E-Infrastructures
H2020-EINFRA-2014-2015

EINFRA-4-2014: Pan-European High Performance Computing Infrastructure and Services

PRACE-4IP

PRACE Fourth Implementation Phase Project

Grant Agreement Number: EINFRA-653838

D7.2
Final Report on Applications Enabling

Final

Version: 1.0
Author(s): Paul Graham, EPCC; Sebastian Lührs, JUELICH
Date: 12.04.2017

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: EINFRA-653838**	
	Project Title: PRACE Fourth Implementation Phase Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D7.2	
	Deliverable Nature: Report	
	Dissemination Level: PU*	**Contractual Date of Delivery:** 30 / April / 2017
		Actual Date of Delivery: DD / Month / YYYY
	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 2991/844/EC.

Document Control Sheet

Document	**Title: Final Report on Applications Enabling**	
	ID: D7.2	
	Version: 1.0	**Status:** ***Final***
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2010	
	File(s): D7.2.docx	
Authorship	**Written by:**	Paul Graham, EPCC; Sebastian Lührs, JUELICH
	Contributors:	Chandan Basu, LiU Thomas Ponweiser, RISCSW Gabriel Hautreux, GENCI Tristan Cabel, CINES Stéphane Lanteri, INRIA Dimitris Dellis, GRNET Volker Weinberg, LRZ Bertrand Cirou, CINES Juan Carlos Cajas, BSC Andrew Emerson, CINECA John Donners, SURFsara Ahmet Duran, ITU Yakup Hundur, ITU Jörg Hertzer, HLRS Bärbel Große-Wöhrmann, HLRS Thomas Boenisch, HLRS Monica De Mier Torrecilla, BSC Jussi Heikonen, CSC Esko Järvinen, CSC Claudio Arlandini, CINECA Raffaele Ponzini, CINECA Vittorio Ruggiero, CINECA

		Isabelle Dupays, IDRIS Lola Falletti, IDRIS Philippe Segers, GENCI Thomas Palychata, GENCI Mariano Vázquez, BSC Beatriz Eguzkitza, BSC Abel Gargallo-Peiró, BSC Ricard Borrell, BSC Michaela Barth, KTH
	Reviewed by:	Nevena Ilieva-Litova, NCSA-BG; Thomas Eickermann, JUELICH
	Approved by:	MB/TB

Document Status Sheet

Version	**Date**	**Status**	**Comments**
0.1	23/March/2017	Draft	Set up document structure
0.2	26/March/2017	Draft	PA and SHAPE overview section were added
0.3	27/March/2017	Draft	All parts included
0.4	30/March/2017	Draft	Correct typos and grammar
0.5	31/March/2017	Draft	Version for internal review
1.0	12/April/2017	Final	

Document Keywords

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

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

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

[2] PRACE-4IP Deliverable 7.1 "Periodic Report on Applications Enabling", April 2016, http://www.prace-ri.eu/IMG/pdf/D7.1_4ip.pdf

[3] PRACE-3IP Deliverable 5.2 "Integrated HPC Access Programme for SMEs", February 2013, http://www.prace-ri.eu/IMG/pdf/d5.2.pdf

[4] PRACE-3IP Deliverable 5.3.1 "PRACE Integrated Access Programme Launch", June 2013, http://www.prace-ri.eu/IMG/pdf/d5.3.1.pdf

[5] PRACE-3IP Deliverable 5.3.2 "Results of the Integrated Access Programme Pilots", June 2014, http://www.prace-ri.eu/IMG/pdf/d5.3.2_3ip.pdf

[6] PRACE-3IP Deliverable 5.3.3 "Report on the SHAPE Implementation", January 2015, http://www.prace-ri.eu/IMG/pdf/D5.3.3_3ip.pdf

[7] Milano, G.; Kawakatsu, T. Hybrid Particle-Field Molecular Dynamics Simulations for Dense Polymer Systems. *J. Chem. Phys.* 2009, 130 (21), 214106.

[8] Zhao, Y.; De Nicola, A.; Kawakatsu, T.; Milano, G. Hybrid Particle-Field Molecular Dynamics Simulations: Parallelization and Benchmarks. *Journal of Computational Chemistry* 2012, 33 (8), 868–880.

[9] De Nicola, A.; Kawakatsu, T.; Rosano, C.; Celino, M.; Rocco, M.; Milano, G. Self-Assembly of Triton X-100 in Water Solutions: a Multiscale Simulation Study Linking Mesoscale to Atomistic Models. *J. Chem. Theory Comput.* 2015, 11 (10), 4959–4971.

