCSS Massimiliano Guarrasi, CINECA Matej Spetko, IT4I Matjaz Subelj, UL Maxime Mogé, SURFSARA Milosz Bialczak, WCSS Miroslaw Kupczyk, PSNC Pedro Alberto, UC-LCA Pedro Ojeda May, SNIC-UU Ricard Borrell, BSC Stelios Erotokritou, CASTORC Stella Paronuzzi, BSC Tiziana Castriganò, CINECA Wahid Rofagha, PRACE aisbl Xu Guo, EPCC
	Reviewed by:	Simone Bnà, CINECA Thomas Eickermann, JUELICH
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	22/September/2020	Draft	Initial draft and set up document structure
0.2	27/September/2020	Draft	Included T7.1 and T7.5
0.3	30/September/2020	Draft	Included T7.2 and T7.3
0.4	02/October/2020	Draft	Draft finalisation
1.0	16/October/2020	Final	Integration of reviewer remarks

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Preparatory Access, SHAPE, HLST

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

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

[2] https://prace-ri.eu/hpc-access/calls-for-proposals/

[3] https://prace-ri.eu/hpc-access/preparatory-access/

[4] PRACE-5IP Deliverable 7.2, https://prace-ri.eu/wp-content/uploads/5IP-D7.2.pdf

[5] PRACE-5IP Deliverable 7.1, "Periodic Report on Applications Enabling Services", http://www.prace-ri.eu/IMG/pdf/D7.1_5ip.pdf

[6] PRACE-3IP Deliverable 5.2, "Integrated HPC Access Programme for SMEs", http://www.prace-ri.eu/IMG/pdf/d5.2.pdf

[7] PRACE-3IP Deliverable 5.3.1, "PRACE Integrated Access Programme Launch", http://www.prace-ri.eu/IMG/pdf/d5.3.1.pdf

[8] PRACE-3IP Deliverable 5.3.2, "Results of the Integrated Access Programme Pilots", http://www.prace-ri.eu/IMG/pdf/d5.3.2_3ip.pdf

[9] PRACE-3IP Deliverable 5.3.3, "Report on the SHAPE Implementation", http://www.prace-ri.eu/IMG/pdf/D5.3.3_3ip.pdf

[10] EaP Connect, https://www.eapconnect.eu/

[11] PRACE white papers, https://prace-ri.eu/training-support/technical-documentation/white-papers/

[12] PRACE white papers – SHAPE Projects, https://prace-ri.eu/training-support/technical-documentation/white-papers/shape-white-papers/

[13] Delft3D Flexible Mesh Suite, https://www.deltares.nl/en/software/delft3d-flexible-mesh-suite/

[14] PRACE Whitepaper 284, "Scalable Delft3D Flexible Mesh for efficient modeling of shallow water and transport processes", https://prace-ri.eu/wp-content/uploads/WP284.pdf

[15] J. Cornelis, S. Cools, W. Vanroose, The communication-hiding Conjugate Gradient method with deep pipelines., 2018, SIAM Journal on Scientific Computing (SISC)

[16] PRACE Whitepaper 293, "Performance Analysisof the BiqBin solver for the maximum stable set problem", https://prace-ri.eu/wp-content/uploads/WP293-Performance-Analysis-of-the-BiqBin-solver-for-the-maximum.pdf

[17] PRACE Whitepaper 288, "Optimisation of the Higher-Order Finite-Volume Unstructured Code Enhancement for Compressible Turbulent Flows", https://prace-ri.eu/wp-content/uploads/WP288.pdf

[18] PRACE Whitepaper 294, "PETSc4FOAM: A Library to plug-in PETSc into the OpenFOAM Framework", https://prace-ri.eu/wp-content/uploads/WP294-PETSc4FOAM-A-Library-to-plug-in-PETSc-into-the-OpenFOAM-Framework.pdf

[19] Marco Kupiainen, Jing Gong, Lilit Axner, Erwin Laure and Jan Nordström, "GPU-acceleration of a High Order Finite Difference Code using Curvilinear Coordinates", The proceeding of CNIOT2020, pp. 41-47, April 24–26, 2020, Sanya, China

[20] PRACE Whitepaper 295, "Keeping computational performance analysis simple: an evaluation of the NEMO BENCH test", https://prace-ri.eu/wp-content/uploads/WP295-Keeping-computational-performance-analysis-simple-An-evaluation-of-the-NEMO-BENCH-test.pdf

[21] PRACE-5IP Deliverable 2.1, "Report on PRACE 2, TNA, DECI and KPIs Year 1", confidential, available only for members of the consortium (including the Commission Services)

[22] Hardware details Archer 2, https://www.archer2.ac.uk/about/hardware.html
[23] Hardware details JOLIOT Curie, https://prace-ri.eu/hpc-access/hpc-systems/#JOLIOTCurie

List of Acronyms and Abbreviations

aisbl	Association International Sans But Lucratif (legal form of the PRACE-RI)
BCO	Benchmark Code Owner
CoE	Center of Excellence
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture (NVIDIA)
DARPA	Defense Advanced Research Projects Agency
DEISA	Distributed European Infrastructure for Supercomputing Applications EU project by leading national HPC centres
DoA	Description of Action (formerly known as DoW)
EC	European Commission
EESI	European Exascale Software Initiative
EoI	Expression of Interest
ESFRI	European Strategy Forum on Research Infrastructures
GB	Giga (= 2^{30} ~ 10^9) Bytes (= 8 bits), also GByte
Gb/s	Giga (= 10^9) bits per second, also Gbit/s
GB/s	Giga (= 10^9) Bytes (= 8 bits) per second, also GByte/s
GÉANT	Collaboration between National Research and Education Networks to build a multi-gigabit pan-European network. The current EC-funded project as of 2015 is GN4.
GFlop/s	Giga (= 10^9) Floating point operations (usually in 64-bit, i.e. DP) per second, also GF/s
GHz	Giga (= 10^9) Hertz, frequency =10^9 periods or clock cycles per second
GPU	Graphic Processing Unit
HET	High Performance Computing in Europe Taskforce. Taskforce by representatives from European HPC community to shape the European HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PRACE project.
HLST	High Level Support Team
HMM	Hidden Markov Model
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
HPL	High Performance LINPACK
ISC	International Supercomputing Conference; European equivalent to the US based SCxx conference. Held annually in Germany.
KB	Kilo (= 2^{10} ~10^3) Bytes (= 8 bits), also KByte
LINPACK	Software library for Linear Algebra
MB	Management Board (highest decision making body of the project)
MB	Mega (= 2^{20} ~ 10^6) Bytes (= 8 bits), also MByte
MB/s	Mega (= 10^6) Bytes (= 8 bits) per second, also MByte/s
MFlop/s	Mega (= 10^6) Floating point operations (usually in 64-bit, i.e. DP) per second, also MF/s
MOOC	Massively open online Course

MoU	Memorandum of Understanding.
MPI	Message Passing Interface
NDA	Non-Disclosure Agreement. Typically signed between vendors and customers working together on products prior to their general availability or announcement.
PA	Preparatory Access (to PRACE resources)
PATC	PRACE Advanced Training Centres
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE 2	The current phase of the PRACE Research Infrastructure following the initial five year period.
PRIDE	Project Information and Dissemination Event
RI	Research Infrastructure
TB	Technical Board (group of Work Package leaders)
TB	Tera (= $2^{40} \sim 10^{12}$) Bytes (= 8 bits), also TByte
TCO	Total Cost of Ownership. Includes recurring costs (e.g. personnel, power, cooling, maintenance) in addition to the purchase cost.
TDP	Thermal Design Power
TFlop/s	Tera (= 10^{12}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also TF/s
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1
UNICORE	Uniform Interface to Computing Resources. Grid software for seamless access to distributed resources.

List of Project Partner Acronyms

BADW-LRZ	Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften, Germany (3rd Party to GCS)
BILKENT	Bilkent University, Turkey (3rd Party to UHEM)
BSC	Barcelona Supercomputing Center - Centro Nacional de Supercomputacion, Spain
CaSToRC	The Computation-based Science and Technology Research Center (CaSToRC), The Cyprus Institute, Cyprus
CCSAS	Computing Centre of the Slovak Academy of Sciences, Slovakia
CEA	Commissariat à l’Energie Atomique et aux Energies Alternatives, France (3rd Party to GENCI)
CENAERO	Centre de Recherche en Aéronautique ASBL, Belgium (3rd Party to UANTWERPEN)
CESGA	Fundacion Publica Gallega Centro Tecnológico de Supercomputación de Galicia, Spain, (3rd Party to BSC)
CINECA	CINECA Consorzio Interuniversitario, Italy
CINES	Centre Informatique National de l’Enseignement Supérieur, France (3 rd Party to GENCI)
CNRS	Centre National de la Recherche Scientifique, France (3 rd Party to GENCI)
CSC	CSC Scientific Computing Ltd., Finland
CSIC	Spanish Council for Scientific Research (3rd Party to BSC)

CYFRONET	Academic Computing Centre CYFRONET AGH, Poland (3rd Party to PNSC)
DTU	Technical University of Denmark (3rd Party of UCPH)
EPCC	EPCC at The University of Edinburgh, UK
EUDAT	EUDAT OY
ETH Zurich (CSCS)	Eidgenössische Technische Hochschule Zürich – CSCS, Switzerland
GCS	Gauss Centre for Supercomputing e.V., Germany
GÉANT	GÉANT Vereniging
GENCI	Grand Equipement National de Calcul Intensif, France
GRNET	National Infrastructures for Research and Technology, Greece
ICREA	Catalan Institution for Research and Advanced Studies (3rd Party to BSC)
INRIA	Institut National de Recherche en Informatique et Automatique, France (3rd Party to GENCI)
IST-ID	Instituto Superior Técnico for Research and Development, Portugal (3rd Party to UC-LCA)
IT4I	Vysoka Skola Banska - Technicka Univerzita Ostrava, Czech Republic
IUCC	Machba - Inter University Computation Centre, Israel
JUELICH	Forschungszentrum Jülich GmbH, Germany
KIFÜ (NIIFI)	Governmental Information Technology Development Agency, Hungary
KTH	Royal Institute of Technology, Sweden (3rd Party to SNIC-UU)
KULEUVEN	Katholieke Universiteit Leuven, Belgium (3rd Party to UANTWERPEN)
LiU	Linkoping University, Sweden (3rd Party to SNIC-UU)
MPCDF	Max Planck Gesellschaft zur Förderung der Wissenschaften e.V., Germany (3rd Party to GCS)
NCSA	NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS, Bulgaria
NTNU	The Norwegian University of Science and Technology, Norway (3rd Party to SIGMA2)
NUI-Galway	National University of Ireland Galway, Ireland
PRACE	Partnership for Advanced Computing in Europe aisbl, Belgium
PSNC	Poznan Supercomputing and Networking Center, Poland
SDU	University of Southern Denmark (3rd Party to UCPH)
SIGMA2	UNINETT Sigma2 AS, Norway
SNIC-UU	Uppsala Universitet, Sweden
STFC	Science and Technology Facilities Council, UK (3rd Party to UEDIN)
SURFsara	Dutch national high-performance computing and e-Science support center, part of the SURF cooperative, Netherlands
TASK	Politechnika Gdańska (3rd Party to PNSC)
TU Wien	Technische Universität Wien, Austria
UANTWERPEN	Universiteit Antwerpen, Belgium
UC-LCA	Universidade de Coimbra, Labotatório de Computação Avançada, Portugal
UCPH	Københavns Universitet, Denmark
UEDIN	The University of Edinburgh
UHEM	Istanbul Technical University, Ayazaga Campus, Turkey
UIBK	Universität Innsbruck, Austria (3rd Party to TU Wien)
UiO	University of Oslo, Norway (3rd Party to SIGMA2)
UL	UNIVERZA V LJUBLJANI, Slovenia
ULIEGE	Université de Liège; Belgium (3rd Party to UANTWERPEN)

U Luxembourg	University of Luxembourg
UM	Universidade do Minho, Portugal, (3rd Party to UC-LCA)
UmU	Umea University, Sweden (3rd Party to SNIC-UU)
UnivEvora	Universidade de Évora, Portugal (3rd Party to UC-LCA)
UnivPorto	Universidade do Porto, Portugal (3rd Party to UC-LCA)
UPC	Universitat Politècnica de Catalunya, Spain (3rd Party to BSC)
USTUTT-HLRS	Universitaet Stuttgart – HLRS, Germany (3rd Party to GCS)
WCSS	Politechnika Wroclawska, Poland (3rd Party to PNSC)

Executive Summary

Applications Enabling and Porting Services in Work Package 7 (WP7) of PRACE-6IP aims to provide HPC enabling support for the applications of European researchers and small and medium enterprises to ensure the applications can effectively exploit the various PRACE HPC systems. Most of the activities are ongoing and were already established and described in deliverables of former PRACE-IP projects. There are four activities described in this deliverable:

T7.1 Applications Enabling Services for Preparatory Access:

This activity provides code enabling and optimisation to European researchers as well as industrial projects to make their applications ready for Tier-0 systems. Projects can continuously apply for such services via the Preparatory Access Calls Type C (PA Type C) and Type D (PA Type D) with a cut-off every three months for evaluation of the proposals. PA Type C provides support and access to a PRACE Tier-0 system while PA Type D provides support and access to a PRACE Tier-1 system to finally reach Tier-0 scalability. In total six Preparatory Access cut-offs have been carried out in PRACE-6IP so far and 16 projects received support within the context of the project.

Beside the statistical overview about the cut-offs and all supported PRACE PA Type C and Type D projects, the report focuses on the optimisation work and results gained by the completed projects in PRACE-6IP. In total eight PA projects have finished their work since the last deliverable D7.2 of PRACE-5IP [4] and are reported within this deliverable.

T7.2 Applications Enabling Services for Industry (SHAPE):

This activity has continued the support for SHAPE (the SME HPC Adoption Programme in Europe) into the PRACE-6IP project. SHAPE aims to raise awareness of HPC within European SMEs by providing them with the expertise necessary to take advantage of the innovation possibilities created by HPC with the aim of increasing the SME's competitiveness. The programme runs a series of regular calls with successful applicants getting support effort from a PRACE HPC expert and access to machine time at a PRACE centre.

One of the key focuses of the last few calls has been to increase the number of countries benefiting from SHAPE support. To this end the SHAPE+ initiative was developed where an allocation of effort is provided to support SMEs from countries new to SHAPE in cases where the local partners don't have existing SHAPE effort. There are good signs that this is working as the number of countries where SMEs have applied to SHAPE, and where SHAPE effort has been awarded, has been increasing more rapidly for the last few calls.

More generally there has been a focus on better dissemination of SHAPE to increase the overall number of proposals. This has been greatly assisted by the employment in September 2019 of a PRACE Industry Liaison Officer who brings a wealth of expertise in communicating with industry. He has first developed a new market proposal for SHAPE which was then tested by him "on the field" by engaging in over 220 face-to-face meetings with managers of European SMEs at 17 industry-specific trade shows. Furthermore, he presented SHAPE to the membership base of several business/innovation clusters in Europe. The feedback he reported back to the SHAPE Project Manager was overwhelmingly positive and indicated a very clear demand for the programme in the R&D-driven business community. This also shows in the increase in the number of proposals SHAPE receives. For the most recent call (SHAPE 11) we received 10 proposals - the largest number since call 2 (back in 2015) and the largest number we have received since the frequency of calls was increased to 2 per year where until this most recent call the average number of proposals had been just under 6.

We report here on the current status of each SHAPE project and highlights of the work carried out. In 2019 [4] we developed surveys for SMEs and PRACE partners where projects had completed around a year ago. These were very successful in gaining both quantitative and qualitative information about the effectiveness of the SHAPE programme and so the exercise was repeated with a slightly updated survey. The results were again positive.

T7.3 DECI Management and Applications Porting:

The DECI (Distributed European Computing Initiative) programme began in 2005 as part of DEISA (Distributed European Infrastructure for Supercomputing Applications), later moving into the PRACE Implementation Phase (IP) projects and then run as a PRACE "Optional Programme". Presently it runs under the Applications Enabling and Support Workpackage (Work Package 7) within PRACE-6IP. The most recent DECI call (DECI-16) closed on 31 January 2020.

The DECI programme provides access to Tier-1 resources across Europe via a resource exchange programme and for the last few years calls have been issued yearly. Calls remain popular with researchers and since the programme began, over 1000 proposals have been received with the last two calls both receiving around 70 proposals resulting in 40 and 48 projects respectively.

Here we focus on the projects from the two most recent DECI calls which have provided around 150M machine core hours each to European researchers from 88 different DECI projects, from 22 different countries across a wide range or architectures provided from 13 different countries.

T7.5 Enhancing the High-Level Support Teams:

T7.5 is a new task implemented in PRACE-6IP. It will provide additional effort to enhance the work of the PRACE High-Level Support Teams (HLSTs). Under PRACE 2, there are coordinated HLSTs at each Tier-0 centre that provide users with support for code enabling and scaling out of scientific applications and methods, as well as for Research & Development on code refactoring on the Tier-0 systems. The HLSTs will enhance the scientific output of the Tier-0 systems through the provision of level 3 (midterm activities) support activities. Task 7.5 supplements those activities and extends this work with specific expertise from the other PRACE centres. This ensures sharing of expertise across PRACE to maximise the benefits to users of the PRACE systems. Task 7.5 will work with the HLSTs to extend and enhance their activities supporting Tier-0 users, and also Tier-1 intensive users targeting Tier-0, in order to maximise the scientific output of the Tier-0 systems.

During the first half of the project, we focused on defining the process of project selection, based on criteria that meet what Tier-0 systems can provide to the users (scalability, new architectures, memory/storage …). The support activity brought to the users was also defined, and two projects have already started within the task.

During the second half of the project we will continue the activity on Petascale and Pre-exascale systems, which is the main target for the activity, and we expect to have an increased effort of the partners on those new systems by the end of the project.

1 Introduction

PRACE offers a wide range of different Tier-0 and Tier-1 HPC architectures to the scientific community as well as to industrial innovative projects. The efficient usage of such systems places high demands on the used software packages and in many cases advanced optimisation work has to be applied to the code to make efficient use of the provided supercomputers. The complexity of supercomputers requires a high level of experience and advanced knowledge of different concepts regarding programming techniques, parallelization strategies, etc. Such demands often cannot be met by the applicants themselves and thus special assistance by supercomputing experts is essential.

PRACE offers such a service through the Preparatory Access Calls. PA Type C and PA Type D are managed by Task 7.1. This includes the evaluation of the PA proposals as well as the assignment of PRACE experts to these proposals. Furthermore, the support itself is provided and monitored within this task. Section 2 gives a more detailed description of PA and some facts on the usage of PA Type C and PA Type D in PRACE-6IP are listed in 2.1. The review process, the assignment of PRACE experts to the projects and the monitoring of the support work are detailed in Section 2.2, Section 2.3 and Section 2.4 respectively. The contents of Sections 2.2 - 2.4 can already be found in deliverable D7.2 of PRACE-5IP [4]. They are repeated here for completeness and the benefit of the reader. Section 2.5 gives some details concerning additional preparatory access supported which was provided towards the EaP connect project [10]. Section 2.6 gives an overview about the Preparatory Access projects covered in PRACE-6IP so far. The dissemination of the PA call is described in Section 2.7. Finally, the work done within the projects along with the outcome of the optimisation work is presented in Sections 2.8 - 2.13.

The second part of this deliverable is the report on SHAPE (SME HPC Adoption Programme in Europe). SHAPE is a pan-European programme set up to support the adoption of High Performance Computing (HPC) by European small to medium-size enterprises (SMEs). It was developed by PRACE under its PRACE-3IP European Union funded project and has continued through to the present PRACE-6IP project.

The SHAPE programme, originally presented in the PRACE-3IP Deliverable 5.2 [6], aims to equip European SMEs with the awareness and expertise necessary to take advantage of the innovation possibilities opened by HPC, with the aim of increasing their competitiveness. The programme aims to help European SMEs to demonstrate a tangible Return on Investment (ROI) from assessing and adopting solutions supported by HPC, thus facilitating innovation and/or increased operational efficiency in their businesses.

The initial SHAPE pilot [7] [8] was launched in 2013 and helped 10 SMEs adopt HPC, with a follow-up exercise to gauge the business impact for the SMEs showing in almost all the cases that the pilot had been of real value to the SMEs, with tangible measures of the ROI for the SHAPE work [9]. Following this pilot, the PRACE Council decided to operate the SHAPE programme as a permanent service. A second call was launched in November 2014 and closed in January 2015, and then from the third call (launched in November 2015, closing in January 2016) calls were launched at 6 monthly intervals. SHAPE has now run 11 calls with the 12th call due to open in October 2020 and close in December 2020.

In section 3 we report on discussions that took place during meetings to replace face-to-face sessions cancelled due to COVID-19, such as the Working with SMEs (WP7-WP2-WP3-WP6-aisbl) event in September where around 75 attendees from SHAPE, other PRACE tasks, and the EuroHPC Competence Centres met and discussed ways in which we can all work together and share experiences in the future.

In section 3.1 we give an overview of the SHAPE programme. In section 3.2 we then present the SHAPE+ programme where we report on the successes of encouraging proposals from countries new to SHAPE. The present status of SHAPE is described in section 3.3. In section 3.4 the survey results for projects finishing around a year ago (sixth and seventh calls) are given with a report for each project followed by an overview of the results for all projects. This is followed in section 3.5 with a description of the more recent projects (calls seven to eleven), and a brief description of the upcoming twelfth call is given in section 3.6. The SHAPE description is then concluded in section 3.7 with discussions about the future of SHAPE and potential improvements.

The third part of this deliverable is the report on the DECI (Distributed European Computing Initiative) resource exchange programme. Section 4.1 gives an overview of the programme and the present status. Most of the effort for DECI goes towards running the programme itself, with a small amount of effort provided for enabling. This is described in section 4.2. Future plans are then described in section 4.3.

The section is followed by section 5, which covers T7.5. T7.5 is a new task implemented in PRACE-6IP. It provides additional effort to enhance the work of the PRACE High-Level Support Teams (HLSTs). Under PRACE 2, there are coordinated HLSTs at each Tier-0 centre that provide users with support for code enabling and scaling out of scientific applications and methods, as well as for Research & Development on code refactoring on the Tier-0 systems. The HLSTs will enhance the scientific output of the Tier-0 systems through the provision of level 3 (midterm activities) support activities. Task 7.5 supplements those activities and extends this work with specific expertise from the other PRACE centres. This will ensure sharing of expertise across PRACE to maximise the benefits to users of the PRACE systems. Task 7.5 will work with the HLSTs to extend and enhance their activities supporting Tier-0 users, and also Tier-1 intensive users targeting Tier-0, in order to maximise the scientific output of the Tier-0 systems.

This deliverable will describe the process that has been set up within this task for targeting the right projects that would benefit from HLSTs and the ongoing activities on the projects accepted as well as the planned activities for the next period.

The deliverable closes with a summary in section 6 highlighting the outcomes of the different tasks.

2 T7.1 Applications Enabling Services for Preparatory Access

Access to PRACE Tier-0 systems is managed through PRACE regular calls, which are issued twice a year [2]. To apply for Tier-0 resources the application must meet technical criteria concerning scaling capability, memory requirements, and runtime set up. There are many important scientific and commercial applications that do not meet these criteria today. To support researchers PRACE offers the opportunity to test and optimise their applications on the envisaged Tier-0 system prior to applying for a regular production project. This is the purpose of the Preparatory Access (PA) Call. The PA Call allows for submission of proposals at any time. Depending on the PA scheme, the review of these proposals takes place directly after the submission of the proposal (Type A and B) or at a fixed date every three months (Type C and D). This procedure is also referred to as cut-off [3]. It is possible to choose between four different types of access:

- Type A is meant for code scalability tests, the outcome of which is to be included in the proposal in a future PRACE regular call. Users receive a limited number of core hours; the allocation period is two months.
- Type B is intended for code development and optimisation by the user. Users also get a small number of core hours; the allocation period is six months.
- Type C is also designed for code development and optimisation with the core hours and the allocation period being the same as for Type B. The important difference is that Type C projects receive special assistance by PRACE experts from the PRACE-IP project to support the optimisation. As well as access to the Tier-0 systems, the applicants also apply for one to six person months (PMs) of supporting work to be performed by PRACE experts.
- Type D allows PRACE users to start a code adaptation and optimisation process on a PRACE Tier-1 system. PRACE experts help in the system selection process. In addition to Tier-1 computing time, the PRACE user will also receive Tier-0 computing time towards the end of the project in the form of a PA Type A project to test the scalability improvements. The work is supported by PRACE experts similar to Type C. The maximum duration of Type D projects is twelve months.

Task 7.1 of the PRACE project covers the expert assistance in context of PA Type C and D. PA Type A and B projects are directly handled by the different hosting sites without involvement of the PRACE-IP project.

2.1 Cut-off statistics

Since the last PRACE-5IP deliverable D7.2 [4] six cut-offs for PA took place in between April 2019 and October 2020 resulting in nine new projects so far. In total, sixteen projects were supported by PRACE-6IP T7.1 within this time frame. Seven of them started already during PRACE-5IP or were established with the frame of a PRACE-5IP cut-off. All of those overlapping projects were taken over by PRACE-6IP. Nine projects were finalized during this time frame. Two of those Type D projects were established in collaboration between PRACE and the EaP Connect initiative [10]. These will be reported in section 2.5. All remaining projects are ongoing and will be reported in an upcoming PRACE-6IP deliverable.

Within the reporting period, only one proposal had to be rejected, because it did not fit into the scope of the preparatory access call. No proposal was rejected due to a lack of available PRACE experts during the review phase.

The September 2020 cut-off is currently ongoing. Here two additional new PA Type C proposals are currently reviewed. One proposal was taken over from SHAPE and could already be accepted in the meantime.