[10] De Nicola, A.; Zhao, Y.; Kawakatsu, T.; Roccatano, D.; Milano, G. Validation of a Hybrid MD-SCF Coarse-Grained Model for DPPC in Non-Lamellar Phases. *Theoretical Chemistry Accounts* 2012, 131 (3), 1167.

[11] PRACE Whitepaper 221, "Parallel I/O Implementation in Hybrid Molecular Dynamics-Self Consistent Field OCCAM Code using fpioLib Library", http://www.prace-ri.eu/IMG/pdf/WP221.pdf

[12] PRACE Whitepaper 222, "Optimising UCNS3D, a High-Order finite-Volume WENO Scheme Code for arbitrary unstructured Meshes", http://www.prace-ri.eu/IMG/pdf/WP222.pdf

[13] http://glaros.dtc.umn.edu/gkhome/metis/metis/overview

[14] http://www.labri.fr/perso/pelegrin/scotch/

[15] http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview

[16] PRACE Whitepaper 223, "Enabling Space Filling Curves parallel mesh partitioning in Alya", http://www.prace-ri.eu/IMG/pdf/WP223.pdf

[17] http://www.ansys.com/products

[18] PRACE Whitepaper 224, “Parallel curved mesh Subdivision for flow Simulation on curved Topographies”, http://www.prace-ri.eu/IMG/pdf/WP224.pdf

[19] PRACE Whitepaper 225, “Particle Transport in a Fluid interacting with an immersed Body with Alya”, http://www.prace-ri.eu/IMG/pdf/WP225.pdf

[20] http://www.prace-ri.eu/white-papers/

[21] https://gitlab.INRIA.fr/solverstack/pastix

[22] https://openfoamwiki.net/index.php/SnappyHexMesh

[23] http://safecast.org/tilemap/

[24] http://blog.safecast.org/history/

[25] https://www.fortissimo-project.eu/

[26] PRACE Whitepaper 211, “Numerical Simulation of complex turbulent Flows with Discontinuous Galerkin Method”, http://www.prace-ri.eu/IMG/pdf/wp211.pdf

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
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.
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 upcoming next 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} ~ 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 UYBHM)
BSC	Barcelona Supercomputing Center - Centro Nacional de Supercomputacion, Spain
CaSToRC	Computation-based Science and Technology Research Center, Cyprus
CCSAS	Computing Centre of the Slovak Academy of Sciences, Slovakia
CEA	Commissariat à l'Energie Atomique et aux Energies Alternatives, France (3 rd Party to GENCI)
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 PSNC)
EPCC	EPCC at The University of Edinburgh, UK
ETHZurich (CSCS)	Eidgenössische Technische Hochschule Zürich – CSCS, Switzerland
FIS	FACULTY OF INFORMATION STUDIES, Slovenia (3rd Party to ULFME)
GCS	Gauss Centre for Supercomputing e.V.
GENCI	Grand Equipement National de Calcul Intensiv, France
GRNET	Greek Research and Technology Network, Greece
INRIA	Institut National de Recherche en Informatique et Automatique, France (3 rd Party to GENCI)
IST	Instituto Superior Técnico, Portugal (3rd Party to UC-LCA)
IUCC	INTER UNIVERSITY COMPUTATION CENTRE, Israel
JKU	Institut fuer Graphische und Parallele Datenverarbeitung der Johannes Kepler Universitaet Linz, Austria
JUELICH	Forschungszentrum Juelich GmbH, Germany
KTH	Royal Institute of Technology, Sweden (3 rd Party to SNIC)
LiU	Linkoping University, Sweden (3 rd Party to SNIC)
NCSA	NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS, Bulgaria
NIIF	National Information Infrastructure Development Institute, Hungary
NTNU	The Norwegian University of Science and Technology, Norway (3rd Party to SIGMA)
NUI-Galway	National University of Ireland Galway, Ireland
PRACE	Partnership for Advanced Computing in Europe aisbl, Belgium
PSNC	Poznan Supercomputing and Networking Center, Poland
RISCSW	RISC Software GmbH
RZG	Max Planck Gesellschaft zur Förderung der Wissenschaften e.V., Germany (3 rd Party to GCS)