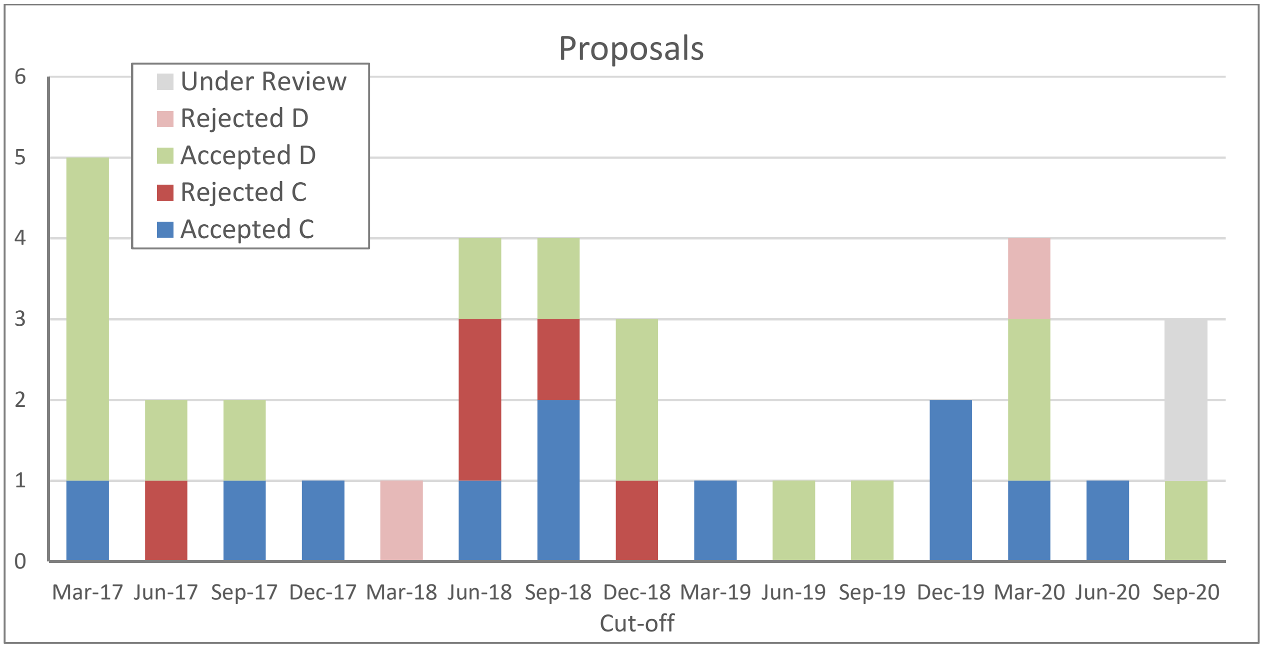

Figure 1: Number of proposals for PA type C and type D per cut-off.

Figure 1 presents the number of proposals that have been accepted or rejected, respectively for each cut-off. So far nine out of twelve proposals were accepted during the reporting period beginning in April 2019 until the cut-off in September 2020.

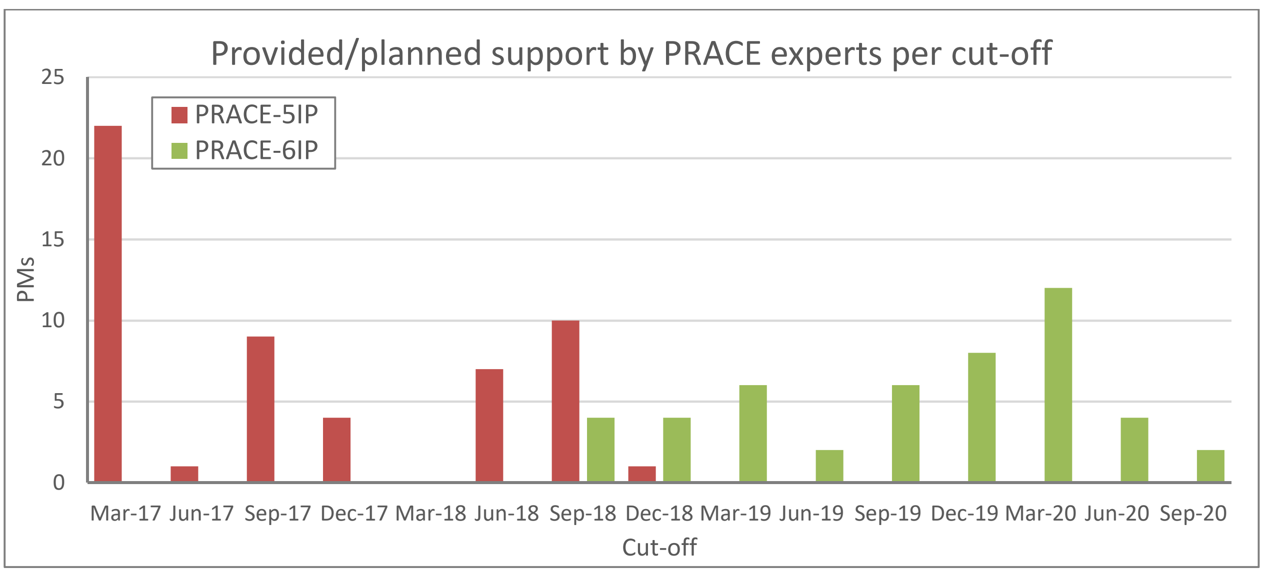

Figure 2: Amount of PMs assigned to PA projects per cut-off.

Figure 2 gives an overview of the number of PMs from the PRACE project assigned to the projects per cut-off. In total 48 PMs were made available to these projects within the context of PRACE-6IP so far. In the September 2018 and December 2018 cut-off of PRACE-5IP, long running projects, which had an overlap with the PRACE-6IP project phase were planned to already receive resources in context of PRACE-6IP. The projects, which were established in the March 2019 cut-off were fully taken over by PRACE-6IP.

Finally, Figure 3 provides an overview of the scientific fields, which are covered by the supported projects in PRACE-6IP.

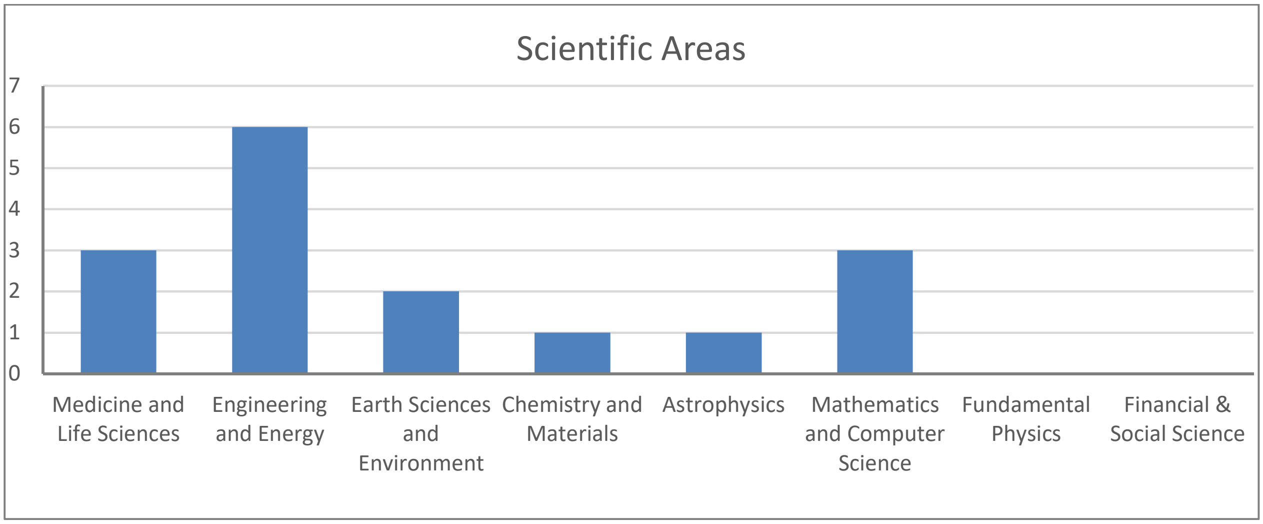

Figure 3: Number of projects per scientific field.

2.2 Review Process

The organisation of the review procedure, the assignment of PRACE collaborators and the supervision of the PA Type C and D projects are managed by task 7.1. In this section, the review process for the preparatory access proposals of Type C and Type D is explained.

All preparatory access proposals undergo a technical review performed by technical staff of the hosting sites to ensure that the underlying codes are in principle able to run on the requested system. For PA C projects, the technical review starts directly after the cut-off. For PA D projects, the technical review is done after the Tier-1 system is finally selected.

In parallel, all projects are additionally reviewed by WP7 staff in order to assess their optimisation requests. Each proposal is assigned to two WP7 reviewers. The review is performed by PRACE partners who all have a strong background in supercomputing. The task leader has the responsibility to contact them to launch the review process. As the procedure of reviewing proposals and establishing the collaboration of submitted projects and PRACE experts takes place only four times a year, it is necessary to keep the review process swift and efficient. A close collaboration between PRACE aisbl, T7.1.A and the hosting sites is important in this context. The process for both the technical and the WP7 review is limited to six weeks. In close collaboration with PRACE aisbl and the hosting sites, the whole procedure from PA cut-off to project start on

PRACE supercomputing systems is completed in less than six weeks as well. Due to different availabilities the PI as well as the PRACE expert can change the final start date if necessary.

Based on the proposals the Type C and Type D reviewers need to focus on the following aspects:

- Does the project require support for achieving production runs on the chosen architecture?
- Are the performance problems and their underlying reasons well understood by the applicant?
- Is the amount of support requested reasonable for the proposed goals?
- Will the code optimisation be useful to a broader community, and is it possible to integrate the achieved results during the project in the main release of the code(s)?
- Will there be restrictions in disseminating the results achieved during the project?

For Type D, the reviewer should also make suggestions for the Tier-1 selection process. The 7.1 task leader finally selects the Type D Tier-1 computing site based on the suggestions of the PI, the potential PRACE collaborator, the reviewers and the computing site availability.

In addition to the WP7 reviews, the task leader evaluates whether the level and type of support requested is still available within PRACE. Finally, the recommendation from WP7 to accept or reject the proposal is made.

The outcome is communicated to the applicant through PRACE aisbl. Here also the PRACE Board of Directors is informed. Approved proposals receive the contact data of the assigned PRACE collaborators, rejected projects are provided with further advice on how to address the shortcomings.

2.3 Assigning of PRACE collaborators

To ensure the success of the projects it is essential to assign suitable experts from the PRACE project. Based on the described optimization issues and support requests from the proposal experts are thus chosen who are most familiar with the subject matter.

This is done in two steps: first, summaries of the proposals describing the main optimization issues are distributed via corresponding mailing lists. Here, personal data is explicitly removed from the reports to maintain the anonymity of the applicants. Interested experts can get in touch with the task leader offering to work on one or more projects.

There is one exception to the procedure when a proposal has a close connection to a PRACE site which has already worked on the code: In this case this site is asked first if they are able to extend the collaboration in the context of the new PA Type C or PA Type D project.

Should the response not be sufficient to cover the support requirements of the projects, the task leader contacts the experts directly and asks them to contribute.

The assignment of PRACE experts takes place concurrently to the review process so that the entire review can be completed within six weeks. This has proven itself to be a suitable approach, as the resulting overhead is negligible.

As soon as the review process is finished, the support experts are introduced to the PIs and can start the work on the projects. The role of the PRACE collaborator includes the following tasks:

- Formulating a detailed work plan together with the applicant;
- Participating in the optimization work;
- Reporting the status in the task 7.1 phone conference every second month;
- Participating in the writing of the final report together with the PI (the PI has the main responsibility for this report), due at project end and requested by the PRACE office;
- Writing a white paper containing the results, which is published on the PRACE web site [1].

2.4 Monitoring of projects

Task 7.1 includes the supervision of the Type C and Type D projects. This is challenging as the projects' durations (six months for PA Type C and twelve months for PA Type D) and the intervals of the cut-offs (three months) are not synchronised. Due to this, projects do not necessarily start and end at the same time but overlap, i.e. at each point in time different projects might be in different phases. To solve this problem, a phone conference and regular status updates take place in task 7.1 every two months to discuss the status of running projects, to advise on how to proceed with new projects and to manage the finalization and reporting of finished projects. In addition the task leader monitors all relevant project deadlines to inform the PRACE experts concerning reporting periods.

In addition, the T7.1 task leader gives a status overview in a monthly WP7 conference call to address all PRACE collaborators who are involved in these projects.

The T7.1 task leader is also available to address urgent problems and additional phone conferences are held in such cases.

Twice a year, a WP7 face-to-face meeting is scheduled. This meeting gives all involved collaborators the opportunity to discuss the status of the projects and to exchange their experience. Also project results are presented within these meetings.

2.5 PA support in context of EaP connect

The Eastern Partnership Connect (EaP Connect) [10] is a project funded by the European Union, launched in 2015 to

- establish and operate a high-capacity broadband internet network for research and education (R&E) across six EaP partner countries;
- integrate the national research and education networks (NRENs) into the pan-European network, decreasing the digital divide;
- facilitate participation of local scientists, students and academics in global R&E collaborations.

To support the EaP projects and scientists, PRACE offered the option to access PRACE Tier-0 and Tier-1 systems in the frame of the PA program. To manage the system access, PA Type B proposals were selected for the Tier-0 and PA Type D proposals for the Tier-1 access. The support within the projects was mostly limited to technical guidance by the involved hosting sites, therefore these projects are not listed alongside the regular projects in the following sections. Task 7.1 coordinated the system access and helped during the PA proposal creation and provided initial guidance. Table 1 lists all supported projects in this context.

Supported projects in context of EaP connect	
Title	**Deep Learning applications for big medical image archives: detection of diseases and generation of artificial images**
Type	D
Project leader	Dzmitry Paulenka
PRACE expert	Damian Podareanu
PRACE facility	CARTESIUS
PA number	2010PA4773
Project's start	01-March-2019
Project's end	29-February-2020
Title	**Development of a parallel algorithm ILST for the iterative reconstruction of 3D image based on the MPI system**
Type	D
Project leader	Sergei Zolotarev
PRACE expert	Miroslaw Kupczyk
PRACE facility	EAGLE
PA number	2010PA5033
Project's start	01-August-2019
Project's end	31-July-2020
Title	**Forecasting convective high-impact weather events in Armenia using the WRF model and WRF Data Assimilation system**
Type	B
Project leader	Artur Gevorgyan
PRACE expert	Sebastian Lührs
PRACE facility	JUWELS
PA number	2010PA5239
Project's start	01-January-2020
Project's end	31-December-2020
Title	**The Parallel Molecular Dynamics Simulation of Mixed Lipid Bilayer modified with αVβ3 integrin**

Supported projects in context of EaP connect	
Type	B
Project leader	Ashot Meltonyan
PRACE expert	Sebastian Lührs
PRACE facility	JUWELS
PA number	2010PA5241
Project's start	01-January-2020
Project's end	31-December-2020
Title	**Computational material discovery for gas sensing application**
Type	B
Project leader	Hayk Zakaryan
PRACE expert	Sebastian Lührs
PRACE facility	JUWELS
PA number	2010PA5243
Project's start	01-January-2020
Project's end	31-December-2020

Table 1: List of supported projects in context of EaP connect

2.6 PRACE Preparatory Access projects covered by this report

The support for Preparatory Access projects has been and is part of all PRACE projects (PRACE-1IP, -2IP, -3IP, -4IP, -5IP, -6IP). Therefore, projects which were not finalized or at least not reported within PRACE-5IP were taken over by PRACE-6IP and will be reported in this deliverable.

The timeline of all projects supported by PRACE-6IP is shown in the chart of Figure 4. The chart shows the time span of each project. Projects supported by PRACE-5IP are shown in red. The projects which received support by PRACE-6IP are shown in green. The slightly different starting dates of the projects per cut-off are the result from the decisions made by the hosting members, the PI and the PRACE expert, which determine the exact start of each individual project.

Beside these projects, one project from PRACE-5IP, which was not reported before will be covered by this deliverable as well.

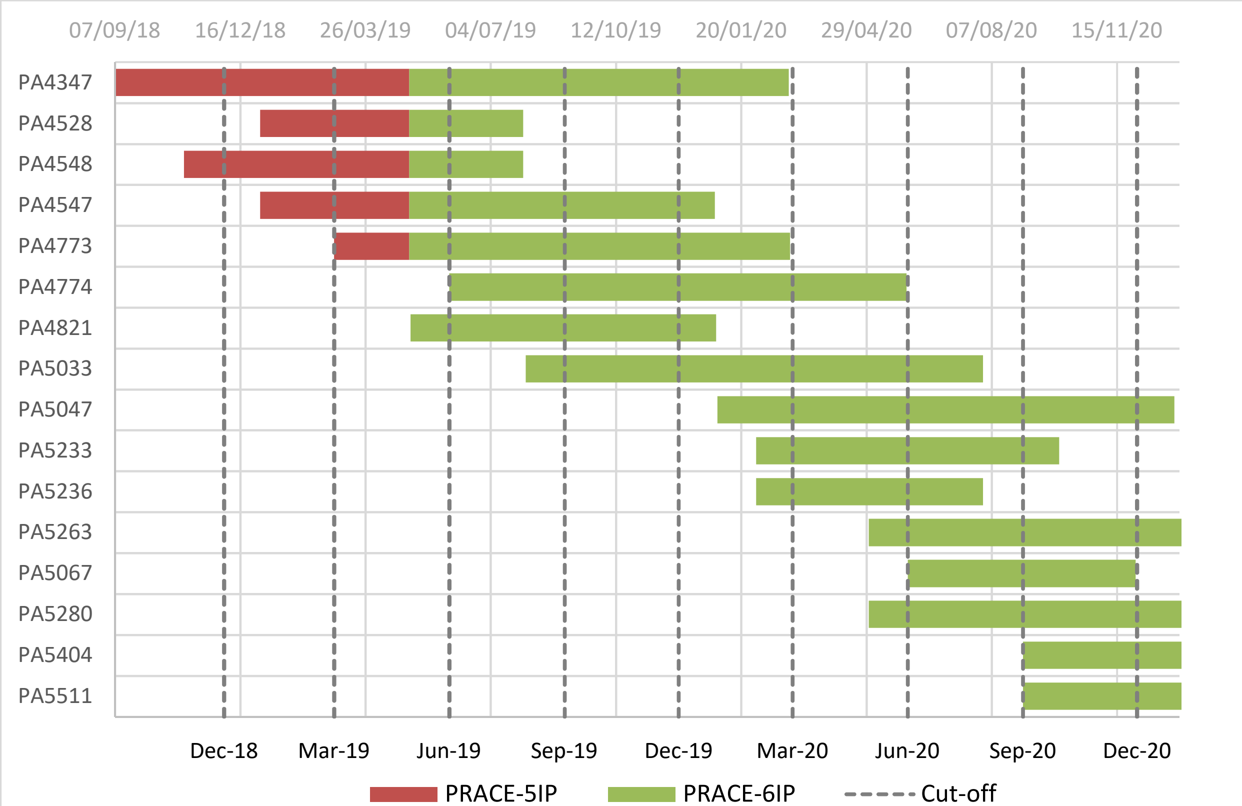

Figure 4: Timeline of the PA projects.

Except for one project from the March 2017 cut-off, all projects before the June 2018 cut-off were already reported in previous deliverables. This report will cover eight projects, which were finalized in between March 2019 and September 2020. Table 2 lists the corresponding projects.

Cut-off March 2017	
Title	**scalable Delft3D FM for efficient modelling of shallow water and transport processes**
Type	D
Project leader	Menno Genseberger
PRACE expert	Maxime Mogé, Martijn Russcher, Andrew Emerson
PRACE facility	CARTESIUS, MARCONI
PA number	2010PA3775
Project's start	01-July-2017
Project's end	30-March-2019

Cut-off June 2018	
Title	**BiqBin solver**

Cut-off June 2018	
Type	D
Project leader	Janez Povh
PRACE expert	Anastasia Shamakina
PRACE facility	SALOMON
PA number	2010PA4347
Project's start	01-September-2018
Project's end	28-February-2020

Cut-off September2018	
Title	**HOVE2 Higher-Order finite-Volume unstructured code Enhancement for compressible turbulent flows**
Type	C
Project leader	Panagiotis Tsoutsanis
PRACE expert	Anastasia Shamakina
PRACE facility	HAZEL HEN
PA number	2010PA4528
Project's start	01-January-2019
Project's end	31-July-2019
Title	**HPC Performance improvements for OpenFOAM linear solvers**
Type	C
Project leader	Mark Olesen
PRACE expert	Ivan Spisso
PRACE facility	MARCONI-KNL
PA number	2010PA4548
Project's start	01-November-2018
Project's end	31-July-2019
Title	**HPC-MetaShot: a multi-node implementation for metagenomics analysis**
Type	D
Project leader	Giovanni Chillemi
PRACE expert	Tiziana Castriganò
PRACE facility	GALILEO

Cut-off September2018	
PA number	2010PA4547
Project's start	01-January-2019
Project's end	31-December-2019

Cut-off December 2018	
Title	**Accelerating a high order accurate CFD solver using GPUs**
Type	D
Project leader	Marco Kupiainen
PRACE expert	Jing Gong
PRACE facility	CARTESIUS
PA number	2010PA4774
Project's start	01-June-2019
Project's end	31-May-2020

Cut-off March 2019	
Title	**NEMO performance**
Type	C
Project leader	Eric Maisonnave
PRACE expert	Kim Serradell
PRACE facility	MARENOSTRUM
PA number	2010PA4821
Project's start	01-May-2019
Project's end	31-December-2019

Cut-off December 2019	
Title	**Intelligent HTC for Committor Analysis**
Type	C
Project leader	Alan O'Cais
PRACE expert	Milosz Bialczak, Mariusz Uchronski, Adam Wlodarczyk
PRACE facility	JOLIOT CURIE KNL, JOLIOT CURIE SKL, JUWELS
PA number	2010PA5236
Project's start	01-February-2020

Cut-off December 2019	
Project's end	31-July-2020

Table 2: Projects, which are reported in this deliverable.

All remaining projects have to be reported in the next deliverable, as those projects are still ongoing or were not completely finalized so far. Table 3 lists the corresponding projects.

Cut-off September 2019	
Title	Next steps for scalable Delft3D FM for efficient modelling of shallow water and transport processes
Type	D
Project leader	Menno Genseberger
PRACE expert	Andrew Emerson, Maxime Mogé
PRACE facility	CARTESIUS, MARCONI
PA number	2010PA5047
Project's start	01-January-2020
Project's end	31-December-2020

Cut-off December 2019	
Title	**Load Balancing of Molecular Properties Calculations In VeloxChem Program**
Type	C
Project leader	Zilvinas Rinkevicius
PRACE expert	Isabelle Dupays
PRACE facility	SUPERMUC-NG
PA number	2010PA5233
Project's start	01-February-2020
Project's end	30-September-2020

Cut-off March 2020	
Title	**Enhancing Parallelism in MAGMA2**
Type	D
Project leader	Stephan Rosswog
PRACE expert	Isabelle Dupays
PRACE facility	BESCOW

Cut-off March 2020	
PA number	2010PA5263
Project's start	01-May-2020
Project's end	30-April-2021
Title	**XDEM4HPC: eXtended Discrete Element Method for High-Performance Computing**
Type	C
Project leader	Alban Rousset
PRACE expert	Isabelle Dupays
PRACE facility	MARENOSTRUM
PA number	2010PA5067
Project's start	01-June-2020
Project's end	30-November-2020
Title	**Scalability of Automatic Design Optimzation Algorithms**
Type	D
Project leader	Zilvinas Rinkevicius
PRACE expert	Joris Coddé
PRACE facility	BRENIAC
PA number	2010PA5280
Project's start	01-May-2020
Project's end	30-April-2021

Cut-off June 2020	
Title	**HOVE3 Higher-Order finite-Volume unstructured code Enhancement for compressible turbulent flows**
Type	C
Project leader	Panagiotis Tsoutsanis
PRACE expert	Anastasiia Shamakina
PRACE facility	HAWK
PA number	2010PA5404
Project's start	21-September-2020
Project's end	20-March-2021

Cut-off September 2020	
Title	**Use of HPC for optimization of drinking water distribution networks**
Type	D
Project leader	Mario Castro-Gama
PRACE expert	Jean-Christophe Pénalva
PRACE facility	OCCIGEN
PA number	2010PA5511
Project's start	01-September-2020
Project's end	28-February-2021

Table 3: Currently running and not fully finalized PA projects.

2.7 Dissemination

New PA Cut-offs are normally announced on the PRACE website [1] and through the official PRACE Twitter channel.

Sites are also asked to distribute an email to their users to advertise preparatory access and especially the possibility of the dedicated support.

Each successfully completed project should be made known to the public and therefore the PRACE collaborators are asked to write a white paper about the optimization work carried out. These white papers are published on the PRACE web page [11] and are also referenced by this deliverable.

2.8 Cut-off March 2017

This section and the following sub-sections describe the optimisations performed on the Preparatory Access Projects. The projects are listed in accordance with the cut-off dates in which they applied. General information regarding the optimisation work done as well as the achieved results are presented here.

2.8.1 Scalable Delft3D FM for efficient modelling of shallow water and transport processes, 2010PA3775

Overview:
Delft3D is used worldwide with a broad application range: from modelling of flooding, morphology to water quality, from coastal and estuarine areas, rivers to lakes, from consultancy work to applied research.

Currently, for Delft3D there is a transition from the shallow water solver Delft3D-FLOW for structured computational meshes to D-Flow Flexible Mesh ("D-Flow FM") for unstructured computational meshes. D-Flow FM is the hydrodynamic module of the Delft3D Flexible Mesh Suite [13].

For typical real-life applications, for instance for highly detailed modelling and operational forecasting, there is urgency to make D-Flow FM also more efficient and scalable for high performance computing. As the solver in D-Flow FM is quite different from the one in Delft3DFLOW, some major steps have to be taken. The main goal of the current project was to make significant progress towards Tier-0 systems for the shallow water and transport solvers in the Delft3D Flexible Mesh Suite.

Scalability results:
For improving and benchmarking the codes, we selected the following model applications, on which Deltares (https://www.deltares.nl/en) is currently working; they impose a computational challenge for the new software:

- global model for predicting tide and storm surges (GTSM test case), to zoom in from global to regional scale. The use of unstructured computational meshes to represent coastal areas in more detail than the open oceans is of importance as much of the tidal energy is dissipated on the shelf.
- Lake Marken model (Lake_Marken test case) and North Sea models (North_Sea_2D, North_Sea_3D, and North_Sea_3D_ST test cases) to have advanced modelling capabilities for integrated assessments of multiple impacts (flooding, ecology, water quality, climate change, infrastructure, sand extractions, fishery).
- Rhine Delta model (Waal_schematic and Waal_40m test cases), for quantifying effects of maintenance policies of the river bed that also affect navigability. The Rhine is a busy navigation route for transport over water with very high economic interest.

We used D-Flow FM Version 1.2.1.62244 compiled with flags '-xHost -axMIC-AVX512 -O3' on Marconi, and with flags '-xAVX2 -O3' on Cartesius. For all test cases there is first a serial initialization phase that we do not want to consider in the scalability analysis. We only consider the time spent in the time loop.

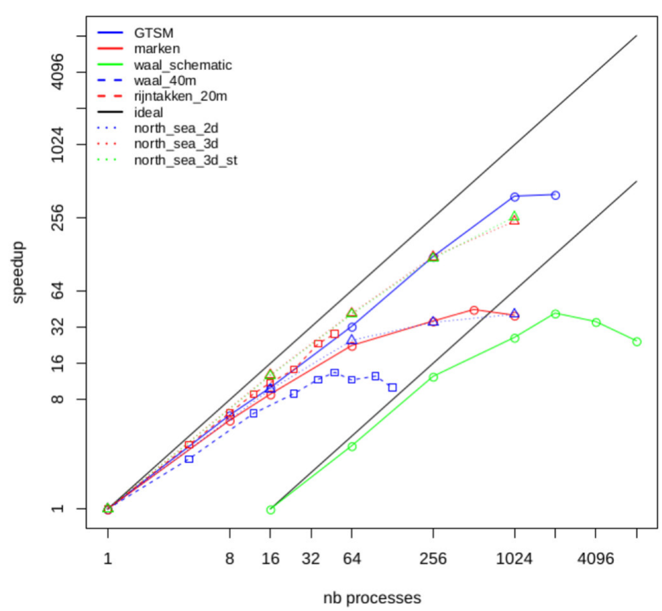

Figure 5: Speedup of all test cases on Cartesius

As Figure 5 shows, some test cases are too small to scale well for larger number of processes.

- For the Lake_Marken (345184 grid elements) and North_Sea_2D (348842 grid elements) test cases parallel efficiency is good (> 0.50) for only up to 64 processes.
- For the Waal_40m (108143grid elements) test case parallel efficiency is good (> 0.50) for only up to 48 processes.
- The biggest test cases (North_Sea_3D and North_Sea_3D_ST with 8721050 grid elements, Waal_schematic with 9000000 grid elements, and GTSM with 9584149 grid elements) show good scalability (efficiency > 0.50) up to 256 processes.

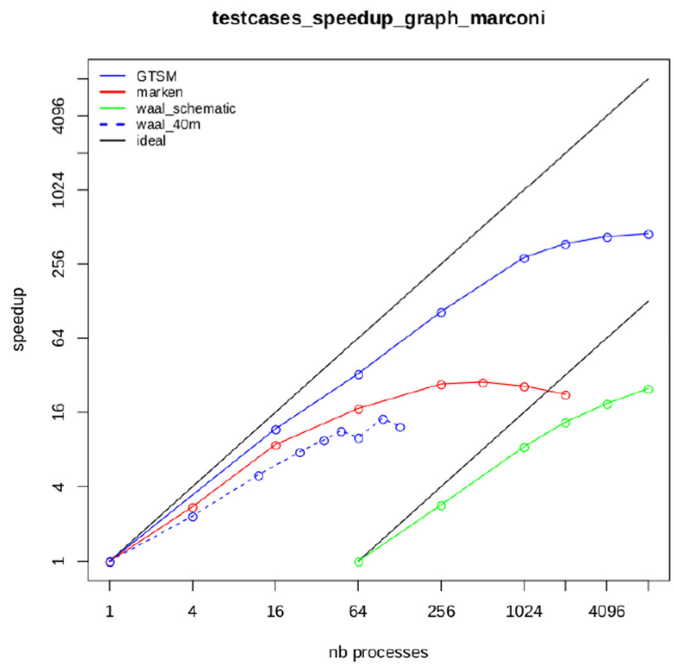

Figure 6: Speedup of all test cases on Marconi

Figure 6 shows that the scaling behaviour is similar on Marconi, but the computation time is much higher than on Cartesius when the number of processes is the same (~4x higher for most test cases, up to ~10x higher for Waal_schematic).

Accomplished work:

Porting on Cartesius and Marconi:

We used the Cartesius Haswell partition, and Marconi KNL. We also used Cartesius Sandy Bridge partition for the Waal_40m test case, since it has higher memory requirements. The nodes on these systems have significantly different architectures, and also different instruction sets (AVX2 for Haswell, AVX512 for KNL). To successfully port D-Flow FM on both systems, we had to install dependencies (METIS, NetCDF) and tune compilation flags. We used D-Flow FM version 1.2.1.62244 on both systems, Intel compiler 18.0.1 and Intel MPI 2018.1.163 on Cartesius, Intel compiler 18.0.5 and Intel MPI 2018.4.274 on Marconi.

Detailed profiling to get an overview of the bottlenecks in D-Flow FM:

With profiling and in-depth analysis of the results we obtained a clear overview of the current performances and bottlenecks of D-Flow FM. The use of different representative test cases ensures that most of the use cases have been covered by our analysis.

This will help decide the future optimization and development effort of Deltares.

The main bottlenecks are: (not ordered)

- cache efficiency, the code is memory bound,
- PETSc linear solver strong scaling limit, the code is MPI bound, most of the MPI time is in global communication,
- mesh imbalance for strong scaling.

Other aspects that were analysed include IO and vectorization.

The performances and scalability are good when the computation load is high enough. For example, with an unstructured mesh consisting of 1000000 cells, the scalability is good up to ~256 MPI processes. As a rule-of-thumb partitions should have more than ~25000 cells.

Attempts at mesh imbalance improvement
We worked on improving the mesh imbalance arising from the mesh partitioning and large halo regions. We investigated two strategies:

- Halo-aware repartitioning using an algorithm inspired by "A. Sarje, S. Song, D. Jacobsen et al. Parallel Performance Optimizations on Unstructured Mesh-based Simulations. Procedia Computer Science Volume 51, pages 2016-2025, 2015."
    - The resulting partitions including halo region are better balanced. However the imbalance on the internal subdomains (excluding the halo regions) increases.
    - This does not lead to improved performances as the linear system solver (PETSc) only uses the internal cells and accounts for a significant portion of the total execution time.
    - The communication load does not improve either as this algorithm does not try to reduce the sizes and the imbalance of the halo regions.
- Using a hypergraph instead of the classic dual graph to represent and partition the mesh. Hypergraphs are able to model the communication volume accurately and allow for more complex partitioning objectives.
    - The resulting partitions including halo region are better balanced. The imbalance on the internal subdomains (excluding the halo regions) increases, but not as severely as with the halo-aware repartitioning. The halo regions are smaller and better balanced.
    - This leads to improved performances for Lake Marken test case with 1024 partitions, which is a very imbalanced case with MPI communication accounting for 69% of the total execution time (and MPI_Wait alone accounting for 32%).
    - Unfortunately we were not able to test this on our large test cases due to the limitation of the hypergraph partitioner hMetis.

The Halo-aware repartitioning approach does not lead to good performances as it increases greatly the imbalance on the inner subdomains.

The hypergraph showed promising results on a very badly balanced case, but the partitioning tool we used (hMetis) is not suitable for our larger test cases so we could not test is further.

Use of a linear solver with better scalability: Pipelined Conjugate Gradient

We did experiments with PETSc's linear solver KSPPIPECG (see [15]). This is supposed to improve performances of the conjugate gradient method when communications are important.

However, we were not able to get any performance improvement or to reproduce the speedup published by the authors.

- We performed micro benchmarks to test our architectures (Cartesius Haswell and Marconi KNL) and MPI implementations for asynchronous collective performances (using OSU benchmarks).
- We tested CG and PIPECG on a simple example, failing to get speedup with PIPECG over CG with single threaded MPI.

In addition, the performances seem highly dependent on MPI implementation and architecture.

Main results:

The current version of D-Flow FM can be used on the Cartesius Tier-1 system for parallel simulations using up to a few hundred cores with good efficiency, as long as partitions have more than ~25,000 computational cells. On Marconi KNL, the performances are not as good. The current version of D-Flow FM is not targeted for good performance with the specific capabilities of the KNL architecture. Moreover, the KNL architecture is being discontinued by Intel. However, it was very useful to have access to two systems with different architectures to compare the behaviour of D-Flow FM, get a good overview of its performance and identify the bottlenecks.

Several previous scalability analyses showed similar results. However, these analyses have been carried out using a lower number of cores. With the assigned PRACE systems the project was able to perform further analysis on Tier-1 and Tier-0 systems. This yielded new insights and also provided the opportunity to test possible improvements (for instance the PETSc's linear solver KSPPIPECG) in D-Flow FM for real-life applications.

As explained before D-Flow FM is the hydrodynamic module of the Delft3D Flexible Mesh Suite. For typical real-life applications, for instance for highly detailed modelling and operational forecasting, there is urgency to make D-Flow FM more efficient and scalable for high performance computing. The selected test cases from model applications Deltares is currently working on imposing computational challenges for Deltares. For these test cases of the current project it has become clear which further steps have to be taken towards Tier-0 systems.

In September 2019, the applicant applied for a new PRACE Preparatory Access Type D project to perform these further steps (2010PA5047). They have ideas to tackle, with a combined effort of SURFsara, CINECA, and Deltares, the most important bottlenecks.

Based on intermediate results of this PRACE project the applicant already applied for computational resources from DECI with an important application of the Global Tide and Surge model to contribute to the next IPCC 6th assessment report.

The project also published a white paper, which can be found online under [14].

2.9 Cut-off June 2018

2.9.1 BiqBin solver, 2010PA4347

Overview:
The BiqBin application is a high-performance solver for linearly constrained binary quadratic problems. This software is open source and available as an online solver.

Strong scaling experiments and performance analysis for the BiqBin v1.0 and v2.0 were made. A hybrid prototype was created. The prototype shows how to decrease a number of MPI communications and how to change the structure of the BiqBin application in general.

Scalability results:
An original BiqBin application v1.0 before optimization had already almost ideal scaling. After changes in a procedure of graph node generation to utilize all involved processors more efficiently (see [16] for more details), the application v2.0 showed superlinear scalability. This can be explained by the non-deterministic behaviour of the involved branch&bound tree traversal as part of the BigBin solver: if one of the workers quickly finds a stable optimal set, this helps pruning branches of another worker, thus acting advantageously on the pruning process in the parallel algorithm.

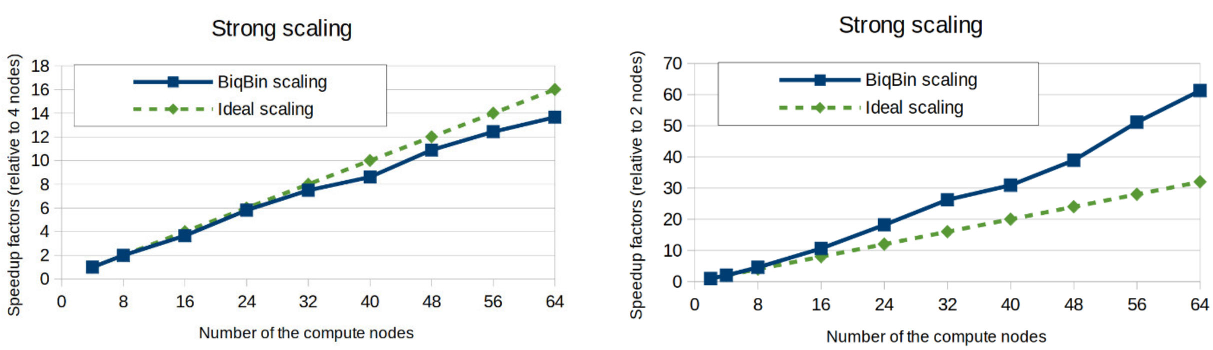

Figure 7: Strong scalability of BiqBin v1.0 (left) and v2.0 (right) for the sparse_graph250_density10 input file

Accomplished work:
The switch from BigBin v1.0 to v2.0 already provided a significant performance benefit, but MPI communication still shows a significant impact (~40% of the total runtime). The BiqBin application is too complex to target this problem directly, thus all the changes were illustrated on a prototype. The programming model was changed from MPI to MPI+OpenMP, overlapping computations and communications were implemented, and MPI P2P calls were replaced with MPI collective calls. Figure 8 and Figure 9 explains the design of the communication scheme.

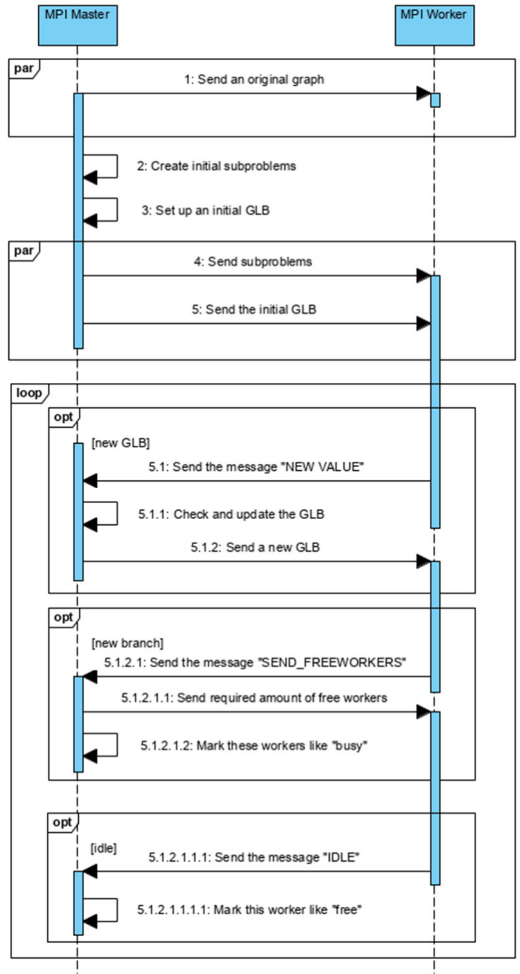

Figure 8: UML sequence diagram of MPI processes interaction in the prototype.

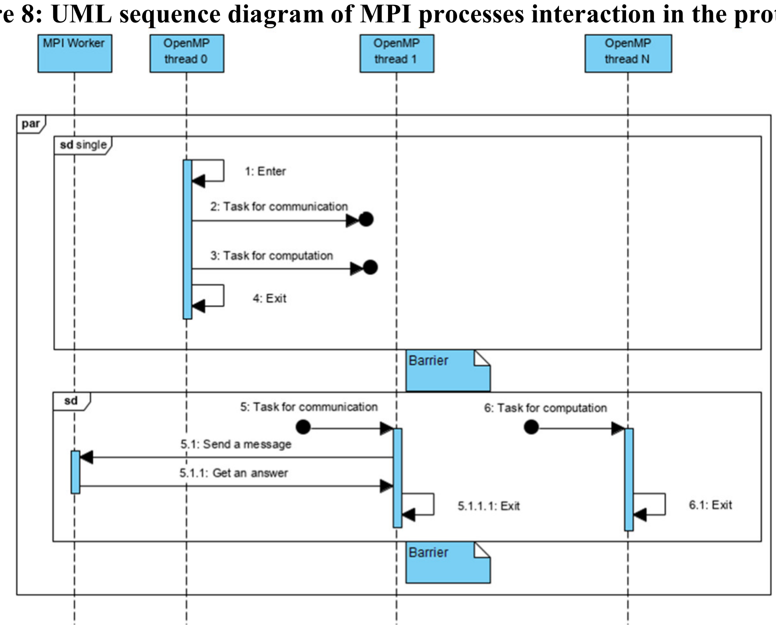

Figure 9: UML sequence diagram of the creation and execution of OpenMP tasks in the prototype.

Main results:
The scalability behaviour of BigBin was analysed, which highlighted the benefit of BigBin v2.0 towards v1.0 but deeper analysis also highlighted a significant communication overhead.

A change of programming model from MPI to MPI+OpenMP should significantly improve the application runtime. This requires the implementation of a new algorithm which was designed in the form of a prototype application but was not benchmarked so far.

The project also published a white paper, which can be found online under [16].

2.10 Cut-off September 2018

2.10.1 HOVE2 Higher-Order finite-volume unstructured grids, code enhancement for compressible turbulent flows, 2010PA4528

Overview:
Unstructured grids are widely used in science and engineering for representing complicated geometries in an accurate and efficient manner. The arbitrariness of the grid itself poses a series of challenges in terms of memory footprint and patterns, as well as development of numerical methods and computing algorithms aiming for high accuracy and computational efficiency.

Previous developments of the UCNS3D CFD code under a PRACE type C project associated with optimisation of the implementation of very high-order numerical schemes for unstructured meshes resulted in an 8.5 speedup. This was achieved by restructuring some of the computational intensive algorithms, employing linear algebra libraries and combining the state-of-the-art parallel frameworks of MPI and OPENMP. These developments have been applied to Large Eddy Simulations of canonical flows and Reynolds-averaged Navier–Stokes simulations of full aircraft geometries during take-off and landing.

This project aimed to enable extremely large scale simulations by focusing on the mesh partitioning algorithms and the I/O of the UCNS3D CFD code in order to perform ILES simulations on unstructured meshes in the scale of billion cells with very high-order finite volume methods. This is going to enable us to improve our understanding of the aerodynamic performance of complicated geometries and therefore enhance their efficiency.

Scalability results:
The scaling results were limited to the performance of the I/O operations of the code, and the performance of some initialisation processes that took considerable time due to the nature of the global arrays employed.

The following test also includes I/O times for a relatively large mesh of 230 million cells in the original version of the code prior to any performance enhancements.

Number of cores	Wall clock time before optimization	Wall clock time after optimization
1680	146 minutes	22 minutes
4800	662 minutes	16 minutes

Table 4: UCNS3D scalability tests on Hazelhen

Accomplished work:

1) Implementation of Parmetis in the code to partition larger meshes than 25 million cells that metis was limited to
2) Change of all the large input files to binary format
3) Rewriting sections of the code and eliminating global arrays that were limiting the maximum number of cells that could be used
4) Implementing MPI-IO for writing and reading checkpoint output files, and paraview output files

Main results:

The results obtained provided speedups of more than 130-250 for the size of problems concerned, for the writing and reading of large files (35GB per file). Also the speedup of the initialisation was 15-30 times with respect to the original version of the code due to the elimination of the global arrays.

The project also published a white paper, which can be found online under [17].

2.10.2 HPC Performance improvements for OpenFOAM linear solvers, 2010PA4548

Overview:

OpenFOAM acts as a major player in the Open Source CFD arena, due to its flexibility, but on the other side due to its complexity it is hard to define performance figures and scaling. One of the main bottlenecks for a full enabling of OpenFOAM for massively parallel clusters is the limit in the parallelism paradigm, that is the Pstream Library, which limits the scalability up to the orders of a few thousands of cores.

The proposed work aimed to create an interface to external solver libraries such as PETSc/Hypre thus providing the users a greater choice and flexibility when solving their problem cases, and to utilise their respective community's knowledge which has been developed over decades and not currently accessible within the OpenFOAM framework.

The main objectives of the current project are:

- *ldutocsr*: implement an interface in OpenFOAM to convert the Matrix from the native LDU format into the CSR format of PETSc
- Quantify the benefit of using external libraries as linear algebra solver (PETSc, Hypre, etc.) vs. OpenFOAM standard solvers
- Examine the effects of lower-level core changes within the matrix structures from the LDU to Compact-Storage-Row (CSR).

Target: increase the performances, in terms of execution time, for very large test-cases (at least 50M of cells) running on a massively parallel clusters (order of thousands of cores).

Scalability results:

See section “Main results”.

Accomplished work:

The test-case is a modification of the XL Lid-Driven cavity flow.

The set-up has been chosen to have the same residual levels among the different solvers with a fixed tolerance (i.e. 1e-04). The Courant number has been fixed at 0.25, well below the stability condition. The number of iterations have been fixed to 100, to have a reasonable execution-time to compare the results.

Only the pressure solver is taken into account in this study. The velocity solver is fixed with the values available in the public repo: DILU-PBiCGStab solver with tolerance = 0, and fixed maxIter = 5. The profiling of such tests-case has shown that the time spent in the pressure solver is around 67%, while the time spent in the velocity solver is around 27%.

The aim of this work was to compare different linear algebra solvers to minimize the execution time

The accomplished work is the following:

- An interface to an external linear solver packages has been created and tested, on a private repo https://develop.openfoam.com/Community/external-solver, accessible upon request.
- The converter modifies the matrix format from the LDU to the matrix data structure specified in the dictionary, that is CSR.
- The whole bunch of solvers and preconditioners can be invoked with a dictionary or a rc type file, for usability purpose
- When using the PETSc, the residuals are computed in the same way as done by OpenFOAM, that is the scaled L1 norm.
- The test-cases used in this study have been made publicly available by means of the OpenFOAM HPC Technical Committee: https://develop.openfoam.com/committees/hpc/tree/patch-1
- Work presented at the OpenFOAM conference: https://www.esi-group.com/it/resources/bna-abstract-openfoam-2019
- Possibility to use external linear algebra solver packages (CSR based) within OpenFOAM
- Mat data structure increases vectorization for a serial test case from the 3,9% of standard OpenFOAM up to 19.8% of the PETSc enabled version.

Main results:

Figure 10 shows the comparison of the execution times in seconds for the different pressure solvers under investigation by increasing the number of cores. Please note the figure is in the semi-log y scale. The data labels are only for FOAM-DIC-PCG solver (dark blue) and PETSc-AMG-CG+CACHING (green).

For low number of cores, i.e. 8 nodes the IC based solvers (FOAM-DIC-PCG and PETSc-ICC-CG) are much slower than the MG based solvers (FOAM-GAMG-PCG, PETSc-AMG-CG). A factor of around 8x is observed among the two classes.

The FOAM-GAMG-PCG does not scale by increasing the number of nodes, with a slight drop in execution time at 16 nodes, and a marginal increase at 128 nodes.

The PETSc-AMG-CG solvers (pink column) scale ideally up to 64 nodes, whereas at 128 nodes the time is almost the same compared to the 64 node test-case. This is, probably, due to the time spent in the set-up of the preconditioner, which does not scale optimal..

In fact, the PETSc-AMG-CG + CACHING (green column) solver decreases the execution time up to 128 nodes. Preliminary profiling shows that most of the coefficients of the preconditioner matrix do not need updates. Hence there is a decrease compared to the non-caching case, where the update is done at every iteration.

At 128 nodes, the IC based solvers scales super-linearly, as shown in Figure 10. In that case the PETSc-AMG-CG + CACHING is still the fastest solver, with a factor of 1.8 x compared to FOAM-DIC-PCG, and a factor 7x compared to FOAM-GAMG-PCG.

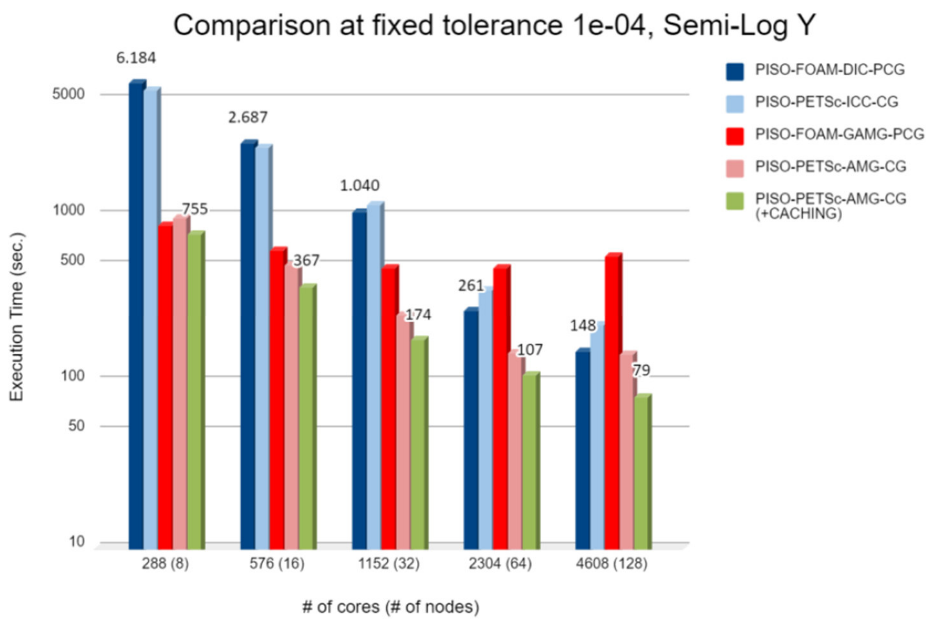

Figure 10: Comparison of execution time at fixed tolerance

Figure 11 shows that the time per time step is very short at 128 nodes. For the PETSc-AMG-CG+CACHING it is less than 1 second of average elapsed time per time step, with around 13,000 cells per core. Such time is too short for further investigation at an increasing number of nodes.

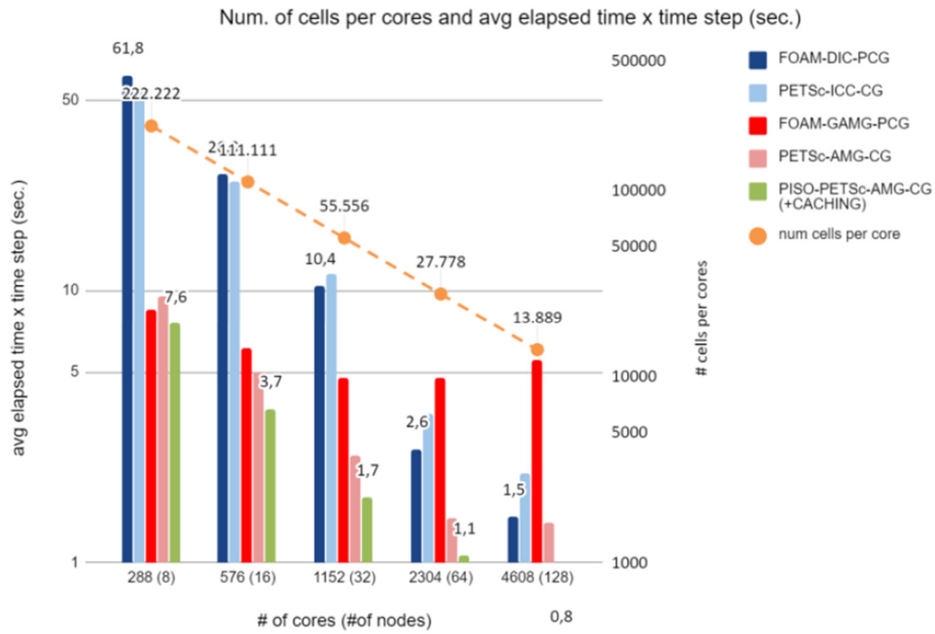

Figure 11: Number of cells per core and average elapsed time x time step (sec.) versus number of cores (nodes). Semi-log Left Y axis.

Further and on-going work is to profile the individual times of different meaningful portions of the code (initialisation, conversion, pressure solver, velocity solver, etc.) and monitor the influence of the available memory bandwidth.

The project also published a white paper, which can be found online under [18].

2.10.3 HPC-MetaShot: a multi-node implementation for metagenomics analysis, 2010PA4547

Overview:

Metagenomics analysis deals with the characterization of the microorganisms living in environmental samples. In particular, shotgun metagenomics by high-throughput sequencing of the total genetic material contained in the investigated samples allows the deep and accurate characterization of total microbiomes, including bacteria, viruses, protists and fungi.

In the recent years, the metagenomic approaches allowed to investigate the microbiome colonizing different types of habitat from the ocean to the acid mine drainage, with a particular interest in the community living on the human body and widely referred as human microbiome.

The investigation of the human microbiome allowed for example to partially discover the tight relationship between the physiological host process and the bugs colonizing the gut, as the digestion and the training of the immune adaptive system. Moreover, we also learned how in some cases these microorganisms became dangerous or how the microbial communities shape changes during chronic inflammatory diseases.

However, the analysis of such sequencing data is still extremely challenging in terms of both overall accuracy in the taxonomical and functional profiling and computational efficiency.

Shotgun metagenomics, in fact, is particularly useful for the identification of unculturable microorganisms and a potential field of application is the investigation of infective pathologies. Particularly interesting is the investigation of the micro-environment in which infective agents act and/or the interactions with other players, especially in case of co-infections (bacteria and virus, for example).

This approach does not require any prior knowledge of the target agents and allows to analyse the whole collection of genomes present in a sample, overcoming the principal limitations of the classical tools for pathogen detection, that are target-oriented.

MetaShot is a bioinformatic workflow designed for assessing the total microbiome composition from host-associated shotgun sequence data, by using sequencing data produced by Illumina platforms, that showed overall optimal accuracy performance (10.1093/bioinformatics/btx036). It is designed to analyse both genomic and transcriptomic data.

In MetaShot all the steps required for the analysis of "Next generation sequencing" data are embedded in a single Python pipeline managing the execution of external tools and the parallelization of the processes. Currently, the tool design forces its application to the analysis of few samples at a time, principally due to limits in the implemented parallelization scheme. This aspect was targeted in this preparatory access project.

Scalability results:
Within a strong scaling analysis, shown in Figure 12, different data distributions were tested, by splitting the global input file into a different amount of data chunks and distributing those amongst a different number of processes.

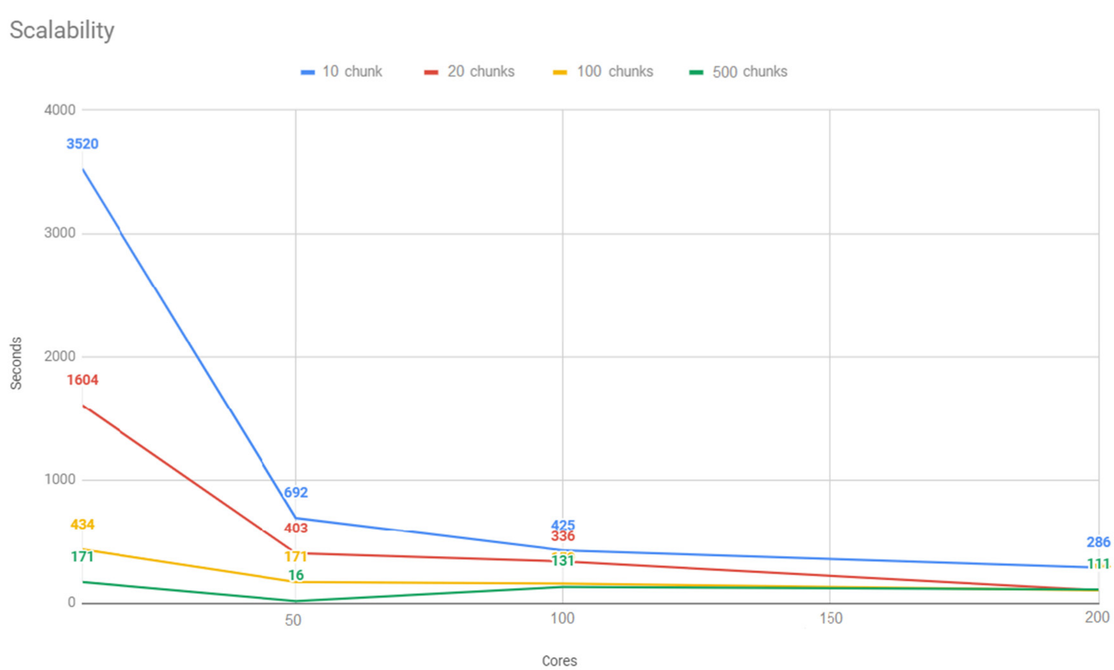

Figure 12: Strong scaling of HPC version of MetaShot utilizing different numbers of data chunks.

Accomplished work:

In order to develop a HPC version of MetaShot, capable of carrying out the analysis of tens/hundreds of samples in a few hours, we first concentrated our efforts on parallelizing the very computationally demanding part of Metashot: the alignment phase of input against the microbial and human genomes. The first genome is about 500 GB in size and the mapper involved in this analysis is bowtie where the second is only 3 GB, and the used mapper star. We developed a MPI program that uses the input of the fastq file of each sample and a configured aligner (i.e. bowtie or star).

A master process M divides the input into n (configurable) smaller equal fastq files and dispatches each mapping against the microbial genomes to n free parallel slave processes by sending a COMPUTE message and until all processes are kept busy. Each slave process S performs the alignment of each smaller file and produces an intermediate temporary mapping result file. To promote process recycling, whenever a slave process S completes its analysis, it notifies the master process M by sending a DONE message which in turn assigns a new alignment to S (if any). When all intervals have been analysed, the master process M sends a FINISH message to each slave process S notifying the end of the computation so that they can gracefully exit. A final procedure is implemented to collect all intermediate mapping results and to create a single, unified output with all potential reads mapped against the genome.

We have achieved our goal of getting a parallel program that can completely exploit the available computational resources of the HPC cluster. Since this was the first goal of our development activity, our expectations were completely met in this first phase. The second part of the work involving the integration of the MPI program into the Metashot pipeline is not yet finished. The cause of the delay in integrating our MPI program for the alignment phase is due to the rewriting of the serial code of the Metashot pipeline by the developers. The rewriting is in progress to adapt the use of Metashot to several sequencing platforms (such as Ion Torrent for input data in single end mode). The original code was written for Illumina sequencer platforms with input data in paired-end format. The tuning of the code for this important modification required more time than expected. It is worth noting that this is an external activity carried out by the Metashot developers and not planned within the framework of the HPC version of the code, proposed in this preparatory PRACE project. Now the goal continues to be the integration of MPI code in the pipeline but it needs to be done in a stable and reliable version of Metashot. For this reason the activity has been limited to the Tier-1 resources and further activity on the HPC version of Metashot is suspended until the serial Metashot is stable and produces reliable results.

Main results:

Successfully parallelization of the first step: alignment of reads. Not complete HPC code due to delay in the tuning of a new version of the code capable of analyzing Ion Torrent reads (external activity carried out by the Metashot developers and not planned within the framework of the HPC version of the code, proposed in this preparatory PRACE project).

2.11 Cut-off December 2018

2.11.1 Accelerating a high order accurate CFD solver using GPUs, 2010PA4774

Overview:
Heterogeneous HPC architectures are increasingly prevalent in the Top500 list (137 systems in the latest released version, up from 110 six months ago), with CPU based nodes enhanced by accelerators or co-processors optimized for floating-point calculations. For GPU accelerators, applications are typically rewritten in a low-level language such as CUDA or OpenCL. On the other hand, OpenACC enables existing HPC application codes to run on accelerators with minimal source-code changes. This is achieved with compiler directives and API calls, with the compiler being responsible for generating optimized code and the user guiding performance only where necessary.

ESSENSE is a computational fluid dynamics (CFD) code for the simulations of compressible flows. The code is used in a broad range of applications, including the study of aerodynamics, aeroacoustics, the climate modelling and in applied mathematics research. In the code, the Navier-Stokes equations are discretized in space by using a high-order finite difference method, which is based on summation-by-part (SBP) operators. The SBP operators can be rewritten as sparse matrix-vector products.

In this project, we took advantage of our previous work and made use of the optimized results from that work to port the code to GPU-accelerated systems. The project focussed on porting and optimizing the most time-consuming parts of ESSENSE to the GPU systems, namely the sparse matrix-vector products and evaluation of inviscid and viscous fluxes using the OpenACC directives and OpenACC API with NVIDIA CUDA libraries (e.g. cuSPARSE). Performance analysis of the ESSENSE code using PGPROF was carried out. We obtained a maximum speed-up of 61.3 using a V100 GPU in comparison to the results using the serial CPU version.

Scalability results:
The performance tests were evaluated on two different GPU systems; a Cray XC40 system with NVIDIA Tesla P100 GPUs and an IBM Power9 system with NVIDIA Tesla V100 GPUs at BSC.

We compared the total execution times in seconds for 10 timesteps using different order SBP operators on different grid points. The speed-up with the GPU P100 increases with the number of grid points and a maximum speed-up of 40.2 was obtained. The speed-up increases from around 11 with 33x33x33 grid points to 40 with 161x161x161 grid points. With the 6th order SBP-scheme, it takes 487.9s on a single Intel Broadwell CPU core with 161x161x161 grid points while it only takes 7.96s on the P100 GPU. A similar trend for speed-up and performance can also be observed on the CTE-IBM Power system with V100 GPUs. For the case of 161x161x161 grid points, the speed-up is 61.3x on NVIDIA V100 GPU in comparison to the IBM Power9 CPU core.

The performance results are as expected, i.e., the performance on V100 GPU is better than that on P100 GPU, and the performance is better using more grid points (larger matrix sizes) on GPU.

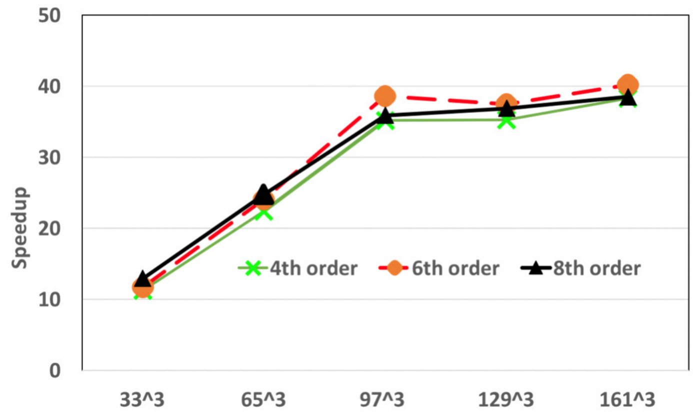

Figure 13: Speed-up on a P100 GPU in comparison with serial CPU version using different number of grid points

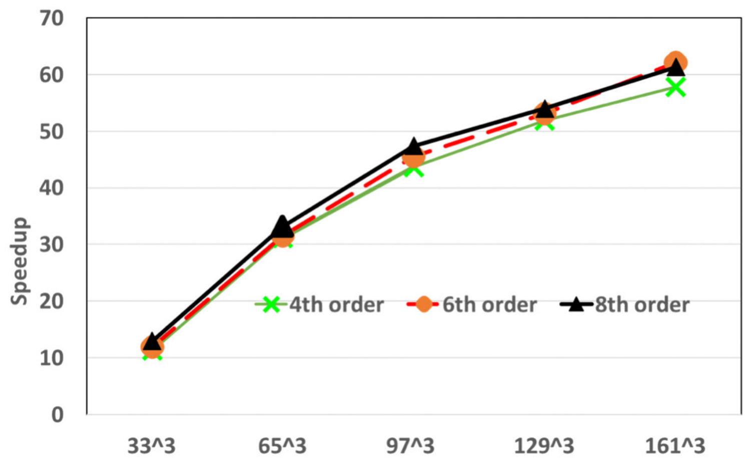

Figure 14: Speed-up on a V100 GPU in comparison with serial CPU version using different number of grid points

Accomplished work:

To achieve the target of scalability and improve those bottlenecks, we have performed:

- Optimization for OpenACC directives for the SpMV on GPU: With the PGI compiler, we used the OpenACC clauses NUM_WORKERS and VECTOR_LENGTH to manually control the occupancy in order to improve the performance of the SpMV. The OpenACC directive $ACC LOOP SEQ has been used in the inner-loop to avoid parallelization within threads. Because we observed that the OpenACC clause REDUCTION(+:sum) and ATOMIC directive performed worse for our cases in practice, they were not used in the studies.
- Data Management: To reduce data movement between hosts and devices, we have created data regions using the OpenACC data directive, by combining OpenACC unstructured enter data and exit data directives, as well as the PRESENT data clause. Then we have indicated to the compiler that the data is already present in the device memory in order to

keep all the data on the device operated by GPU. This would allow for data to be initialized on the device at the start of simulation, and for data to be retained in the device memory for the lifetime of the application. In this OpenACC implementation, most arrays and multi-dimensional pointers, including temporary arrays, have been created only once in the global memory within the subroutines setInitialDerivedConditions before the main loop. After that, these arrays and pointers will be reused and updated in the device global memory. When the preset number of timesteps or convergence conditions have been achieved, the variables in the device memory will be updated (copyout) to the host. In addition, ESSENSE has FORTRAN F95 features with object-oriented programming models. With the data movement of multi-dimension pointers within the object-oriented programming models, the compiler must assume that other pointers point to the same memory, thus making the inner loops non-parallelizable. We have used the CUDA Unified Memory (`-ta=tesla:managed`) for the deep copy in the port. The CUDA Unified Memory does have the caveat that only dynamic data (allocatable) is managed. With the CUDA Unified Memory, one pointer still needs to be manually managed using data regions but all allocatable members of the point can be managed by the run-time.

- In order to avoid a significant modification of the original code, we have added the ACC ROUTINE directive to instruct the compiler to build a device version of the subroutine, which can be called from an OpenACC parallel loop.

Main results:
The ESSENSE solver has been ported to GPU systems using OpenACC compiler directives. The work was focused on porting the most time-consuming parts of ESSENSE, namely the sparse matrix vector multiplications and evaluation of inviscid and viscous fluxes, to the GPU systems. A performance analysis of the ESSENSE code using pgprof was carried out. We obtained a maximum speed-up of 61.3 using a V100 GPU in comparison to the results using the serial CPU version. In the future, we plan to implement a multi-GPU version and to continue optimizing the computation kernels. The main results have been published in [19].

2.12 Cut-off March 2019

2.12.1 NEMO performance, 2010PA4821

Overview:
NEMO (Nucleus for European Modelling of the Ocean) is a state-of-the-art modelling framework for research and forecasting activities.

It allows various configurations based on the different modules available: OCE (core ocean model) SEA-ICE (thermodynamics) and TOP (tracers), with various horizontal and vertical resolutions (grids). It is parallelized by means of the MPI library: each processor computes the model equation over a subdomain (or tile) of the whole model domain.

Since the configuration of this model can vary widely a reliable performance analysis should comprehend different resolution/module sets. It would also be desirable to keep such a study as simple as possible: with this goal, the BENCH test was created.

Anyway, to understand if such a usage is feasible a preliminary study is needed.

The study is subdivided in the following tasks, agreed upon with E. Maisonnave, author of the BENCH test:

- The BENCH configuration offers the possibility to modulate the amount of MPI communications (or even to totally disable them) so as to evaluate their impact on NEMO's scalability. The first goal was to extend this feature to the SEA-ICE module, since it was available just for the OCE.
- The second task was to study the feasibility of extrapolating conclusions on NEMO's computational performance starting from studies conducted on the BENCH test.
- A third task consisted in assessing the cost of storing output when running at ORCA12 resolution, when the output is handled by the XIOS I/O server.
- The last assignment was to evaluate the weight of the north fold boundary condition computation (created to handle the north boundary of the ORCA tripolar grid) on the overall performance of a run.

Scalability results:
The scalability tests performed within the scope of this project had the objective of evaluating if the scalability of the BENCH configuration mimics, in some sense, the one of a real NEMO configuration.

What is remarkable here is that the two configurations differ in a fundamental aspect: the BENCH is not performing the 'land removal process', i.e. the removal of sub-domains entirely constituted of land points, since it has no bathymetry nor continents. This results in a different domain decomposition using the same number of parallel resources, thus in a different load of cells per core and, consequently, the number of Simulated Years Per Day (SYPD) per core will not be the same.

What it is instead expected to be similar is the amount of SYPD per domain decomposition size.

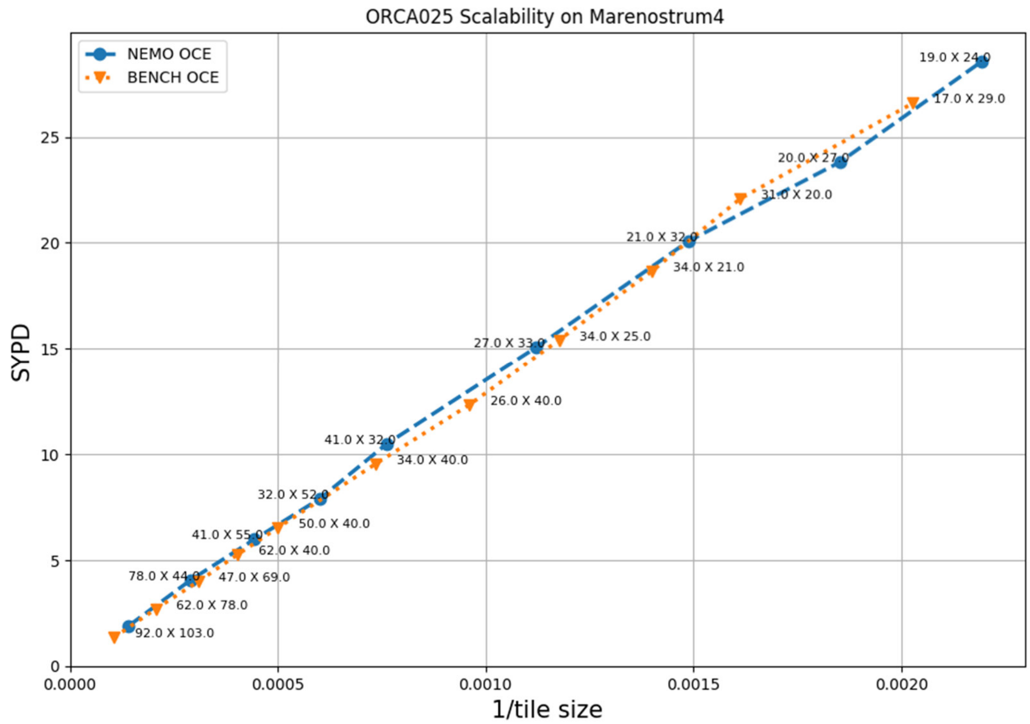

Figure 15: Scalability of NEMO using ~28 km grid resolution on MareNostrum4. Comparison of strong scaling of NEMO code with OCE component on an ORCA025 grid (75 vertical levels, 200 s. time step) with the BENCH test run on the same grid.

The results (Figure 15) are shown as SYPD per tile size, and not as SYPD per core or node. The reason is that, on the same number of cores, the two configurations will show a different domain decomposition, thus a different load per core, resulting in scaling curves with different slopes. This way of presenting the data allows showing that the two models have the same scaling properties.

We took the chance to investigate the duration of the time step length, which variability limits the performance. We ascertained that a fundamental role in this is played by the number of resources used per node: fulfilling all the available cores leads to processes competing for resources. The best compromise between performance and machine usage, on MareNostrum 4, is 46 processes per node. Using this layout the variability is reduced, and does not increase as the number of cores increases, pointing the finger at the operating system as the probable culprit for the observed noise.

Accomplished work:
We present the obtained results following the same order of the aforementioned planned tasks:

The BENCH test, as is in the trunk of the NEMO repository, could be deployed without particular problems on the MareNostrum 4 platform. After following the instruction that E. Maisonnaive gave us, the module mpp_lnk_generic, and mpp_nfd_generic were modified. The goal was to run the model without calling the Routine_LNK: these are general purpose interfaces that manage all the communication between the subdomains in the time step loop. Since the SEA-ICE communication pattern does not include collective operations, no more actions were needed to extend the no-communications option to this configuration.

To assess whether the BENCH test can be considered computationally similar to a standard NEMO run, and in this case, whether it can be used to study its performance, we generated the scaling plot presented in Figure 15. Both configurations can be considered similar as long as we measure the throughput as SYPD with respect to grid points per core.

All the measurements and results shown so far are not taking into account the burden implied when the output process is enabled. We measured how costly it would be to run the BENCH configuration on an ORCA12 grid, without ICE. We increased the output frequency to one record per time step on purpose, so as to increase the impact on performance. We tested both XIOS2.0 and its newer version, XIOS2.5, and reported here the steps per second measured with the different configurations, fixing the number of BENCH processes to 1200.

NODE-BENCH-XIOS processes	No XIOS	XIOS2.0	XIOS2.5
30 - 40 - 1	0.4	0.02	X
30 - 40 - 2	0.4	0.03	X
60 - 20 - 1	0.7	0.04	0.1

Table 5: Steps per second, testing different XIOS configurations

All the reported results are taken in “one file” mode, i.e. not using NetCDF parallel library, as this would bring a further slowdown of ~30%.

The results marked with an ‘X’ are experiments which run out of memory.

To evaluate the influence of the north fold execution we disabled the communications just there, and compared the resulting scaling curve on an ORCA1 grid with a real one: at 25 nodes the difference between the two curves measures around 2%, at 60 nodes already around 10%.

Main results:

We have shown how the BENCH test can accurately represent the computing performance of NEMO itself, given the correct way of interpreting the data, and spares the experts, working on optimization tasks, to struggle with the search for input data, or with domain decomposition issues. What we think is still lacking, is a working ORCA2 grid configuration, which is another of the five build-in ORCA configurations of NEMO.

This study also leads to the following conclusions:

- The best way to run NEMO, in all its flavors, on MareNostrum 4 is not to completely fill the nodes, but instead leaving some room for the OS to work. To give some numbers: running an ORCA025 resolution on 160 nodes (7680 cores) brings to 45 SYPD using all the available resources, while to 67 SYPD when leaving one core empty per socket in each node. Moreover, in the latter configuration we can detect good scaling performance up to 9,000 cores, while in the first case already at 6,000 cores the code does not scale anymore.
- When running an ORCA12 configuration, the time-steps involved in output operations will slow down, with respect to a computation-only one, from seven to twenty times. The XIOS

version is also crucial: it looks like the newest version needs around double the memory compared to the old one. Not exploiting the NetCDF library parallel capabilities also seems like a good idea.

- The north fold extra communications become a burden for the model at high numbers of cores, a region where the communications in general are already the bottleneck. In this sense it is a further remark of the fact that whichever attempt to optimize the model should target this topic.

The project also published a white paper, which can be found online under [20].

2.13 Cut-off December 2019

2.13.1 Intelligent HTC for Committor Analysis, 2010PA5236

Overview:
Committor analysis is a powerful, but computationally expensive tool, to study reaction mechanisms in complex systems. The main goal of this project is to facilitate an implementation of committor analysis in the software application OpenPathSampling (OPS, http://openpathsampling.org/) that is performance portable across a range of HPC hardware and hosting sites.

The committor can also be used to generate initial trajectories for transition path sampling, a less-expensive technique to study reaction mechanisms. The software developed here is being used to generate initial trajectories to study a conformational change in the main protease of the SARS-CoV-2 virus, which causes COVID-19 (this is a modification from the original goal when the proposal was written). This conformational change may regulate the accessibility of the active site of the main protease, and a better understanding of its mechanism could aid drug design.

Committor analysis is essentially an ensemble calculation that maps straightforwardly to an HTC workflow, where typical individual tasks have moderate scalability and indefinite duration. Since this workflow requires dynamic and resilient scalability within the HTC framework, OPS is coupled to a custom HTC library (jobqueue_features, https://github.com/E-CAM/jobqueue_features) that leverages the Dask (https://github.com/dask) data analytics framework and can manage MPI-aware tasks.

The project targets porting the custom HTC library to a wider variety of HPC platforms, specifically additional resource managers and MPI runtimes. Both OPS and jobqueue_features are Python-based and we have also developed more sophisticated support for Python-based tasks, including direct access to the memory space of executed tasks to avoid the use of filesystems for data transfer. We also investigated the use of the UCX protocol for communication in the HTC framework to reduce the overall overhead of the HTC framework.

Scalability results:
The jobqueue_library had not really been stress-tested for the number of tasks it was capable of supporting. Previous efforts had looked at up to 2000 tasks, in this case we scaled out to 1M tasks on all available architectures, with each individual task using 2 nodes worth of resources. Each set of two nodes forms a worker with the workers reused by different tasks. The number of workers

that the framework can use is entirely dependent on the maximum number of allowed simultaneous jobs by the specific site.

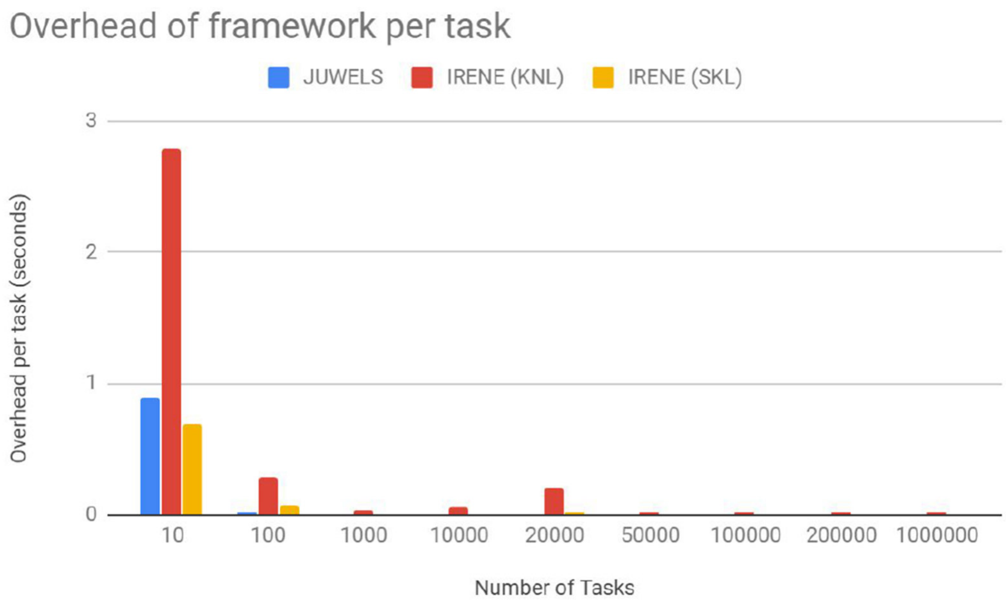

Figure 16: Framework overhead

The overhead of the framework is negligible at ~1ms for the Skylake architecture and ~10ms for the KNL architecture.

To test the efficiency of the UCX protocol in comparison to TCP, we used several test cases with different kinds of computation, from basic hand written functions to those based on functions provided by Dask. On most of these tests, UCX performance is not very different to TCP performance (which uses IPoIB). In general, UCX does give slightly better results for tasks based on Dask functions and TCP performs better for hand written tasks. Critically, UCX has been seen to lead to some errors related to the fact that it is not yet a fully supported technology within Dask. In particular, during tests there were problems in the case of restarting the cluster while UCX was in use, which means that we lose resiliency in this scenario. Given the limited performance improvement and the significant impact of sacrificing resiliency, we would not recommend the use of UCX with Dask at this time.

Accomplished work:

Leveraging the Dask framework involved significant software development within OPS. Two main problems had to be addressed: (1) OPS had no mechanism to gather results reported from parallel workers into a single result set, and (2) tasks based on OPS could not be serialized by Dask. Both of these problems were solved by building atop an ongoing overhaul of the OPS storage and serialization subsystem. These have been combined with a new implementation of the committor simulation to make a usable parallel committor simulation for OPS.

For jobqueue_features, we did the following:

- Support for MPI-enabled tasks (i.e., tasks which leverage MPI4PYfor parallelisation) was added. The framework has the ability to directly access the memory space of the root MPI process.
- Support for the PBS resource scheduler was added.
- Continuous integration support for SLURM and PBS was created. This means that the library features are fully tested (automatically) on both of these schedulers through the use of custom Docker containers.
- Support for OpenMPI, Intel MPI and MPICH MPI runtimes was added (including runtime specific process distribution).
- A Docker-based tutorial infrastructure was created which includes a SLURM scheduler. The infrastructure includes two computation nodes and one head node. On the head node next to SLURM was installed a JupyterLab instance with which the users can try our package using their browser. We intend to create some tutorial Jupyter notebooks to facilitate use of this infrastructure.
- Support for Dask and Dask.distributed 2.x (tested up to version 2.20, released in July 2020).
- Configuration of the HTC library for JUWELS and IRENE
    - IRENE uses both a custom scheduler (based on SLURM) and a custom MPI runtime, we were able to override the more generic classes of jobqueue_features to support these

Main results:
MPI-enabled task support is now tested in continuous integration and was the main test case used for our scalability studies. The framework reuses the MPI environment created when the worker initialises on the node for each task. We have tested the framework with PyLAMMPS to show that each task does indeed have access to, and use, all available resources (this information is something that LAMMPS reports in its output).

We were able to easily scale our task workload out to 1M tasks with the MPI-enabled task overheads plateauing at about 1ms for Skylake architectures (JUWELS and IRENE) and 10ms for KNL (IRENE). Such overheads are negligible even for short task durations. Node configuration times on JUWELS is about 10s, and on IRENE is about 6s for the SKL, and 29 second for the KNL. CPU time savings for short task durations are, therefore, potentially substantial when using the framework.

UCX support was found to be somewhat unstable (currently) and beneficial only in specific scenarios that make heavy use of Dask objects. Resiliency is (currently) compromised when using UCX, as such we would not recommend its use until this issue is resolved.

OPS now has an experimental parallelized committor simulation. This will be utilized in a project studying the SARS-CoV-2 main protease. In the future, the approach used here can be adapted to parallelize other simulations performed in OPS, such as replica exchange transition interface sampling.

3 T7.2 Applications Enabling Services for Industry

In this section the progress and present status of the T7.2: Applications Enabling Services for Industry ("SHAPE") task will be described. Section 3.1 gives an overview of the processes involved in SHAPE along with some important statistics right up to the eleventh (most recent) call. The SHAPE+ initiative is described in section 3.2 together with its impact on introducing new countries to SHAPE.

A summary of the present status is given in section 3.3. Projects from the sixth and seventh calls completed around a year ago, and follow-up requests have been sent to these SMEs and associated PRACE centres to see how the work performed with the assistance of SHAPE has affected the SMEs' business – these are reported in section 3.4. Summary reports for the more recent eighth to eleventh call projects are reported in section 3.5 along with an overview of the lessons learned with regards to the implementation of the SHAPE programme. The twelfth call, soon to be opened, is described in section 3.6.

In September 2020 the Working with SMEs (WP7-WP2-WP3-WP6-aisbl) workshop was held. The aim of this workshop was to replace the session that would have taken place at the cancelled PRACE 2020 "All Hands" meeting. Participation at this meeting was opened up to participants from the EuroHPC Competence Centres (EuroCCs) and talks were given as follows:

- Welcome (Florian Berberich, JUELICH, PRACE-6IP)
- Overview of the PRACE SHAPE programme (Chris Johnson, EPCC, PRACE-6IP)
- Overview of EuroCC (Bastian Koller, HLRS, CASTIEL)
- Experiences with SHAPE and setting up a Competency Centre (Claudio Arlandini, CINECA, PRACE-6IP)
- Experiences with engaging with SMEs (Wahid Rofagha, PRACE aisbl, PRACE-6IP)
- Discussion (Moderator: Chris Johnson, EPCC, PRACE-6IP)

In the discussion of the final session of the workshop, we discussed ways in which SHAPE and the EuroCCs can work together as well as discussing improvements for SHAPE. This is discussed in section 3.7.

When applicants are unsuccessful in being awarded a SHAPE project, they are informed of this together with the reasons why. Where possible, SMEs are directed to other calls within PRACE or beyond. In the case where SMEs are requesting significant HPC resources, they may be directed towards calls focussing specifically on offering CPU time, such as PRACE Project Access or DECI. In other cases they may be more appropriate for Preparatory Access. For the most recent SHAPE calls, one SHAPE proposal was instead moved to Preparatory Access without the SME needing to be involved. Future smoothing of this mechanism is also discussed in section 3.7.

3.1 SHAPE Overview

SHAPE (SME HPC Adoption Programme in Europe) is a pan-European initiative supported by the PRACE project. The programme aims to raise awareness and provide European SMEs with the expertise necessary to take advantage of the innovation possibilities created by High-Performance

Computing (HPC), thus increasing their competitiveness. The programme allows SMEs to benefit from the expertise and knowledge developed within the top-class PRACE Research Infrastructure.

There are many challenges in getting SMEs to consider adopting HPC, most notably a lack of in-house expertise or a lack of available staff effort. SMEs often have little or no ready access to suitable hardware, and an SME may consider that the risk of committing time and effort to HPC when they have no prior experience is too high when there is no guarantee of a suitable return on investment. However, utilising HPC gives an SME the potential to improve product quality via an enhanced performance and accuracy of their models, or by reducing time to delivery, or by providing innovative new services to their customers. Ultimately, this can in turn increase their competitiveness. SHAPE is there to help overcoming the barriers faced by SMEs and allow them to get a foot on the HPC ladder.

Successful applicants to the SHAPE programme get dedicated effort from a PRACE HPC expert as well as access to a suitable supercomputing resource, or other appropriate hardware, at a PRACE centre. In return the SME commits a comparable amount of effort and provides their domain expertise. In collaboration with the SME, the PRACE partner helps them try out their ideas for utilising HPC to enhance their business. So far, SHAPE has awarded 62 SMEs effort across 11 calls and the 12th call for applications will open in October 2020.

Table 6 shows the calls, applications, approved projects and person months committed from PRACE so far in SHAPE and the chart in Figure 17 shows the countries of the SMEs of awarded projects.

Call	Call open	No. Proposals	No. Awarded	PMs
Pilot	June 2013	14	10	35
2	Nov 14 – Jan 15	12	11	45.25
3	Nov 15 – Jan 16	8	8	30.75
4	Jun 16 – Sep 16	8	4	17
5	Mar 17 – Jun 17	7	6	20.75
6	Oct 17 – Dec 17	5	2	9.5
7	Apr 18 – Jun 18	6	3	16.5
8	Oct 18 – Dec 18	1	1	3
9	Apr 19 – May 19	7	5	20.5
10	Oct 19 - Dec 19	5	5	18.75
11	Apr 20 – Jun 20	10	7	28
	Totals	83	62	245

Table 6: SHAPE proposals received and awarded by call

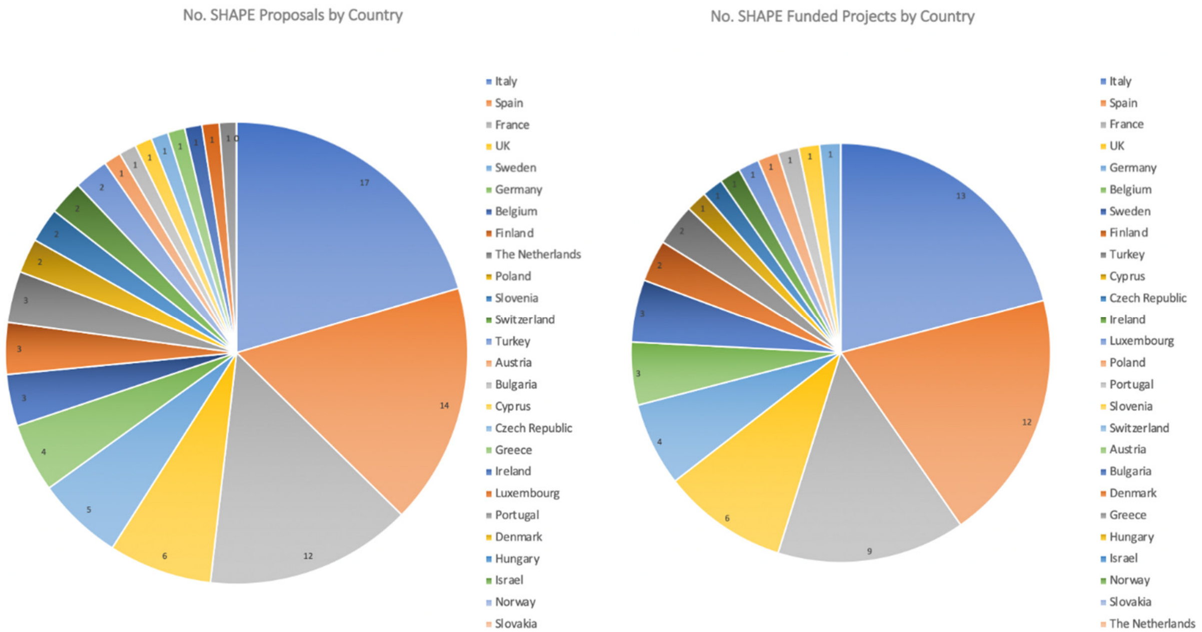

Figure 17: SHAPE proposals received and awarded by country

3.1.1 SHAPE Process

The process for SHAPE is as described in [5]:

- **The SHAPE Call is launched**: There is a form to be filled in (by the SME, although assistance from PRACE is available if required) which can be completed online or downloaded and submitted via email – the form contains suggestions and information on how to complete it.
- **The Call closes**: The applications are reviewed by the SHAPE review panel (see below).
- **Recommendations**: The review panel makes their recommendations on which projects to approve, and the PRACE BOD ratifies this.
- **Pairing**: the successful projects are matched to PRACE partners who have effort available and relevant expertise.
- **Machine time**: with help from the PRACE partner the SME applies for machine time on an appropriate system – this is usually via Preparatory Access type D, but other arrangements are possible where agreements can be reached between the PRACE partner and the SME.
- **Coordination**: The SME and PRACE partner do the project work, and the SHAPE coordinator monitors progress.
- **Deliverables and output**: During the project the SME is expected to publicise their interaction with PRACE and SHAPE. This can be via their own website, press releases, publications in their field, and so on. In addition, they are expected to contribute to providing information for PRACE deliverables (such as this document) on the status of the project, in collaboration with their PRACE partner.
- **Conclusion**: At the end of the project, it is expected that a white paper will be produced detailing the technical work and results, which will be made publicly available on the

PRACE website following internal review. In addition, approximately 12 months after the completion of the work, a follow-up report will be requested from the SME to try and evaluate the Return on Investment of the SHAPE work.

3.1.2 Review of Applications

The composition of the review panel was described in PRACE-5IP D7.2 [4]. The exact membership of the panel will vary with each call mainly due to availability, but typically it is now:

- One member from the Industry Advisory Committee,
- Two Business Development Officers (replacing two members of the PRACE Board of Directors)
- One member from the PRACE Peer Review Team,
- Two members representing PRACE-6IP-WP7,
- One member representing the preparatory access team.

The process for applying for SHAPE is two-stage. The SME applies to SHAPE, then if successful applies for machine time, usually via Preparatory Access (PA). This final panel member is included to give a preliminary opinion on the technical suitability of the application from the PA point of view and then to facilitate the application for machine time. This two-stage process has been streamlined for the last few calls to minimise the burden on the SME.

The criteria used in the review remain as described in PRACE-5IP D7.1 [5] but are included here for completeness:

- ***Fit with the goals of SHAPE***
  - The aim of SHAPE is to assist SMEs in overcoming the barriers to adopting HPC, such as risk, initial cost, lack of in-house expertise and lack of access to resources.
  - The target audience for SHAPE is SMEs with no or little experience of HPC, who can be assisted with both expertise from PRACE, and time on a PRACE machine.
  - It is not a way of getting processor cycles for production runs for existing codes – there are other avenues for this via PRACE such as the project access calls.
- ***Strength of the business case***
  - The expertise and resources provided through the SHAPE programme are expected to produce a significant Return on Investment for the company. In the mid-term, the SME should be able to build on the results to, for instance, increase its market share, renew its investment, offer new products or services, or recruit dedicated staff. The solution implemented should be part of a business plan to further engage in HPC in the long term.
- ***Technical Achievability***
  - The proposals are expected to be realistically achievable in the timescales described and with the resources made available.
- ***Other aspects considered***
  - The commitment of the SMEs to co-invest with PRACE in achieving the project goals. The effort for the project should be at least approximately equally split between the SME and PRACE.

- The innovative aspects of the proposal.
- The social and economic impact on society as a whole.

The applications are reviewed and ranked according to these criteria, then the final recommendations are put forward to the PRACE Board of Directors for approval. Once the BOD gives approval, the SMEs are informed of the decisions, together with the reasons for the decisions made and, if unsuccessful, any future recommendations for resubmission or redirection to another call.

3.2 Increasing the number of participating countries: SHAPE+

PRACE recognises that SHAPE should assist as many SMEs in as many eligible countries as possible and this has been raised by PRACE reviewers at PRACE-5IP project reviews. With regard to the countries of original, there has been a recognition that SHAPE projects were far more likely to originate from countries where SHAPE effort from PRACE partners was already available than from countries where PRACE centres did not have SHAPE effort assigned (or had little effort). Figure 17 shows the countries where the awarded SHAPE projects come from. The correlation between the number of proposals and effort available is not surprising as a PRACE centre is far more likely to encourage a proposal from their country if they have the effort to then work with the SME. This was discussed at IAC in November 2018 where the idea of the SHAPE+ initiative was born.

The idea of SHAPE+ is that an allocation of person months is made available to centres for projects received where

- the SMEs are from countries where no (or very few) SHAPE projects have been awarded so far;
- the centre in question does not presently have enough effort to carry out the SHAPE project.

At the management board in May 2019 it was agreed that 30 person months would be made available for this initiative. The SHAPE+ initiative was implemented for the SHAPE 10th and 11th calls and will be continued for at least the 12th call.

There are good signs that this initiative is working. Call 10 received the first proposals from Slovenia and Turkey, and call 11 received the first proposals from Cyprus and Portugal. Slovenia had some existing SHAPE effort to use, so this resulted in projects from Turkey, Cyprus and Portugal being considered to be eligible for SHAPE+ using 14.5 of the available person months so far.

Figure 18 shows the number of countries of origin of SMEs in terms of proposals received and projects awarded. The number of countries represented has increased over time and that increase has accelerated for the last few calls. As of call 11, we have now received proposals from 21 of the 26 eligible countries, resulting in projects from 17 different countries. This leaves just SMEs from Denmark, Hungary, Israel, Norway, and Slovakia to produce proposals and for proposals from Austria, Bulgaria, Greece, The Netherlands to turn into projects.

Within SHAPE we will continue to focus on encouraging proposals from countries so far under-represented and look forward to building new links with SMEs from such countries.

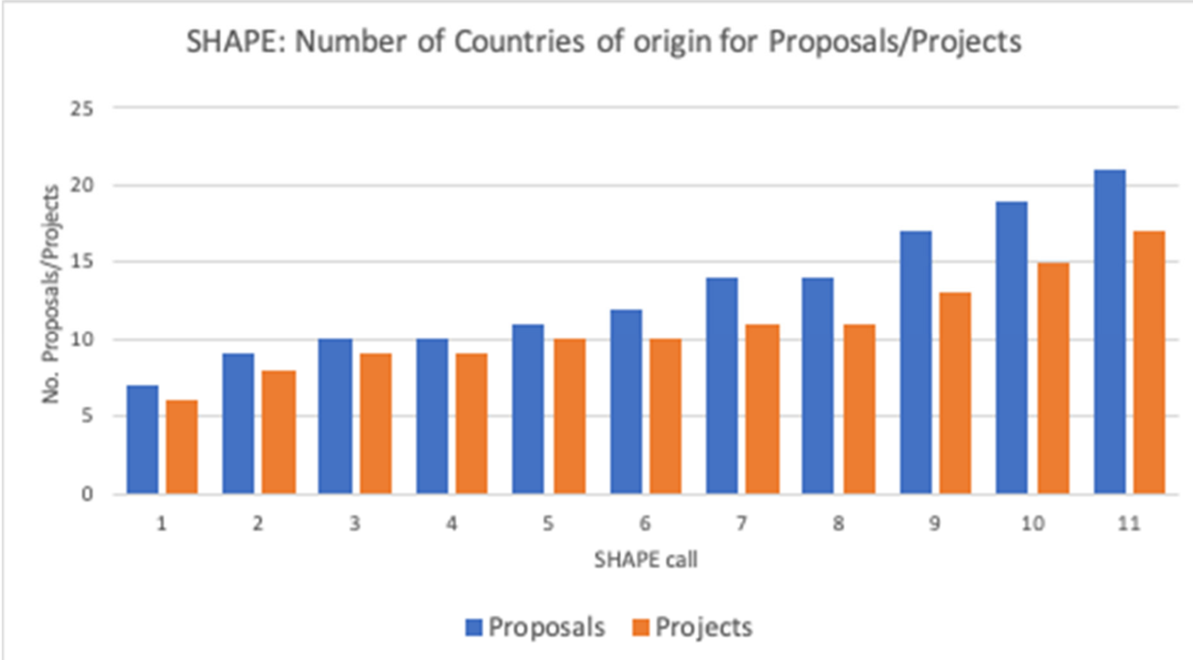

Figure 18: Graph of the increasing number of countries of SHAPE proposals and awarded projects

3.3 SHAPE Programme status

In this section, an overview of the status of all the SHAPE projects since the first (pilot) call are reported. SHAPE projects from the second call onwards fall into three categories:

1. Projects that finished 2 years ago or more which were fully reported in previous deliverables (e.g. PRACE-5IP D7.2 [6]). These are projects from the second to the fifth call.
2. Projects that finished around a year ago (sixth and seventh calls). Both the SMEs and PRACE partners from these projects have each filled in a survey to assess the impact of the SHAPE work on their business and the results are reported on in section 3.4.
3. Projects that are ongoing or in the process of finishing (eighth to the eleventh calls). PRACE partners from these projects have each provided a summary of their project. These are reported on in section 3.5.

On conclusion of the work, each project is expected to produce a white paper for review and publication on the PRACE website. These can be found here [12]

As of September 2020, the status of all of the projects is shown in Table 7.

Call	SME	PRACE Partner	White Paper	Report
2	WB-Sails	CSC	Published	See [5]
2	Principia	CINES	Internal technical report	Not Provided, see [5]
2	Algo'tech	INRIA	Published	Not Provided, see [5]
3	ACOBIOM	CINES	Published	See [5]

3	Airinnova AB	KTH	Published	See [5]
3	Creo Dynamics AB	KTH	Published	See [5]
3	AmpliSIM	IDRIS	Published	See [5]
3	ANEMOS SRL	CINECA	Published	See [5]
3	BAC Engineering Consultancy Group	BSC	Published	See [5]
3	FDD Engitec SL	BSC	Published	See [5]
3	Pharmacelera	RISC	Published	See [5]
4	Artelnics	BSC	Published	See [4]
4	Milano Multiphysics	CINECA	Published	See [4]
4	Renuda UK Ltd	EPCC	Published	See [4]
4	Scienomics	IDRIS	Published	See [4]
5	Disior Ltd	CSC	White Paper not expected due to change in focus of SME	See [4]
5	Invent Medical Group, s.r.o.	IT4Innovations	Published	See [4]
5	AxesSim	CINES	Published	See [4]
5	E&M Combustion S.L.	BSC	Published	See [4]
5	Svenska Flygtekniska Institutet AB	KTH	White Paper not expected due to SME unavailability	See [4]
5	Shiloh Industries Italia s.r.l	CINECA	Published	See [4]
6	Axyon AI SRL	CINECA	Published	See section 3.4.1
6	Vision-e S.r.l.	CINECA	Published	See section 3.4.2
7	Briggs Automotive Company Ltd	Hartree	White Paper awaiting final approval	See section 3.4.3
7	Polyhedra Tech SL.	BSC	Published	See section 3.4.4
7	FLUIDDA	BSC	Published	See section 3.4.5
8	Energy Way Srl	CINECA	White Paper not expected due to change in focus of SME	See section 3.5.1

9	Neuralbit Technologies	PSNC	This project has not yet started.	See section 3.5.2
9	NIER Ingegneria S.p.A.	CINECA	Tech work still ongoing	See section 3.5.3
9	LVD Biotech, S.L.	BSC	Tech work still ongoing	See section 3.5.4
9	Submer Immersion Cooling	BSC	Tech work still ongoing	See section 3.5.5
9	SPARC Industries SARL	University of Luxembourg	White Paper in production	See section 3.5.6
10	OPEN ENGINEERING	IDRIS	Tech work still ongoing	See section 3.5.7
10	Global Surface Intelligence	EPCC	Tech work still ongoing	See section 3.5.8
10	Enlighting Technologies SL	BSC	Tech work still ongoing	See section 3.5.9
10	SPARK inovacije, d.o.o.	UL	Tech work still ongoing	See section 3.5.10
10	Tarentum	UHeM	Tech work still ongoing	See section 3.5.11
11	VIPUN Medical N.V.	UAntwerp	Project starting	See section 0
11	Mobildev	UHeM	Project starting	See section 3.5.13
11	Global Maritime Services Ltd (G.M.S.)	EPCC	Project starting	See section 3.5.14
11	Offshore Monitoring Ltd. (OSM)	CaSToRC	Project starting	See section 3.5.15
11	d:AI:mond GmbH	HLRS	Project starting	See section 3.5.16
11	CENTIMFE	UC-LCA	Project starting	See section 3.5.17
11	SmartCloudFarming GmbH	HLRS	Project starting	See section 3.5.18

Table 7: Complete list of SHAPE projects awarded to date

3.4 SHAPE Sixth and Seventh call: Follow up for completed projects

This section focuses on SHAPE projects completed in 2018-2019 and have thus had time in the 12-18 months since the projects finished to assess the impact of HPC on their business. For the last similar deliverable [4] in consultation with the PRACE IAC, two on-line surveys were constructed, one for the PRACE partner to fill in, the other for the SME. The PRACE partner survey gathers information about the technical improvements arising from the project while the SME survey gathers information about the business improvements. The template for the surveys were given in the appendix of [4]. This survey was a success and provided useful quantitative and qualitative information and so was repeated for this deliverable. The only changes that were made were to add the following question to the SME's survey:

"Please could you describe where you first heard about SHAPE?"

And for existing questions, to include popular responses given in "other" responses last time as selectable options.

Five relevant PRACE partners and SMEs were surveyed. All five of the partners responded and four of the five SMEs responded. The full responses are presented below in sections 3.4.1 to 3.4.5 below and contain information gathered in both the PRACE partner survey and the SME survey. The results are then summarised in section 3.4.6.

3.4.1 Axyon AI (Italy): Axyon Platform Optimisation

Project Partners

- Jacopo Credi , Axyon AI
- Eric Pascolo, CINECA
- Claudio Arlandini, CINECA

Project Overview

Axyon AI brings deep learning-powered solutions to finance. Their main product is a proprietary deep learning platform specifically built for financial problems, developing highly-accurate predictive models. Axyon developed successful proofs of concept in several fields, from credit risk to wealth management, from churn-rate prediction to fraud detection.

The project aims to improve the performance of the Axyon platform in the following fields:

- Support of large dimension datasets;
- Optimize learning jobs distribution among several computational nodes;
- Optimize the learning process to better leverage hardware capabilities.

Project Successes

Axyon runs multiple instances of Axyon neural network training, so the performance increases thanks to HPC. The work had a big impact on the Axyon company because it has allowed them to use supercomputer and cloud remote resources for neural network training in an easy way.

HPC usage

Axyon reported that they are continuing to use HPC, in particular Cloud-based services and that they plan to use HPC systems mainly to pursue computational tasks to take advantage of multi-GPU and multi-node infrastructures. Before the project they had no staff with HPC skills. After the project they hired two people skilled in HPC.

Business impact of the SHAPE project

Axyon has reported a measurable return on investment with an estimated 15% increase in revenue. This has come about due to cost reductions, faster time to market, and improved R&D processes, including Faster Deep learning models development.

3.4.2 Vision-e (Italy): Deep Tile Inspection

Project Partners

- Rudy Melli, Vision-e
- Eric Pascolo, CINECA
- Claudio Arlandini, CINECA

Project Overview

Vision-e is a Spin-Off of Università degli Studi di Modena e Reggio Emilia focused on studying, designing and developing computer vision systems and algorithms for custom industrial application of quality inspection. In the last 10 years, it has studied and developed a tiles inspection system software based on computer vision to detect surface defects of ceramic tiles on production lines.

One of the main problems for tile inspection systems, is the big variety of ceramic tiles that change in terms of colours, glossiness, texture or structures and each single product has dozens or hundreds of different surfaces (the shape above the ceramics). So, it's critical to classify correctly which surface belongs to each tile captured before to find defects.

The biggest challenge is the training process in which the system creates a memory (model) of each different surface.

Current training algorithms are based on classic computer vision and for tiny variations they cannot group correctly the surfaces and need human action to verify and correct the tile's surfaces classification, a "semi-automatic" process.

The goal of this project was to automate completely the training process, using a deep learning approach that should increase significantly the performance and the effectiveness with a big impact on defects detection.

Project Successes

The project has led to improved scalability of code with a speedup of 3.2 on 4 GPUs on a single node of the DAVIDE HPC Cluster compared to an office workstation. The project helps the SME to evaluate the computing power needed to improve the accuracy of results without a total runtime increase.

HPC usage

Vision-e was not using HPC before the project and is still not yet using it. However, they now have one new employee with HPC skills. They particularly noted that the communication with the HPC staff has been useful giving them access to technologies otherwise not considered that have increased their skills.

Business impact of the SHAPE project

Vision-e notes that their R&D process has been improved including industrial automation process control thanks to the SHAPE project.

3.4.3 Briggs Automotive Company (UK): Parallel CFD for single-seat supercar design

Project Partners

- Neill Briggs, Briggs Automotive Company Ltd (BAC)
- Greg Cartland-Glover, Hartree Centre, STFC Daresbury Laboratory
- Andrew Sunderland, Hartree Centre, STFC Daresbury Laboratory

Project Overview

Briggs Automotive Company (BAC) is the UK manufacturer behind the Mono, the world's only road-legal, single-seat supercar. The Mono has been designed, engineered and managed by experts with backgrounds in several leading automotive, high-end supercar and motorsport brands. The company has already achieved several world firsts – such as the use of graphene on a car's body and the implementation of hybrid carbon-composite wheels. BAC's integration of advanced design and engineering with ground-breaking innovation will enable it to follow the success of Mono with other focused, intelligent and exciting products.

This project aimed to apply computational meshing tools to a CAD-based representation of the Mono supercar in order to develop an accurate high-quality unstructured mesh to enable computational fluid dynamics (CFD)-based analysis. Parallel CFD software was then be applied to this dataset, exploiting large-scale computational resources to compute the high-resolution aerodynamic flow around the car. The parallel runs focused on analysing certain design features of the Mono supercar, with a view to improving the performance and efficiency of future models.

Project Successes

The project has provided a demonstration for a company that has not previous used HPC or a dedicated fluid dynamics solver.

BAC did not fill in the survey due to other business pressures.

3.4.4 Polyhedra (Spain): HPC optimisation of SDLPS distributed simulator

Project Partners

- Pau Fonseca i Casas, Polyhedra Tech SL
- Iza Romanowska, BSC

Project Overview

Polyhedra Tech SL builds simulation tools to aid small and medium size businesses partners in optimising their day-to-day operations and increasing their performance. Their expertise includes: optimisation of energy systems in terms of energy efficiency, consultancy for construction industry, and sustainability services and other. The core of the business model is the use of SDL as a modelling language and the SDLPS, a distributed simulator enabling automatic translation of models defined using SDL into working simulations which can then be parallelised.

The SDLPS is a simulator that allows the definition of the models using SDL language. This definition is complete and represents the model behaviour and structure, allowing its simulation without the need for implementing the model, while at the same time simplifying the validation and verification processes. The aim of the project was twofold. First, to investigate novel parallelisation strategies that can leverage an HPC platform such as MareNostrum 4 at BSC to accelerate the simulator runtime. Second, to test, benchmark and optimise the simulator on social science models requiring different parallelisation paradigms. Therefore, opening the HPC capacity to a new group of users while at the same time increasing the robustness of their software by improving the level of documentation, verification and validation of the models.

Project Successes

During the project, two tools were developed:

1. compiler source to source xml to C++
2. a library that allows parallel execution of SDL defined models

The SME now has the ability to run a previously serial simulation on HPC systems. The software was developed from scratch but it enables running of the SME´s code in an HPC environment so potentially improving performance by several orders of magnitude.

HPC usage

Polyhedra reported that they were using HPC before the SHAPE project and continue to do so now using Cloud-based HPC Services. They plan to continue to use HPC and are building an infrastructure that integrates with MareNostrum at BSC and allows the execution of models in a desktop environment on top of the MareNostrum infrastructure. Before the project they had one employee with HPC skills. Now they have three such employees.

Business impact of the SHAPE project

Polyhedra has acquired one new customer due to the SHAPE work and it has improved R&D process.

3.4.5 FLUIDDA (Belgium): HPC simulation of particles deposition in human airways

Project Partners

- Wim Vos, Fluidda
- Beatriz Equzkitza, Constantine Butakoff, BSC

Project Overview

FLUIDDA is a world leader in the field of Functional Respiratory Imaging (FRI) research and development. The company's proprietary FRI technology offers pharmaceutical companies and healthcare providers a unique entry point into personalised medicine for patients suffering from respiratory diseases and sleep-related breathing disorders. Implementation of FRI in the clinical practice creates significant added value to the current healthcare standard in the respiratory field.

In the proposal, HPC simulations of particle deposition in human airways were performed and validated. Different patient specific cases were considered. Computational meshes for each case will be constructed from medical images, simulations for different kinds of particles and inflow conditions were performed. The final goal of the project is to optimise the HPC simulation code, to validate its results and to establish a pipeline to enable the use of HPC resources in the company workflow. The code used is Alya, which is part of the benchmark suite of PRACE.

Project Successes
The SME now experiences improved resolution of simulations and the ability of a previously serial simulation to run on HPC systems as well as an improved workflow.

HPC usage
FLUIDDA reported that they were using HPC before the SHAPE project and continue to do so now using their own HPC system. In future they will probably use Cloud-based HPC services. Before the project they had two employees with HPC skills and this has now risen to three.

Business impact of the SHAPE project
The work carried out here has not yet noticeably changed the SME’s business model.

3.4.6 Summary of follow-up responses

The follow-up feedback from the SMEs has again generally been very positive with records to SHAPE. Here we present some overall analysis of the feedback provided by SME in terms of their previous and current HPC usages and the effects the SHAPE project has had on their business process.

SMEs were asked the following question

“With regard to your current HPC usage, are you

- *Not using HPC at all,*
- *Using Cloud-based HPC Services,*
- *Using your own HPC Systems,*
- *Using resources from a PRACE HPC centre,*
- *Other”*

As can be seen from Table 8 all SMEs have report an increase in HPC skills. Most are still using HPC (Cloud-based or the SME’s own system) and will continue to do so in the future.

SME Name	Using HPC before SHAPE?	Using HPC after SHAPE	No. Employess with HPC skills before SHAPE	No. Employess with HPC skills after SHAPE	Rise in employees with HPC skills	Do you have plans to do more with HPC in the future?
Axyon AI	Yes	Yes - Using Cloud-based HPC Services	0	2	2	Yes
Vision-e	No	No	0	1	1	No
Polyhedra	Yes	Yes - Using Cloud-based HPC Services	1	3	2	Yes
FLUIDDA	Yes	Yes - Using their own HPC Systems	2	3	1	Yes
	3 (Yes); 41(No)	3 (Yes); 1 (No)	0.75 (Average)	2.25 (Average)	1.5 (Average)	3 (yes); 1 (No)

Table 8: Responses of SMEs to SHAPE survey

SMEs were asked the following question

"How has the SHAPE project affected your business process?

- *Cost reduction,*
- *Faster time to market,*
- *More sales,*
- *Improved R&D process*
- *Improved service to customers*
- *Not at all,*
- *Other."*

From Table 9 the majority of SMEs have reported improvement in the research and development methodologies, as well as reports increased sales (one with a 15% increase in revenue) and cost reductions.

	Research and Development	More sales	Cost reduction
Axyon AI	✓		✓
Vision-e	✓		
Polyhedra	✓	✓	
FLUIDDA			
Total	**3**	**1**	**1**

Table 9: SHAPE survey response regarding improvements to business model

All SMEs surveyed reported that they would recommend SHAPE to other SMEs and gave positive comments. The surveys appear to be a good way of gaining feedback from SMEs and we plan to continue this with minor refinements for future deliverables.

3.5 SHAPE eighth-eleventh call: Project summaries

The projects in the section are mostly underway at present unless otherwise indicated. At the time of writing, the projects from the eleventh call are only just starting.

3.5.1 Energy Way Srl

This project will not go ahead due to a change of priorities of the SME.

3.5.2 Neuralbit Technologies

This project has not yet started and is awaiting availability of the SME.

3.5.3 NIER Ingegneria S.p.A.

Projects Partner:

- Cesare Sassoli (NIER)
- Laura Rocchi (NIER)
- Eric Pascolo (CINECA)
- Roberto Da Vià (CINECA)
- Claudio Arlandini (CINECA)

NIER is an engineering consultant company founded in 1977, with approximately 125 employees (15% of them with a Ph.D.) and offices in Bologna (headquarters), Milano and Napoli. NIER's main activities are in the Railway, Nuclear, Automotive, Oil&Gas, Biomedical and Manufacturing sectors.

The project aims at implementing a Deep Neural Network (DNN) which classifies human tasks based on eye-tracking (ET) during driving tests via video analysis. Deducing where a person is looking and which elements capture his attention can help in a different situation (e.g., knowing which elements of a dashboard capture a driver's attention can help in improving the dashboard design). Fast processing of this large amount of data is challenging: image analysis is data-intensive and highly compute-intensive mostly because of the volume of data; eye-tracking data are related to time and frequency. Hence, a parallel computing paradigm (e.g. data parallelism) was proposed, which would benefit from computationally efficient hardware such as HPC. The DNN model is being trained using data available thanks to the collaboration with DICAM (University of Bologna).

Activity performed

The project is still in progress. We have parallel pre-processing software to divide the video into frames and label them as DNN input. We are now focused on developing the neural network to classify where the driver is looking in every single frame. The classes to be identified are: road, cars on the road, sky, car mirror, car interior. The technologies used are: Python for programming, Keras and Tensorflow as Deep Learning libraries and for parallelisation we use python threads.

PRACE cooperation

For the given topics of the project, CINECA created a specific one-day training format for NIER partners. The course was focused on how to use HPC machines (environment, application, scheduler, remote visualisation) and on parallelism technologies and techniques. Furthermore, CINECA has been supporting NIER to parallelise both the pre-processing and Deep Learning applications and in the use of the Galileo HPC cluster.

Benefits for SME

With the present project, NIER means to develop the skills and the experience that will allow the integration of data-driven technology to the already consolidated services provided by the company, and also the inclusion of new services. Image and data analysis, based on AI methods, are particularly crucial for the development of the company since they will provide useful tools to optimise the solutions NIER offers to clients in different industrial and scientific domains.

For these reasons the impact of the participation in the PRACE project is seen as particularly relevant by company managers. First of all, it will allow them to explore the advantages of HPC in our applications, such as model's code-optimisation, data parallelism and applications' scalability. The activities foreseen during the PRACE project will facilitate the acquisition of new skills related to HPC and Machine Learning. The predicted impact includes the acquisition of new clients, provision of new activities and consequently the increasing of the company's potentials. The development of new services and products in NIER will potentially affect the organisation, even with the set-up of a new Business Unit specialized in ML with new personnel hiring. In addition, the SME is considering signal analysis as a further service they can offer.

Lessons learned

The project is still running, so it is premature to give feedback. For the company the participation in the SHAPE project was useful for the acquisition of new skills both learned during the initial course, but also working with the support of HPC specialists. Currently the company is finding it challenging to create a DNN that satisfies the needed accuracy, so in the last period, even if the interest is still high towards HPC technologies, the company has focused more on developing the DNN.

3.5.4 LVD Biotech, S.L.

Overview

LVD Biotech is a fast growing company founded in 2010 in Barcelona (SPA) with the aim of developing advanced medical devices and therapies for the treatment of disorders of the cardiovascular system. The team that started LVD Biotech has more than 20 years of experience in the field of cardiovascular devices, being responsible for achieving one of the first CE marks for a drug eluting stent more than 12 years ago. LVD Biotech was born to become a reference in the cardiovascular field by empowering the value of technology. They have an extensive knowledge in biopolymers, coatings, drug delivery, mechanical engineering, medicine and pharmacy, hence bringing together recognized experience in the development of medical devices and implants.

The contact person in LVD Biotech is Arnau Vidal, product manager (avidal@ivascular.es) and the BSC contact is Eva Casoni, postdoc researcher (eva.casoni@bsc.es).

The project aims to demonstrate the capability of advanced numerical simulation based on Finite Element modelling to investigate transcatheter aortic valve implantation (TAVI) designs, deployment geometries and the capacity of the stent of shape setting. Numerical simulation appears as a cheap option to study the behaviour of the devices and the different factors that may influence it, and in this study it is proposed as a support and guidance for the design process.

Activity performed

The Alya code has been adapted in order to simulate the physical response of the loading step of the stent into the deployment device, using a contact algorithm and large deformations approach.

First, a shape setting phase is performed simulating the process of deforming and cooling the stent in order to obtain the expanded geometry. Experimentally a mandrel and several heating and cooling steps are needed but with the developed code the final shape can be obtained in a single run, hence avoiding the multi-step process.

Second, the loading step in which the stent is loaded into the deployment device, a sheath, typically at a similar diameter to the original one. The shape of the sheath is one of the bottlenecks of the design process and one of the objectives of the study was to determine its optimal shape in order to guarantee a smooth deployment. Several shapes have been tested and it has been observed that the problem is very sensitive, especially because of the large deformations and the dependence on the boundary conditions. The initial mesh for the stent is an unstructured mesh of 167,809 tetrahedra. The material model is a hyperelastic model with the properties of the martensite phase of Nitinol. Each of these simulations has been performed in MareNostrum 4 with 4 nodes (192 cores) and the Alya code. One of the optimal shapes found will be described in the White Paper. The computation of radial and axial forces of the stent within the optimal deployment device is still a work in progress. Moreover, in order to improve the results, the implementation of the complex and highly non-linear material Nitinol or a shape memory alloy is also needed. To do so, more resources are needed.

PRACE cooperation

PRACE provided the expert support to adapt the Alya code for the application. The code has been adapted in order to deal with large deformations and a contact algorithm, developed in BSC. The approach puts special focus on its computational implementation and the use of parallel resources.

Despite not all the requested hours being employed, the machine and time needed to perform the simulations was also provided by PRACE. The assigned machine was Marenostrum IV at BSC and 3 TBytes of storage were available. Due to the non-linearity of the problems, the wide range of tests to be performed and the amount of intermediate steps that need to be stored, the use of PRACE resources have been essential for the development of the project.

Benefits for SME

The results of this SHAPE project are still preliminary but they provide guidance and support to the company in the design process. So far, the company has been using commercial simulation software to help the design process. However, it gives limited flexibility in the modelling part and presents difficulties in the convergence of large deformations and sensitive cases. The simulations performed with the Alya code for the loading case present good convergence and give more flexibility to adapt the code to the requirements of the SME. In addition, commercial codes fall short when very fine and accurate meshes are used (that will be needed for the stent modelling) and a wide number of test cases are required (for instance by considering the influence of several parameters in the design process). The HPC-code Alya covers these issues.

This project presents numerical simulation and High Performance Computing, a new tool for the SME to consider to include in its design process. Once the technical problems are solved and the experimental results are properly reproduced, we expect to perform a complete simulation of the loading process of the stent with a refined mesh in order to obtain more accurate results. In order to do so, more computational resources will be needed.

Once the ongoing work is finished, a meeting between the SME and BSC will be scheduled in order to consider future collaboration. The obtained results are very promising and it is expected that future studies under the HPC framework will improve the outcomes.

Lessons learned

The successful experience of this collaboration in the SHAPE projects framework has allowed the SME to make contact with the computer resources and to consider computer modelling and HPC as a tool to help in their design process. The flexibility of the in-house code and the adaptation to parallel computing of commercial codes is one of the main advantages. Nonetheless, since the code needed to be adapted to the application, technical problems appeared, which have somehow delayed the simulation plan. In future collaborations this complexity shouldn't be underestimated.

The unexpected health crisis (COVID-19) situation has made the communication between the SME and BSC more difficult and the expected outcomes have been more difficult to cope with. Despite periodic reports and teleconferences being scheduled, not all of them could be carried out and this delayed the planned work. A better and more precise definition of the problem may diminish the communication requirements. Moreover, periodic reports in a non-deep technical language between both sides should be considered in possible future collaboration.

3.5.5 Submer Immersion Cooling

Overview

Submer (https://submer.com), the clean tech company, designs, builds and installs Liquid Immersion cooling solutions for HPC, hyperscaler, datacenters, Edge, AI, deep learning and blockchain applications. Submer Immersion Cooling is changing how datacenters are being built from the ground up, to be as efficient as possible and to have little or positive impact on the environment around them (reducing their footprint, reducing their consumption of precious resources such as water). Submer was founded in 2015 by Daniel Pope and Pol Valls, to make operating and constructing datacenters more sustainable. Regarding this project, the contact person in Submer is Daniele Rispoli, distinguished engineer of the company (daniele@submer.com) and the PRACE contact is Guillaume Houzeaux, Team Leader at the Department of Computer Applications in Science and Engineering from the Barcelona Supercomputing Center (guillaume.houzeaux@bsc.es).

This project aims to increase the knowledge and expertise of Submer employees in the realm of CFD simulations in an HPC environment. Together with the solver availability, the educational factor, both on the application side and on how to work with HPC resources, is of the utmost importance for the company.

The project also aims to demonstrate tangible benefits in terms of improved design, efficiency and effectiveness of their immersion cooling solution. The vision of Submer is that, being able to model a reasonably close estimation to the operational behavior of their systems is key to avoiding costly mistakes, providing a better product to their clientele and greatly decreasing time to market.

Activity performed

The final objective of the project was the simulation of a complete cooling system in operating conditions, including 22 boards of CPUs and GPUs. To achieve this, different intermediate activities have been undertaken. First, CPUs and GPUs were characterised individually and a mesh convergence analysis was conducted. Then, we developed a heat sink and anisotropic porosity model for the heat and Navier-Stokes equations, respectively, to model the effects of the CPUs, GPUs and heat dissipators. The goal of such approximation is to greatly simplify the meshing and reduce the number of elements and thus the cost of the simulation. Heat sink and porosity were then calibrated against the first individual simulation. We then carried out single board simulations and validated the results by comparing the temperature peaks observed through experiments conducted by Submer. Then Submer decided to consider a new board that was validated as well. This new board was then placed into the cooling system and duplicated to obtain a complete configuration of the cooling System. We then carried out simulations with a 70M mesh, and to ensure mesh convergence, the same simulation on the multiplied mesh (70*8) of 560M elements. Results have enabled Submer to clearly identify the zone where heat is not properly evacuated.

All these activities were decided by agreement with Submer and BSC. During the project, we had weekly meetings in order to present the work carried out, but also to decide on the next steps, according to the results obtained. These meetings have also been used to educate the Submer team on the use of an advanced CFD solver for the numerical simulation of their applications on HPC resources.

PRACE cooperation

We obtained 200k CPU hours through a PRACE preparatory access (type B) on the MareNostrum IV supercomputer at BSC. Most of the hours were used for the fine mesh simulation of the complete cooling system.

Benefits for SME

Previously to the project, the SME was using a commercial code, offering limited access to modelling options and computing hours. During the project we have demonstrated that using well-calibrated models for the heat sinks and the porosity, we were able to reproduce the flow inside a complete cooling system.

The code Alya is now under an "available source licence" and can be used by the company on any supercomputer. We are indeed presently in contact with them to further collaborate on new designs. Through this collaboration, Submer could identify hot spots in their design. They are now evaluating modifications of the setup to optimise the heat evacuation. The numerical validation of their designs is for sure an additive value for their clients.

Lessons learned

What worked: A specific and adapted numerical strategy has been developed (simplified models for CPUs and GPUs) to carry out the simulation of a complete setup. This has been possible thanks to weekly telcos with the company and a straight follow-up of the advances of the project.

What could be improved: During the project, the company decided which simulations should be carried out and in which functioning conditions. But BSC had to be in charge of executing the simulations. However, since the beginning of the project, we have been updating a shared document, we detail all the computational details including the cost of the simulations, the number of cores used and the time to solution. In future projects we expect that the company will end up running the simulations independently.

3.5.6 SPARC Industries SARL

Overview

SPARC Industries is a 2 years old Startup and has two core activities:

1. development of a novel space propulsion system.
2. development and commercialisation of a plasma simulation tool for space and non-space flows at very low pressures.

The latter is in the focus of this project and is applicable specifically in the context of satellite propulsion design and satellite contamination assessment.

This SHAPE project investigates how an HPC platform, and the GPU specifically, can significantly increase the performance of an existing application, VSTRAP, owned and developed by SPARC Industries. The performance objective is to decrease runtime by 2 (or more) orders of magnitude. The application VSTRAP is an electrostatic plasma 3D simulation tool.

The project team is composed of Rabab Bouziane, Ullrich Siems and Dejan Petkow (CEO and contact for the SME). The PRACE assigned person is Frederic Pinel.

Activity performed

The activity was essentially the porting to GPU (CUDA) of a tree-code algorithm for electrostatic plasma simulation. The tree-code algorithm was first ported to GPU in a standalone prototype version. Once this version was satisfactory, performance-wise, it was then integrated in SPARC Industries' own plasma simulation tool VSTRAP.

The GPU implementation is verified via two benchmark test cases, which are a simple cube test case and a two-stream instability test case. The performance gain is measured for the simple cube test case to be a factor of 140-500 times faster compared to a serial execution on one CPU core. The two-stream instability test case for 1 million particles runs on a GPU 5.4 times faster than a parallel execution on a 20 core CPU and 100 times faster than a serial CPU execution.

PRACE cooperation

PRACE was involved in prototyping the GPU code for the electrostatic plasma simulation tool. Specifically, PRACE's role was to:

- Assemble a standalone version of each new algorithmic version of the plasma tool, from source extracts provided by the Sparc team.
- Port to GPU the standalone version, improve performance as much as possible without sacrificing code readability. This was performed on a PRACE HPC machine (uni.lu).
- Present the code to the Sparc team so that they could integrate the GPU standalone version back into their product.
- Provide guidance directly, or from other PRACE partners as required.

The Sparc team provided much help to explain the plasma simulator, and worked very closely with PRACE.

Benefits for SME

Sparc's business plan includes the increase of performance of its VSTRAP product, to meet future client requests (bigger models), and to offer several software delivery options (cloud, HPC, workstation). One possible solution is a GPU-enabled version. The SHAPE project is intended to explore this acceleration option. The successful port to GPU of a key algorithm of the plasma simulation tool allows such an evaluation, and offers a working solution. The GPU version is 145-500x faster compared to the sequential version, and once integrated into Sparc's product, 5.4x the performance of a 20-core machine. The SHAPE project helps Sparc's management defining a software architecture roadmap that will meet future needs, now that performance numbers are available on GPU.

Lessons learned

Overall, the project went very well. The organisation within Sparc and together with PRACE worked well. This is largely due to Sparc, which gave the right people the availability needed. The work was well organised, such that little overhead was necessary, and much could be performed independently. The proximity of the 2 groups (Sparc and PRACE) helped a lot: the PRACE representative could easily move between work locations (PRACE uni.lu HPC and Sparc offices). As always, the timescale was slightly off, due to various events, however the effort planned was enough.

A risk, which did not materialise, was the lack of domain knowledge of the PRACE representative (specifically: physics/plasma simulation). This was not a problem given the work organisation of the SHAPE project, yet, domain knowledge in addition to HPC skills would have been a plus.

3.5.7 OPEN ENGINEERING

Overview

The Open Engineering is a technology company that provides Computer-Aided Engineering (CAE) software tools and services in the field of multi-physics. This has applications in Sensors and actuators, including MEMS (Micro-Electro-Mechanical Systems) and Microsystems, Optomechanical systems including MOEMS (Micro-Opto-Electro-Mechanical Systems) and Multidisciplinary systems involving interaction between a fluid medium and other structures. Working together with IDRIS based in France, the SME wants to develop a multi-GPU version of their code based on a single-GPU and MPI versions that they already have. The plan is then to analyse the performance of the new code. Developments carried out during this project should improve performance and reduce memory usage which should then lead to faster, more accurate results thus improving the services provided to customers.

The main partners for this project are the following :

- Pascal De Vincenzo (Open Engineering – CEO)
- Santiago Costarelli (Open Engineering - Engineer/Technical contact)
- Isabelle Dupays (Idris – HPC Engineer)
- Rémy Dubois (Idris – HPC Engineer)
- Karim Hasnaoui (Idris/MdlS – HPC Engineer)

Activity performed

The main activities which have been performed till now are the following :

- Creation of the project and account at Jean Zay machine, with 1,000 GPU core hours and 50,000 CPU corehours. The PI and his collaborator have been informed of how to create accounts, allowing them to start the project (Task 2).
- Initial benchmarking and dynamic analyses (profiling and memory leak checks) of the code has been performed in order to identify bottlenecks (Task 3).
- Static analyses of the code have been performed in order to check the portability of the code (through static analyzers and different compilers). Therefore, recent C++ norms have been achieved by adding few corrections to the initial code (Task 3).
- A demonstrator for the initial implementation (asynchronous multiple devices execution and peer communications between devices) has been proposed to the SME partner (Task 4).

PRACE cooperation

Machine access was granted to the SME. Assistance to set up a working environment was provided by the IDRIS contacts at the request of SME partners. Realisation of analysis, benchmarking and demonstrator are being carried out by the IDRIS partner. The IDRIS partner is providing guidance for the parallelisation strategy.

Lessons learned

The initial delivery of the code has been greatly delayed due to the Covid-19. Once the project started, availability of the SME partners was extremely high. Moreover, the IDRIS team is also working intensively on the project in order to catch up with the original timetable.

3.5.8 Global Surface Intelligence

Overview
Global Surface Intelligence (GSI) is an earth observation company that uses satellite data to assess natural resources. Machine Learning and Artificial Intelligence tools are used. The SME has used a number of high-performance platforms (such as Cirrus, a Tier-2 system at EPCC) and primarily wants to improve their workflow to enable ease of use as well as some additional parallelism.

The project is being worked on by Neelofer Banglawala, and various members of the GSI operational team. Guidance is provided by Kevin Stratford at EPCC and Mark Howie (Chief Engineer) at GSI.

Activity performed
The work to date has focused on understanding the existing GSI workflows, and recasting them in a more robust and sustainable form, which includes the addition of unit and regression testing, and standard performance evaluation. We have also helped GSI to run their data science workflows on the NextGEN-IO [http://www.nextgenio.eu] machine in Edinburgh, a machine tailored specifically to problems with large data input and output requirements.

PRACE cooperation
The PRACE partner has arranged the access to the NextGEN-IO platform and assisted on using the system as required.

Benetits for SME
The work undertaken will form the basis of a more flexible framework for data scientists at GSI to perform their work without specific knowledge of the underlying computational resource. This is to be combined with other strands of development at GSI to enable the use of, e.g., Amazon web services.

Lessons learned
We have no specific feedback to the SHAPE process at this time.

3.5.9 Enlighting Technologies SL

Overview
Partners:

- SME:
    - Adrià Huguet
    - Pelayo Méndez
- PRACE (BSC):
    - David Vicente
    - Cristian Morales

This SME is a spin-off company from the University of Barcelona developing simulation software for the creation of Printed Circuit Boards (PCBs) and light scenarios (spectra) using different LEDs available in the market. The project here focuses on the installation of bio adapted luminaires at a hospital facility, where the SME plans to simulate various lighting scenarios.

The SME's codes are not multi-node parallelised and are mainly running on desktops and cloud computing platforms. The aim of the collaboration is work on the code to scale up on multi-node to running on HPC machines like MareNostrum 4 at BSC. Their codes are written in Python and they use the multiprocessing Python library as their main parallelisation strategy.

Activity performed

At the beginning of this activity, an analysis and profile of the code was performed using different python tools as PySpy and cProfile. Also, the code was studied in order to understand its structure and the possible approaches to maximize its performance and to find the best option for its parallelisation. On these early steps carried out by BSC specialists, it was detected that the code required some small modifications to accept command line arguments to choose some basic parameters like the resolution, the number of processes or the algorithm used. Those modifications were done by BSC in order to facilitate the testing and profile of the application.

In the following table is shown the first tests run with the code using their basic multiprocessing parallelisation strategy, it is important to remark that the code works as an embarrassingly parallel workload doing brute force to calculate all the possible combinations.

Processes	Time (seconds)	SpeedUp	Efficiency
1	226.56	1.00	100.00%
2	118.34	1.91	95.72%
4	59.32	3.82	95.48%
8	32.14	7.05	88.11%
16	16.61	13.64	85.27%
32	9.88	22.93	71.64%
48	8.26	27.44	57.17%

Table 10: Performance with multiprocessing

As we can see in Table 10, there is a loss of efficiency using multiprocessing. Additionally, the multiprocessing library only permits running on a single node. So, we decided to develop a version of the code using MPI with the mpi4py Python library, and to remove the multiprocessing sections. Also, the developers from the SME introduced a new computational part in the source code, named "target_engine" that only calculates some combinations depending on some input parameters. This new module was not parallelised and it is independent of the brute force part. The results from both modules are combined at the end of the simulation.

We developed a new version of the code using only mpi4py to parallelise the brute_force and the target_engine parts. The brute_force part can easily be parallelised by dividing the total number of combinations to be calculated into the MPI processes. In the case of target_engine, it can be parallelised by dividing each parameter of the input into the MPI processes. But, most of the time there will be more MPI processes than parameters to be calculated. Furthermore, there is imbalance between the time needed for each parameter calculation of the target_engine. So, we decided to implement a new argument for the code to specify how many MPI processes are dedicated to compute the target_engine part, and the rest of the processes run the brute_force.

We tried to reduce the imbalance by developing a new version of the code to run multiple instances of the code using GREASY (a workload management tool developed at BSC); on each instance it will run target_engine or brute_force depending on one argument. We saw that with this version we have almost the same imbalance and it is also more difficult for the SME to use it.

PRACE cooperation

BSC has done all the runs on MareNostrum, including the performance analysis with the profilers and the benchmarks to measure the performance of the newly developed versions.

Also, BSC has guided the SME to include the arguments options of the code and introducing MPI and its usage to the SME, as they were not familiar with it.

Benefits for SME

With the new parallelisation strategy, from multiprocessing library to mpi4py, the SME has enabled their code to be used on a Tier-1 or Tier-0 Call. Thanks to this evolution they can go from running on their server with only 8 processes to running on multiple nodes. Also a clear improvement in their single-node parallelisation has been achieved vs the previous version using only the python multiprocessing package.

In addition to being enabled for Tier-1 or Tier-0 Calls, the possibility of running the code with more resources will allow the SME to run larger inputs with a higher resolution which would take months of computation with the first version of their code.

Lessons learned

The Collaboration between BSC and the SME worked quite well. It is also important to remark that the code required to be parallelised was not a typical HPC code, as it was an embarrassingly parallel workflow using python. Also the company was using cloud environments to run, so the parallelisation was kept as simple as possible using mpi4py and GREASY.

3.5.10 SPARK inovacije, d.o.o.

Overview

Project partners:

- Martin Pečar, SPARK inovacije d.o.o.
- Gregor Mrak, SPARK inovacije d.o.o.
- Matjaž Šubelj, UL

SPARK inovacije is a startup with 25 people, working on logistics optimisation. They are working on route optimisation problems for logistics companies with 10+ vehicles, doing local distribution. They are investigating the possibilities of further savings beyond classic vehicle routing problems, employing multi-objective route optimisation / alternative routes. They are also a part of a consortium that won a pilot/demonstration project of next-generation Traffic Management Center for the city of Maribor, Slovenia. They will be responsible for improving the distribution of traffic flow in order to minimise traffic congestion. The basic idea is to use an alternative routes modules to generate multiple possibilities, and then select which ones are the best, in combination with others. The final goal is to achieve the Nash equilibrium.

Route optimisation is a process with high computational costs, so it is a perfect match for HPC. Higher computing power leads to better optimisation and less cost. 5-15% savings on distance can be achieved with relatively simple and fast approaches. However, the expectation is that higher numbers can be achieved by using more sophisticated techniques, and above all, more computing power. Apart from that, the next-generation Traffic Management Center could also benefit from HPC capabilities.

Activity performed
The software was installed on the HPC FS cluster and test cases were successfully run. Profiling and parallelisation efforts are ongoing.

PRACE cooperation
UL as PRACE partner provided access to HPC FS cluster and initial coaching. Further collaboration is expected with parallelisation efforts.

Benefits for SME
With the multi-objective optimisation approach, we can optimise all kinds of transport operations, mobility services and other logistic operations. In the field of distribution activities, we can typically decrease the travelled distance by 10+%. Distance is highly correlated with financial costs, CO2 and other emissions. Our solution can be used for improving centralized traffic management solutions, which will be very important in future use of autonomous vehicles. The plan is to guide individual vehicles so that the traffic flow will be distributed better, not only in traffic arteries, but also using possible detours where beneficial.

Lessons learned
The SME is very interested and responsive, but their developers are very busy with daily operations. Hence, we are still in the early stages.

3.5.11 Tarentum

Overview
Tarentum delivers AI powered products for businesses across multiple industries. The Wind Power Forecasting product provides production forecasts for Energy Companies. In this project with PRACE, Tarentum combines Numerical Weather Prediction (NWP) models with Artificial Intelligence/Machine Learning Algorithms to improve prediction accuracy and optimize NWP simulations. The team at UHEM in Turkey will work with Tarentum to develop their HPC plans.

Activity performed

- In the first part of the project, machines provided by UHEM were accessed.
- The NWP model (WRF v4.1) was installed on the machines.
- The software and libraries needed by the model were installed with the help of UHEM.
- The input data required by the model was provided to the machines.
- How the model performs with a different numbers of processors was tested.
- The model was run with the settings to be used in the project.

PRACE cooperation
PRACE was involved with HPC machine access.

Benefits for SME
- This project provided Tarentum access to high performance computer systems.
- Tarentum was able to run NWP models with different settings by using this computer power.
- Different outputs of models with different settings will provide more data opportunities to machine learning models. Machine learning models will minimize errors with more data.

Lessons learned
- We had the opportunity to work with an HPC center to adapt and run the operations of Tarentum.
- In addition to our previous experiences, we had the opportunity to work with different compilers and libraries. In this way, we decided to improve ourselves at different points that we had not experienced before.

3.5.12 VIPUN Medical N.V.

Overview
VIPUN Medical develops a novel method to feed critically ill and other vulnerable patients.

The VIPUN Gastric Monitoring System (GMS) aims to make it easy for medical staff to make a well-informed and faster nutrition therapy decision thus reducing malnutrition and feeding-related complications.

Activity performed
As of September 2020 this project has yet to begin.

Benefits for SME
SHAPE will help VIPUN to simulate various parameters in the models they have which is where HPC will come in.

3.5.13 Mobildev

Overview
Mobildev, founded in 2002, offers corporate bulk messaging and mobile marketing services. It offers data management, productivity, permission collection appropriate for KVKK (Turkey's regulation on privacy protection) and ETK, interactive SMS, bulk SMS and bulk e-mail services.

Activity performed
As of September 2020 this project has yet to begin.

Benefits for SME
This SME works in the data management sector, offering bulk messaging and mobile marketing services. They use machine learning (ML) techniques and wish to use a CPU or GPU system. The SHAPE project will be used to develop this work.

3.5.14 Global Maritime Services Ltd

Overview
SME Contact: Helong Wang - Project Analyst

PRACE Partner: EPCC

GMS was incorporated in 2008 and provides risk and cost reduction services for the high-value maritime and offshore domains. Its main focus is to develop services that improve safety and efficiency for navigation and offshore operations across the global maritime industry.

The aim of the SME is to develop products and services with significant added value to the shipping and offshore industries. Therefore, its research and development team has many years of operational shipping experience and a thorough understanding of the high-value maritime market domain.

Together with its partners, GMS has been developing a state-of-the-art navigation system for the international maritime industry, which will enable greater fuel optimisation and make the global maritime transport sector more sustainable, and also reduce emissions, limit structural fatigue, and improve passenger and crew comfort

GMS is developing a sail planning service system, VoyOpt, which provides optimal routes for shipping, reducing fuel consumption and greenhouse gas emission, and also avoiding heavy weather conditions which could cause ship / cargo damage or crew injury. The system comprises the weather forecast and in-situ/satellite observations data access and processing, naval architecture ship models algorithmic implementation, and optimal route calculation, all connected to each other in a flexible IT infrastructure through micro-services.

VoyOpt's prototype has achieved 5% fuel saving and emission reduction for the shipping companies which act as its project partners. GMS wants to further enhance the system by improving the ship performance modelling techniques, which are strongly dependent on computational techniques and resources, with the following goals:

- Ship data-driven update, development, and update of ship performance model;
- Validation of the ship performance model using computational fluid dynamics (CFD) approach.

Activity performed
As of September 2020 this project has yet to begin.

PRACE cooperation
The PRACE partner will help GSM with the more accurate computational modelling of ship performance, using OpenFOAM for CFD analysis, parallelised using MPI.

Benefits for SME
One of the project's targets is to accurately estimate ship performance through simulations as this is difficult to obtain in real operating conditions. Having simulations as efficient as possible is obviously important, but improving the simulation through more accurate real world conditions and use of more accurate numerical methods can allow for more accurate ship performance predictions.

These are vital in maritime related activities as it can allow to plan a better route optimised ship voyage, thus allowing for GMS to be more competitive by being able to provide its customers with a more accurate, efficient, and safe voyage planning.

3.5.15 Offshore Monitoring Ltd

Overview
SME Contact: Waqas Qazi - MetOcean Manager

PRACE Partner: CaSToRC, Cyprus

Offshore Monitoring Ltd. (OSM) is an SME founded in 2007 that aims to improve efficiency, environmental sustainability, safety, security, and navigation for commercial shipping and related offshore activities through the transference of non-maritime emerging technologies and by establishing new operational or business procedures, services, and solutions for relevant offshore monitoring technologies.

A unique characteristic of OSM is its focus on solving specific challenges where end-users/clients see potential to improve their operations in terms of safety, efficiency, security, and environmental compliance, working jointly with real clients and end-users to ensure the introduction of the exact product or service the market needs.

VoyOpt OSM's automated, all-weather sail planning service guides ships from port to port in real time in the most efficient way. VoyOpt connects various traditionally non-maritime technologies and integrates them in a unique service which delivers the optimum route, according to multiple user criteria. The system comprises the weather forecast and in-situ/satellite observations data access and processing in the Data Merging System (DMS), naval architecture ship models algorithmic implementation, and optimal route calculations, all connected together in a flexible IT infrastructure through micro-services.

While VoyOpt is already at an advanced stage of development, there are certain bottlenecks in the data handling and optimisation parts, which are goals of this project to solve in the SHAPE collaboration:

- Increased parallelization in the DMS (more than 200 processes) for interpolation of timesteps for the intermediate weather products;
- Multi-parameter route optimization (algorithm made to run exhaustively instead of with heuristics).

Activity performed
An initial informal meeting between CaSToRC and OSM has taken place once this project has been accepted as a SHAPE+ project, but the project and work towards it has not officially begun.

PRACE cooperation
CaSToRC will analyse and profile the software of OSM to identify bottlenecks, will parallelise code where possible and also seek to identify parts of the code which can be optimised.

Benefits for SME
The prototype VoyOpt service has already achieved on-average 5% fuel and emissions savings and further enhancements could increase this fuel efficiency. Furthermore, the SME would like to add

more parameters in its optimisation (beyond fuel efficiency) as well as targeting other sea specialised markets - such as short sea shipping and coastal shipping. The R&D from this project will allow the SME to:

- Increase product quality and service delivery through technical innovation;
- Reduce the time-to-market and increase market uptake through product development;
- Increase system capability for scale-up - Serving in real-time hundreds of ships sailing in different areas of the global oceans, while also collecting metocean data and ship performance information;
- Increase service innovation by high-quality route optimization in shorter timeframes, enabling delivery of route updates at shorter temporal intervals;
- All of these will lead in the future towards growth and expansion of the SME, as it acquires more clients and gets the VoyOpt service fully operational.

Lessons learned
This will be CaSToRC's first SHAPE project and further to the benefits for the Cypriot SME, it will also provide CaSToRC with experience in working and helping SMEs which will be beneficial towards future SHAPE+ applications, as well as forthcoming Industrial, SME and governmental engagement within the confines of the EuroHPC Cyprus National HPC Competence Center activities.

3.5.16 d:AI:mond GmbH

Overview
Partners:

- SME contact: Thilo Krueger
- PRACE contact: Anastasiia Shamakina (HLRS)

d:AI:mond GmbH is a services company from Germany. The company supports its customers in a wide range of Data Science projects. A focus is on data driven problem solving, where a team analyses data gathered by their customers and delivers tailored solutions.

The goal of the project is to optimise the production sequence in a company that produces general cargo.

Activity performed
As of September 2020 this project has yet to begin but will begin in November 2020.

PRACE cooperation
The PRACE partner will help d:ai:mond with accelerating the pilot project for a local general cargo company. The simulation is computationally more costly than expected and PRACE experts will implement running multiple simulations in parallel.

Benefits for SME
d:AI:mond company plans to expand their business by introducing a new, very concrete service, namely the simulation and the optimisation of production sequences by Reinforcement Learning. Their ultimate goal is to grow the company by installing such a product-like service. Once the

company proved that the optimisation strategy works, it wants to create a division for those optimisation projects. Therefore, d:AI:mond would hire new Data Science experts.

3.5.17 CENTIMFE

Overview
CENTIMFE, Technological Centre for Mouldmaking, Special Tooling and Plastics, is a Private Non-Profit Association, composed by companies (mostly SMEs), universities, local authorities, public agencies, etc. CENTIMFE is dedicated mostly to applied R&I+D, technology transfer and engineering services and training for Mouldmaking, Special Tooling and Plastics industries.

This SHAPE project will investigate how an HPC platform will improve the airflow generated by automotive heating, ventilation, and air conditioning (HVAC) systems using CFD simulations. The purpose of the simulation is to find out how to reduce the noise produced by this airflow. The improvement will be done by scaling the simulation up to 64 cores in the Navigator cluster of UC-LCA and is expected to be twofold:

1. to reduce substantially (at least an order of magnitude) the time of execution of each simulation;
2. to have finer mesh sizes and reduce the time steps in order to produce more accurate simulations, capable of reproducing current and future experimental results.

The software used is Ansys Meshing and Ansys Fluent.

The project team is composed of João Caseiro from CENTIMFE and from PRACE/UC-LCA Pedro Alberto and Miguel Oliveira. This will be UC-LCA's first SHAPE project and the first project from Portugal receiving an allocation of effort via the SHAPE+ programme.

Activity performed
The activity is just at the beginning. We are in the first stage of the working plan of the SHAPE proposal. The software is already installed in the Navigator cluster and the next steps will be running testing cases and benchmarks. These will be verified and profiled to identify bottlenecks.

Benefits for SME
- Product optimization in which computer simulation is being used to understand and minimize the airflow generated noise in automotive HVAC systems.
- Understand the challenges and advantages of parallelisation of large problems, endogenous knowledge and spread it to CENTIMFE SME community.
- Reduced lead time, due to minimization of a posteriori revisions, and, consequently, reduced manufacturing costs.

3.5.18 SmartCloudFarming GmbH

Overview
Partners:

- SME contacts: Michele Bandecchi, Prof. Dubravko Culibrk
- PRACE contacts: Anastasiia Shamakina (HLRS), Sameed Hayat (HLRS)

SmartCloudFarming started as a project of three university graduates in Berlin/Germany, in 2018. Their aim was to make use of earth observation to make agriculture productive, resilient, and sustainable.

Over the next twelve months many changes and challenges followed: The focus settled on soil health and productivity, and the initial project team of three was reduced to just one.

During this challenging period, Michele Bandecchi guided SmartCloudFarming forward by building a new team, generating visibility, and being accepted to the prestigious Copernicus Incubator. In 2019, SmartCloudFarming was selected by Indigo Ag as one of three European semi-finalists for their global Terraton Challenge. Today, SmartCloudFarming's new and growing team is supported by an international team of specialist collaborators and experienced mentors.

The SmartCloudFarming company is developing a machine learning model capable of accurately predicting Soil Organic Carbon by deploying deep learning models on a large amount of satellite data. Their target is to train and evaluate the applicability of VGG16, Resnet-50 and other architectures on hyperspectral satellite imagery with the goal of determining key parameters related to soil texture and soil organic carbon. They would start from pretrained hyperspectral models trained to classify land use and finetune them to estimate first soil texture, then soil organic carbon.

Activity performed
As of September 2020 this project has yet to begin but will begin in October 2020.

PRACE cooperation
The PRACE partner will help SmartCloudFarming with setting the deep learning environment suitable to the hardware provided; profiling data; providing strategies for code optimisation; profiling the solution to ensure maximum efficiency.

Benefits for SME
Having an optimal computational power will allow the SmartCloudFarming company to processed larger databases of soil samples, satellite images and multi-temporal dataset and enable them to establish solid causations that deliver accurate and reliable predictions as well as accelerate the release, test and validation the prototype version and Minimum Viable Product (MVP) version for pilot customers.

3.6 The twelfth SHAPE call

The twelfth SHAPE call opens on 1st October 2020 and closes on 1st December 2020. The call will follow the same procedures as for the eleventh call and SHAPE+ effort will be available for working with SMEs from countries new to SHAPE.

3.7 SHAPE: future

The results of the SHAPE+ initiative have been very encouraging, and the success of the last call both in attracting new countries and the overall number of proposals received shows that the SHAPE programme is presently strong. We intend to build on this success and are looking at ways of improving the experience for SME applicants. At recent PRACE IAC meetings, and the replacements to the face-to-face meetings, we discussed potential improvements to both the timescales and the proposal forms.

With regard to the timescales, it is felt that 6-monthly call deadlines can result in a long process for SMEs, especially if they just miss a deadline and have to wait several months for the next call, followed by awaiting a decision on the award and then set up of the project. As a result of this, we are looking at the possibility of having more regular call cut-offs, probably at the level of 3-4 per year. This will be discussed further at IAC and within PRACE-6IP.

With regard to the proposal form, the form is already fairly lightweight in nature and the technical section is optional. However, the presence of the technical section (even if optional) is likely to put some SMEs off. One possible improvement is to have a 2-stage process where the initial application form asks mainly for administrative details. This then starts a dialogue, where at a second stage we can gather more technical information, with help provided by the local PRACE centre. Again, this will be discussed over the coming months.

A major advantage of SHAPE is the unique Europe-wide network of national HPC-centres which represent the membership base of PRACE. The proximity of these centres to regional SMEs, often based on long personal relationships, cultural closeness and without linguistic barriers, is a priceless market entry for SHAPE. Strategies are currently formulated as how to better use these existing connections for the overall business development of SHAPE.

We are also looking at ways to provide a smoother transition between various PRACE programmes, encouraged by the success of seamlessly moving one SHAPE proposal from SHAPE to Preparatory Access for the most recent call. Looking outside of PRACE, we will continue to build links with the EuroHPC Competence Centres.

4 T7.3 DECI Management and Applications Porting

At the top of the pyramid of HPC services are the Tier-0 level resources available to researchers via PRACE via Project Access. The PRACE Distributed European Computing Initiative (DECI) programme provides access to resources at the next level down: the Tier-1 level. Underneath this sit Tier-2 systems operated at the regional level within each country.

In this section we focus on the projects from the DECI-15 and DECI-16 calls. An overview of the DECI is given in section 4.1 which includes the processes used for awarded resources. When the DECI programme began under DEISA applicants had the option to request enabling effort from DEISA staff to assist with their project. This sometimes consisted of small amounts of effort to install codes or minor porting assistance, but could sometimes be more significant work on restructuring the code or reworking the parallelisation. Such effort was not available during the years that DECI ran under the PRACE Optional Programme, but some enabling effort has been brought back into the project now that DECI is part of the Work Package 7. The effort here is only small and so any enabling work consists mainly of the installation of packages and giving advice to users about optimal ways to run their codes. Section 4.2 describes the present status including any enabling work carried out. Projects requiring more significant enabling effort are more appropriate for Preparatory Access Type D. Finally, in section 4.3 we describe the future of DECI including plans for the next call.

4.1 DECI Overview

The DECI-15 call opened on 15 January 2019 and closed on 28 February 2019 receiving 73 proposals, 40 of which were awarded resources. Projects ran from summer 2019 to summer 2020 with a few running later into 2020 where extensions were considered acceptable. The DECI-16 call opened on 16 December 2019 and closed on 31 January 2020 receiving 68 proposals, 48 of which were awarded resources. Projects from this call started in summer 2020 and will continue to run until summer 2021.

The process of DECI submission and award was described in [21] and has remained broadly the same since. A summary of the evaluation and assignment process is provided here for completeness:

1. **Assigning of home sites**: DECI defines "internal proposals" as all proposals where the Principal Investigator (PI) is from a country which both participates in and contributes resources to DECI calls. "home sites" are institutions of a country contributing resources to DECI and also dealing with the administrative side of applications. After a call closes, each internal proposal is assigned to its home site - with the country of a PI being the same as the home site. "External proposals", which are DECI proposals with a PI from a country which is not contributing resources to the DECI call, are also assigned to a home site from the available pool of DECI home sites.
2. **Administrative evaluation**: Once applications have a home site, all applications are checked for the consistency of their information and completeness of the applications. If any clarifications or extra details are required, home site representatives may contact the

PIs of applications. All applicants will receive confirmation that their proposals have been correctly submitted and that they meet the eligibility criteria.

3. **Technical evaluation**: Applications then undergo technical review by a member of DECI staff at the relevant home site to assess that the technical requirements of applications conform to DECI expectations. This means that applications should be scalable for the requested resources and that these resources will be managed and used in an efficient manner. The technical review should find any obstacles that would prevent the project succeeding but also find out if the required architectures suggested by the candidate are suitable. These reviews are carried out within the DECI submission tool.
4. **Scientific evaluation**: For internal proposals, these are usually carried out by a national panel of experts assembled from the same country as an application's PI. For external proposals, this is carried out by scientific reviewers from the Tier-0 Project Access pool of scientific reviewers provided by the PRACE aisbl. Occasionally, internal proposals are reviewed by the Tier-0 pool. This usually happens if a country is new to DECI or if there are very few applications from the country in question. National panels and the external reviewers all rank their respective subset of applications based on scientific excellence. Scientific reviews are also carried out within the DECI submission tool.
5. **Awarding (DAAC)**: Once all reviews are completed, the DAAC (DECI Access and Allocation Committee) meets to agree which proposals will be awarded resources with the key principle being scientific excellence. These decisions are made based on the rankings provided by scientific reviewers. Allocation of resources follows the juste retour principle where each country receives approximately the same amount of computing time the country has contributed, minus 15% which is dedicated to external projects.
6. **Allocation of resources**: Taking into consideration each application's technical requirements, projects are then assigned to the DECI Tier-1 machine where the project will run. Applicants are then informed of their awarded computational allocation and their assigned machine, and they are expected to confirm their acceptance of their award or inform us of any potential problem with the chosen machine.

DECI projects are awarded resources in units of DECI standard hours. This concept was introduced under DEISA to allow for the transfer of resources between sites where the conversion from core hours to DECI standard hours depends on the approximate expected performance of codes on the given cores.

For the most recent two calls, systems have included: Cray XC30/XC40, Intel clusters (various processor and memory configurations) and hybrid systems (clusters with GPGPU accelerators or Xeon Phi Co-processors (KNC))

4.1.1 DECI-15

Projects from DECI-15 were awarded under the PRACE Optional Programme, before PRACE-6IP started. However, they were running during PRACE-6IP so are reported here. Proposals were received from 26 different countries and awarded to proposals from 21 different countries. Figure 19 shows the number of proposals received from each country alongside the same data for awarded projects. Figure 20 shows the subject areas for received and awarded projects. As is typical for HPC projects, Materials and Bio Sciences are the most popular.

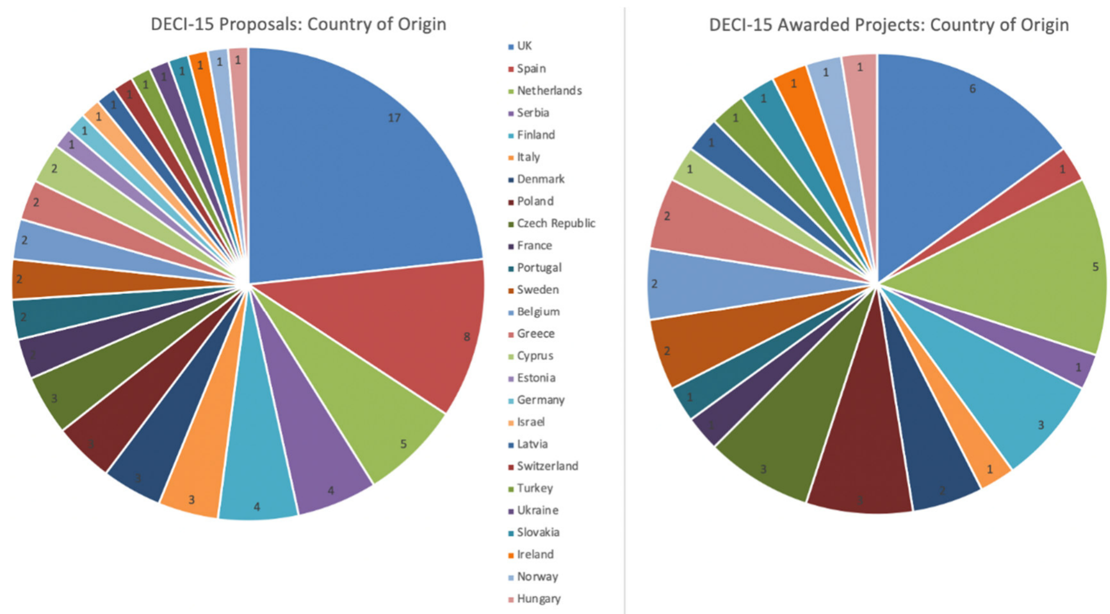

Figure 19: DECI-15 Proposals received and Projects awarded by country

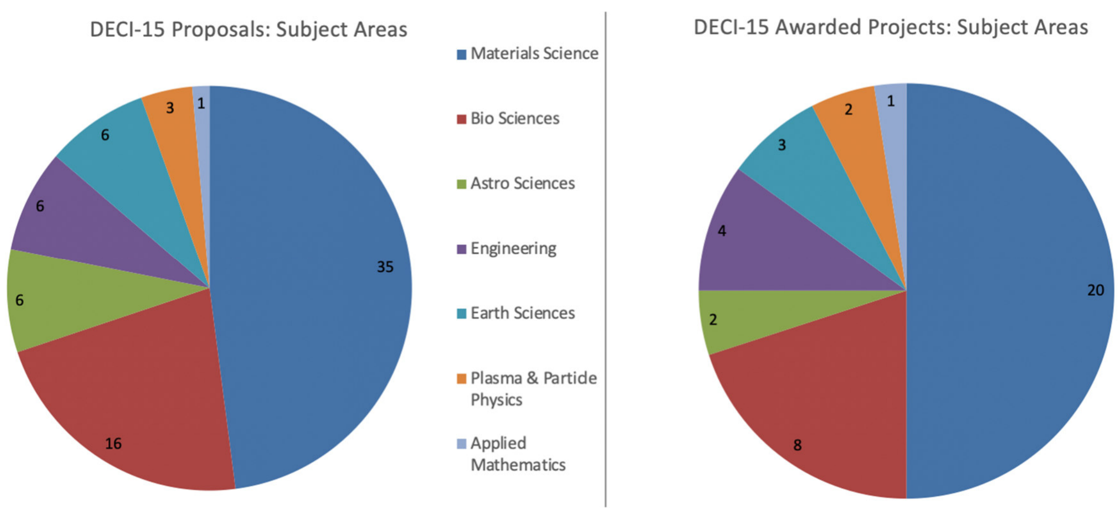

Figure 20: DECI-15 Proposals received and Projects awarded by subject area

4.1.2 DECI-16

Proposals were received from 21 different countries and awarded to proposals from 19 different countries. Figure 21 shows the number of proposals received from each country alongside the same data for awarded projects. Figure 22 shows the subject areas for received and awarded projects. As for DECI-15, Materials and Bio Sciences are the most popular.

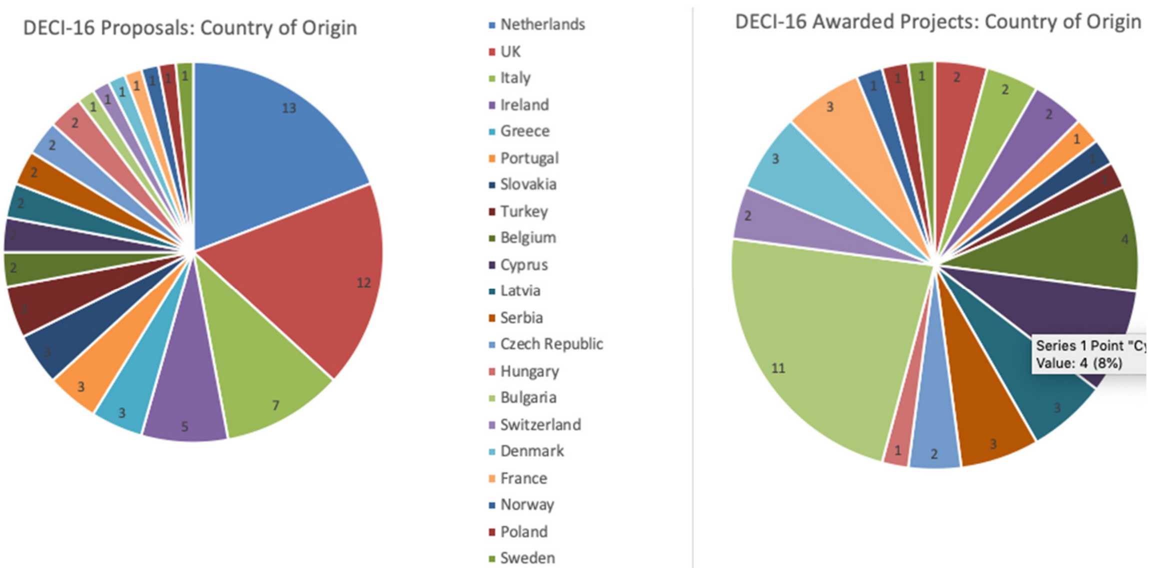

Figure 21: DECI-16 Proposals received and Projects awarded by country

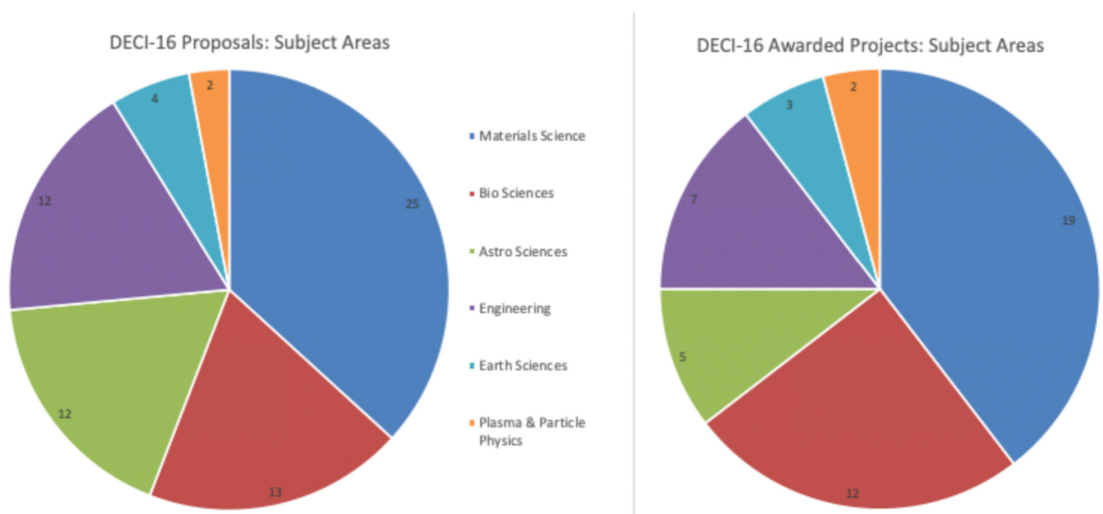

Figure 22: DECI-16 Proposals received and Projects awarded by subject area

4.2 DECI Programme Status and Project Enabling

4.2.1 DECI-15

Table 11 shows the entire list of DECI-15 projects awarded together with the country of origin of the proposal, the home site looking after the project, the subject area, and the amount of DECI standard hours awarded.

DECI-15 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
OFFSPRING	Belgium	SURFsara	Materials Science	4,016,250
pylight	Belgium	SURFsara	Materials Science	9,957,294

DECI-15 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
LHCPSSalt	Cyprus	CaSToRC	Bio Sciences	11,812,500
BOLERO	Czech Republic	VSB-TUO	Materials Science	9,100,000
PERMON2	Czech Republic	VSB-TUO	Applied Mathematics	6,000,000
MD4V2O5	Czech Republic	VSB-TUO	Materials Science	10,050,000
ANNEX	Denmark	CSC	Bio Sciences	14,400,000
MELMUT	Denmark	CSC	Bio Sciences	12,000,000
PSurf	Finland	CSC	Bio Sciences	20,000,000
QMCPIS	Finland	CSC	Materials Science	20,000,000
TFLL	Finland	CSC	Bio Sciences	20,000,000
CatDesign	France	UIO	Materials Science	8,000,000
AT1R	Greece	GRNET	Bio Sciences	2,100,000
ELSTELL	Greece	GRNET	Plasma & Particle Physics	8,750,000
LocalCorr	Hungary	KIFU	Bio Sciences	7,400,000
MetalAgg	Ireland	ICHEC	Materials Science	7,337,472
ToQTS	Italy	ICHEC	Materials Science	5,400,000
WS2MoS2AIMD	Latvia	KTH	Materials Science	3,564,000
Flamingo	Netherlands	SURFsara	Astro Sciences	20,000,000
TIO2HDS	Netherlands	SURFsara	Materials Science	2,380,000
RAMPA	Netherlands	SURFsara	Engineering	10,000,000
DDD	Netherlands	SURFsara	Engineering	19,867,200
HIPERGLOFCC	Netherlands	SURFsara	Earth Sciences	3,600,000
BHAO2	Norway	UIO	Earth Sciences	10,400,600
NEUROENZ	Poland	PSNC	Bio Sciences	38,325,000
molfragprop	Poland	PSNC	Materials Science	25,200,000
MuMo4PEC	Poland	WCSS	Materials Science	20,000,000
COIMBRALATT5	Portugal	CESGA	Plasma & Particle Physics	7,560,000
SACandM	Serbia	GRNET	Materials Science	2,400,000
TITANIA	Slovakia	VSB-TUO	Materials Science	23,920,000
4lessCH4	Spain	BSC	Materials Science	7,956,480

DECI-15 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
AMCVD	Sweden	KTH	Materials Science	31,200,000
DYNAMAT	Sweden	KTH	Materials Science	31,320,000
COSHOIP	Turkey	EPCC	Materials Science	6,000,000
PACWBTF	UK	EPCC	Engineering	13,771,296
DOMINOS	UK	EPCC	Earth Sciences	2,800,000
subgridEoRevol	UK	EPCC	Astro Sciences	21,000,000
PINNTFlows	UK	EPCC	Engineering	10,077,804
QTProp	UK	EPCC	Materials Science	4,375,000
SimMolManip	UK	EPCC	Materials Science	5,809,344
			Total	497,850,240

Table 11: Full list of DECI-15 projects

Enabling work

For DECI-15, which started before PRACE-6IP, enabling effort consisted of setting up projects and getting users access to the various systems. Projects were not selected if there was any further enabling work required.

4.2.2 DECI-16

Table 12 shows the entire list of DECI-16 projects awarded together with the country of origin of the proposal, the home site looking after the project, the subject area, and the amount of DECI standard hours awarded.

DECI-16 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
RemEPI	Belgium	SURFsara	Materials Science	20,230,720
Mag1	Belgium	SURFsara	Materials Science	16,912,000
DynLHCX	Cyprus	CaSToRC	Bio Sciences	19,250,000
aiGRALIB	Cyprus	CaSToRC	Materials Science	15,720,000
CABLE	Czech Republic	VSB-TUO	Materials Science	3,657,500
TriboVan	Czech Republic	VSB-TUO	Materials Science	7,700,000
LC3bind	Denmark	CSC	Bio Sciences	6,709,248
PeMaD	France	KTH	Materials Science	13,447,680
ELSTELL2	Greece	GRNET	Plasma & Particle Physics	5,267,000

DECI-16 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
LOXIRB	Greece	GRNET	Bio Sciences	2,400,000
KrasMem	Greece	GRNET	Bio Sciences	4,800,000
SomatoDock	Hungary	KIFU	Bio Sciences	5,598,720
2DQubits	Hungary	KIFU	Materials Science	10,400,000
BlackOut	Ireland	ICHEC	Astro Sciences	17,680,000
MetalDep	Ireland	ICHEC	Materials Science	10,463,232
HIGHWAVE	Ireland	ICHEC	Earth Sciences	15,600,000
CO2MET	Ireland	ICHEC	Materials Science	19,336,200
DyMoBoNT	Italy	CINECA	Bio Sciences	4,680,000
RADOBI	Italy	CINECA	Engineering	4,160,000
HIRESBHAC	Netherlands	SURFsara	Astro Sciences	20,000,000
Beehive	Netherlands	SURFsara	Astro Sciences	20,000,000
FunIe3C	Netherlands	SURFsara	Materials Science	18,000,000
HISIGEM	Netherlands	SURFsara	Engineering	8,000,000
BHAO3	Norway	Sigma2	Earth Sciences	8,800,000
COMPENZ	Poland	PSNC	Bio Sciences	31,500,000
UCORecyclingGP	Portugal	UC-LCA	Materials Science	2,000,000
COIMBRALATT6	Portugal	UC-LCA	Plasma & Particle Physics	4,000,000
JaFerMat	Portugal	UC-LCA	Materials Science	2,000,000
SimSIT	Serbia	ICHEC	Bio Sciences	1,722,240
THEOCHAL	Slovakia	CCSAS	Materials Science	6,760,000
NANOJOINING	Slovakia	CCSAS	Materials Science	2,600,000
FluxMol	Slovakia	CCSAS	Bio Sciences	3,380,000
HIFIPIPE	Sweden	KTH	Engineering	36,000,000

DECI-16 Project Acronym	Country of PI	DECI Home site	Subject Area	DECI standard Hours Awarded
GCOUP	Switzerland	KTH	Bio Sciences	28,003,334
2DCSP	Turkey	UHeM	Materials Science	36,400,000
CCAN	Turkey	UHeM	Bio Sciences	1,344,000
HEMLAB	Turkey	UHeM	Engineering	8,000,000
WOBBLYDUST	UK	EPCC	Astro Sciences	17,500,000
DOMINOS2	UK	EPCC	Earth Sciences	11,904,000
WETONB	UK	EPCC	Engineering	16,640,000
GDFoEGF	UK	EPCC	Astro Sciences	3,675,000
DSMCsimulations	UK	EPCC	Engineering	2,111,760
SDAnOx	UK	EPCC	Materials Science	16,611,840
2TRPM8	UK	EPCC	Bio Sciences	7,000,000
ADoBE	UK	EPCC	Materials Science	12,250,000
LESHTlowPrt	UK	EPCC	Engineering	7,000,000
FEINO	UK	EPCC	Materials Science	8,400,000
MATROMUGR	UK	EPCC	Materials Science	14,000,000
			Total	559,614,472

Table 12: Full list of DECI-16 projects

Enabling work

For DECI-16 some enabling effort was available, although the amount of effort was small. This did enable sites to install newer versions of software (e.g. a newer version of Quantum ESPRESSO installed by UHeM), help with more complicated licensing arrangements (e.g. PDC helping users with VASP) and some help in running packages optimally (e.g. CSC helping to get GROMACS and LAMMPS running optimally). Other help was given to the use of the queue (tools such as SLURM) as well as a large number of account set ups for all of the above projects which typically have 1-4 users per project.

4.3 DECI Future

For DECI-16 work was carried out on the DECI submission form provided by CINES. This included restructuring the sections to make it easier to input information about predicted job runs. For the DECI-17 call, we plan to improve this further to better handle projects with multiple codes.

On the longer term, the database we use - the DECI Project Management Database (DPMDB) – has become outdated and difficult to maintain or update and is not easy to use. Over the next year decisions will be taken on its future, specifically whether to replace it with a similar tool or to integrate completely with the submission tool which already captures a lot of the data needed. Decisions on this will be taken in conjunction with any decisions about the submission tool used for PRACE Project Access and the DECI team will liaise with the Work Package 2 team and the aisbl here.

The DECI-17 call is likely to open on a similar timescale but one year later than the DECI-16 call (e.g. opening Dec 2020 closing Jan 2021). Over the next few weeks an audit of available DECI-17 resources will be collected in order to make a final decision on timescale on the opening of the call.

5 T7.5 Enhancing the High-Level Support Teams

T7.5 is a new task implemented in PRACE-6IP. It will provide additional effort to enhance the work of the PRACE High-Level Support Teams (HLSTs). The HLSTs will enhance the scientific output of the Tier-0 systems through the provision of level 3 (midterm activities) support activities. Task 7.5 supplements those activities and extends this work with specific expertise from the other PRACE centres. This ensures sharing of expertise across PRACE to maximise the benefits to users of the PRACE systems. Task 7.5 will work with the HLSTs to extend and enhance their activities supporting Tier-0 users, and also Tier-1 intensive users targeting Tier-0, in order to maximise the scientific output of the Tier-0 systems.

As a new activity within this work package, the first steps are to clearly define the scope of this task. As the task definition was to offer level 3 support to "large" users of the PRACE platforms, we agreed on the fact that specific acceptance criteria, as well as a very well defined project process are the key to enhance the global throughput of the systems.

Once the basis was set, the first projects were able to start.

The main concern for this activity, and to make sure that there is no overlap with other activities (such as Preparatory Accesses), was to make sure that we provide a level 3 support such as:

- HPC Co-development (specific libraries, new architectures, …)
- Enabling workflows (HTC requirements, dataflow management, …)

5.1 Project acceptance criteria

The criteria for project acceptance are the following:

- Scalability: targeting a scalability that only Tier-0 can provide;
- Architecture: HPC hardware targeted not accessible on Tier-1 platforms;
- CPU hours volume: exceeding what could reasonably by provided on a Tier-1;
- Memory/storage requirements: exceeding what could reasonably be provide on a Tier-1.

On top of that, a maximum of 2 projects per partner will be accepted (as only 10 PMs are available per partner).

5.2 Project process

A clear process was agreed with all partners within this task during the Face-to-Face meeting in Antwerp (October 2019).

This process includes the following points:

- Partners contact users that are using a huge fraction of their available system to offer them level 3 support
- Users answer the partners with specific project and technical issues to be addressed
- A monthly teleconference
- Project selection during the teleconference based on the criteria

- If the project is accepted, then:
  - The project starts on local Tier-1 system
  - Apply for a Preparatory Access on targeted Tier-0 system
  - Move to Tier-0 and validate the developments
  - Collaborate with local HLST to validate the developments
  - Activity reporting

5.3 Projects and plans for each partner

5.3.1 EPCC

Within the T7.5 in PRACE-6IP, EPCC aims to support UK/European users to work on the new UK Tier-1 system ARCHER2 and beyond.

The UK national supercomputing service will move to the next generation soon. The new UK Tier-1 system, ARCHER2, will be an HPE Cray Shasta system with an estimated peak performance of 28 PFLOP/s. The full machine will have 5,848 compute nodes, each with dual AMD EPYC Zen2 (Rome) 64 core CPUs at 2.2GHz, giving 748,544 cores in total and 1.57 PB of total system memory. ARCHER2 should be capable, on average, of over eleven times the science throughput of its predecessor, ARCHER. This estimate is based on benchmarks of five of the most heavily used research software packages on ARCHER, including CP2k, OpenSBLI, CASTEP, GROMACS and HadGEM3 climate model. Further hardware and software details on ARCHER2 can be found here: [22].

Due to the COVID-19 pandemic, the start date of the new ARCHER2 service had to be postponed and therefore we have not been able to start the project's technical work at EPCC in T7.5 by the time of preparing this deliverable. The ARCHER2 service is due to commence operation in late 2020 and currently the preparation work is ongoing on the first phase of ARHCER2, which is a 4-cabinet pilot system. Once the 4-cabinet ARCHER2 pilot system is ready for access, we plan to start the project at EPCC as soon as possible. There will be many useful topics/applications to work on based on the initial experience and analysis on ARCHER2. However, our main purpose for this project is to provide in-depth support to the UK/European users for their better transition and usage of the new UK Tier-1 service. Depending on the users' requirements, we will be able to provide in-depth support such as porting and optimising the widely used scientific application(s) on ARCHER2. The efforts made on this system will directly benefit to PRACE Tier-0 systems such as Joliot-Curie (CEA-TGCC, France), using the same hardware components (see: [23]), and enable large cpu.hours consumers to move directly to Tier-0 systems.

5.3.2 CaSToRC

CaSToRC started a project on Speeding up Full Waveform Migration in Geo-Imaging within the Task 5 of WP7 which will focus on Geo Imaging using Full Wave-Form Migration in collaboration with the project partner Prof. Eric Verschuur (University of Delft, Delphi Research Consortium). The project will focus on the optimization of the application code of the Delphi Research Consortium which uses full waveform migration with phase-shift plus interpolation.

The project was proposed at the virtual WP7 T7.5 face-2-face meeting on May 28th 2019, and discussed in the following teleconferences involving the Task members.

The project targets to enable the Geo-Imaging software to utilize larger HPC resources in a faster and more efficient way. Currently a typical seismic survey takes around 1.25 Mio Corehours using a dataset of 2.5 TeraByte.

The internal work on the project started within M12 of PRACE 6IP with the performance profiling of the software using MPI-timings or/and Score-P on local resources and JUWELS. The profiling results show that the scalability is hindered by the communication overhead (blocking MPI functions), and strong scaling is breaking down if 32 Cascadelake nodes are exceeded. Note that this is the size of our upcoming local machine and pushing the strong scaling regime towards larger systems will justify the utilization of larger HPC systems, such as PRACE Tier-0 systems. Moreover to reach higher resolutions by an efficient use of larger computational resources, there is a need to extend the scalability towards 100 nodes and more. The optimization of the communication patterns is currently ongoing, namely by switching from point-to-point to global communication routines and by enabling non-blocking communication functions. The current plan is to have results on that work by M18 which will be shared in the monthly teleconference of WP7 T7.5.

In a second step the project will focus on optimization of the compute kernel. Multiple directions are considered: adding flexibility to the dominant computational kernels by offloading FFT or using FFT-libs, or adding multi-threading or restructuring of memory allocations.

The results of this project will be disseminated within the Delphi Research Consortium. Results and improvements of the project will be regularly reported in the monthly teleconference. A final report will be included in the final Deliverable which will include a possible white paper reporting on the progress made within the project.

5.3.3 IT4I

We are discussing a project with the head of Material Design group at IT4I – Dominik Legut. This group has the highest CPU hours consumption at our infrastructure. They would like to run molecular dynamics simulations with LAMMPS framework with reactive force field (ReaxFF package) on a GPU. There are already existing implementations of the accelerated versions of this code available: the KOKKOS package of LAMMPS framework and PuReMD application but the group has no experience with them.

The goal of this project would be to set up the environment for these two GPU accelerated applications. After that we will run some small simulations and compare the speedup to the CPU version. Subsequently we will do some more performance measurements to see if there are any bottlenecks. We will also determine the size of the simulation when the scalability drops and whether these applications can utilize multiple GPUs or even multiple GPU nodes.

This project will probably be a form of capability computing –a very large number of smaller jobs run on the GPU accelerated nodes. The discussion with the Material Design group is still ongoing and there are some details which need to be sorted out. For example, what is the purpose of these simulations and what research will it help and what is the intended scale of the simulation. We also need to pick a targeted Tier-0 machine with GPU nodes which will be used for the final benchmarking.

5.3.4 SNIC-UU

High Performance Computing Center North (HPC2N), which is part of SNIC, has contributed to T7.5 with the Preparatory Access Type B project 2010PA5221 (3 months). The project was entitled "Monitoring proteins conformational changes through multiscale simulations". We work in collaboration with a local PI, Dr. Kwangho Nam, who is a core developer of the Chemistry at HARvard Macromolecular Mechanics (CHARMM) package.

CHARMM is the "workhorse" for quantum chemistry calculations as it includes the most popular methodologies. A shortcoming of the standard CHARMM version is its speed w.r.t. other molecular mechanics software. In order to achieve a better performance, the CHARMM-DOMDEC was released which provides GPU support. In the present project, we aimed at porting and monitoring the performance of CHARMM-DOMDEC on the Piz Daint Tier-0 machine at CSCS.

The purpose of using CHARMM-DOMDEC was for running multi-scale String Method simulations where multiple copies of the system are run in parallel. For the present project, we studied two systems called CASE1 and CASE2 with 49,000 and 23,558 numbers of atoms respectively. We tested CHARMM-DOMDEC with up to 106 nodes (106 copies) and it displayed a smooth weak scaling behaviour. The timings for the two cases run or Piz Daint where similar to those obtained on our local cluster for a single copy of the system (image). However, our local user cannot run there more than a few images as the number of GPU nodes is much smaller than on Piz Daint. Thus, our local users could benefit from using it. A report with the results of this project was written and sent to PRACE.

The progress on this line was interrupted by the pandemics as the local user couldn't apply for a larger PRACE project for data production in a more realistic system such as enzymes. The local user was also interested in using a different MPI protocol specially for the GPUs.

6 Summary

Four parallel tasks on application enabling and porting services in Work Package 7 of PRACE-6IP have been described including reports on the supported applications and activities. These two activities have been organised into support projects formed on a basis of either scaling and optimisation support for Preparatory Access, DECI, SHAPE and the HLSTs.

6.1 Applications Enabling Services for Preparatory Access

During PRACE-6IP Task 7.1 successfully performed six cut-offs for preparatory access including the associated review process and support for the approved projects. In total sixteen Preparatory Access projects have been supported or are currently supported by T7.1 . Eight projects were finally reported within this deliverable.

The projects try to produce white papers by the end of the project phase. Approved white papers are published online on the PRACE RI web page [11]. Table 13 gives an overview of the status of the white papers for all finalized projects.

Project ID	White paper	White paper status
2010PA3775	**WP284**: Scalable Delft3D Flexible Mesh for efficient modeling of shallow water and transport processes	Published online [14]
2010PA4347	**WP293**: Performance Analysis of the BiqBin solver for the maximum stable set problem	Published online [16]
2010PA4528	**WP288**: Optimisation of the Higher-Order Finite-Volume Unstructured Code Enhancement for Compressible Turbulent Flows	Published online [17]
2010PA4548	**WP294**: PETSc4FOAM: A Library to plug-in PETSc into the OpenFOAM Framework	Published online [18]
2010PA4547		No white paper produced
2010PA4774		No white paper produced
2010PA4821	**WP295**: Keeping computational performance analysis simple: an evaluation of the NEMO BENCH test.	Published online [20]
2010PA5236		White paper planned

Table 13: White paper status of the finalized PRACE-6IP PA projects

6.2 Applications Enabling Services for Industry

The SHAPE programme has continued to assist SMEs in trying out HPC and feedback from SMEs remains positive. We are very encouraged by the results of the SHAPE+ programme in attracting new countries to SHAPE as well as seeing an increase in the number of proposals overall, partly due to the efforts of Wahid Rofagha, the recently appointed PRACE Industry Liaison Officer. High quality White Papers continue to be produced from completed projects and a set of Success Stories can be found on the PRACE website. The twelfth call is about to open and plans are underway to streamline the process for SMEs to apply. We look forward to working alongside and sharing experiences with the EuroHPC Competence Centres.

6.3 DECI Management and Applications Porting

The DECI programme continues to be very popular with researchers around Europe with an almost constant number of around seventy proposals received each year. Improvements will continue to be made to the application process and discussions are underway regarding the smoothing of the process for gaining access to resources once projects are awarded. The seventeenth DECI call is expected to open in late 2020.

6.4 Enhancing the High-Level Support Teams

During the first part of the project, the T7.5 set up the processes that will be necessary to achieve the targeted goals. The following table is an overview of the task definition:

Activities	WP7 Task 7.5
Support teams	• Non-hosting members • 1 team per partner
Targeted users	• Any user aiming for PRACE (see criteria in 5.1)
Targeted machines	• National (Tier-1) machines • Then PRACE Tier-0 machines

Activities	• Contact “large” users
	• Optimisation (lvl3, 2-6PMs)
	  o co-development (specific libraries, new architectures, …)
	  o code refactoring, workflows
	• Max 2 projects per partner

Goals	• Improving global efficiency of Tier-0 machines • Collaboration HLSTs-WP7.5 to validate developments on Tier-0 machines

Table 14: T7.5 task definition overview

Currently two projects are ongoing within this task.

However with upcoming systems within Tier-1 partners (10’s Pflops systems), we all agreed within the task that it would make more sense to increase the effort within the second half of PRACE-6IP to make sure that we are preparing the users for pre-exascale architectures.