SIGMA2	UNINETT Sigma2 AS, Norway
SNIC	Swedish National Infrastructure for Computing (within the Swedish Science Council), Sweden
STFC	Science and Technology Facilities Council, UK (3rd Party to EPSRC)
SURFsara	Dutch national high-performance computing and e-Science support center, part of the SURF cooperative, Netherlands
UC-LCA	Universidade de Coimbra, Labotatório de Computação Avançada, Portugal
UCPH	Københavns Universitet, Denmark
UHEM	Istanbul Technical University, Ayazaga Campus, Turkey
UiO	University of Oslo, Norway (3rd Party to SIGMA)
ULFME	UNIVERZA V LJUBLJANI, Slovenia
UmU	Umea University, Sweden (3 rd Party to SNIC)
UnivEvora	Universidade de Évora, Portugal (3rd Party to UC-LCA)
UPC	Universitat Politècnica de Catalunya, Spain (3rd Party to BSC)
UPM/CeSViMa	Madrid Supercomputing and Visualization Center, Spain (3rd Party to BSC)
USTUTT-HLRS	Universitaet Stuttgart – HLRS, Germany (3rd Party to GCS)
VSB-TUO	VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA, Czech Republic
WCNS	Politechnika Wroclawska, Poland (3rd party to PSNC)
SNIC	Swedish National Infrastructure for Computing (within the Swedish Science Council), Sweden
STFC	Science and Technology Facilities Council, UK (3rd Party to EPSRC)
SURFsara	Dutch national high-performance computing and e-Science support center, part of the SURF cooperative, Netherlands
UC-LCA	Universidade de Coimbra, Labotatório de Computação Avançada, Portugal
UCPH	Københavns Universitet, Denmark
UHEM	Istanbul Technical University, Ayazaga Campus, Turkey
UiO	University of Oslo, Norway (3rd Party to SIGMA)
ULFME	UNIVERZA V LJUBLJANI, Slovenia
UmU	Umea University, Sweden (3 rd Party to SNIC)
UnivEvora	Universidade de Évora, Portugal (3rd Party to UC-LCA)
UPC	Universitat Politècnica de Catalunya, Spain (3rd Party to BSC)
UPM/CeSViMa	Madrid Supercomputing and Visualization Center, Spain (3rd Party to BSC)
USTUTT-HLRS	Universitaet Stuttgart – HLRS, Germany (3rd Party to GCS)
VSB-TUO	VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA, Czech Republic
WCNS	Politechnika Wroclawska, Poland (3rd party to PNSC)

Executive Summary

Task T7.1 “Enabling Applications Codes for PRACE Systems” in Work Package 7 (WP7) of PRACE-4IP aims to provide application enabling support for the HPC applications which are important for the European researchers and small and medium enterprises to ensure the applications can effectively exploit HPC systems. There were two activities in T7.1:

T7.1.A Petascaling & Optimisation Support for Preparatory Access Projects:

This activity provided 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 Call Type C (PA C) with a cut-off every three months for evaluation of the proposals. Nine Preparatory Access cut-offs have been carried out in PRACE-4IP.

The report focuses on the optimization work and results gained by the completed projects in PRACE-4IP if those were not reported already as part of the earlier deliverable D7.1 [2]. In total five PA C projects have finished their work since the last deliverable. The statistics about the PA C calls in PRACE-4IP as well as a description of the call organization itself is also included. The results of the completed projects have also been documented in white papers, which were published on the PRACE-RI website [1].

In addition to the results achieved in the PA C activity, this deliverable will introduce a new type of PRACE Preparatory Access Call: PA Type D (PA D) which was established within PRACE-4IP and allows applicants to start their optimization work on a Tier-1 system with support from PRACE experts to ultimately reach Tier-0 system scalability. The document will present the call structure and provide initial statistics from the first cut-off, which took place in March 2017.

All the applications that have been accepted by this program have seen an increase in terms of performance or scalability of the application. These tasks have enabled computing at the Petascale, in preparation for the road to Exascale.

T7.1.B SHAPE:

This activity continued the support for SHAPE (the SME HPC Adoption Programme in Europe). SHAPE 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. It holds regular calls, and successful applicants to the SHAPE programme get support effort from a PRACE HPC expert and access to machine time at a PRACE centre. In collaboration with the SME, the PRACE partner helps them try out their ideas for utilising HPC to enhance their business.

This report focusses on the second and third call of SHAPE, looking at the results of the projects and lessons learned from the perspective of both the SMEs and the PRACE partners. In addition, it provides follow-up reports from the SME projects, which were completed last year, in order to evaluate the effectiveness of the SHAPE programme in terms of the SMEs’ ongoing engagement with HPC.

SHAPE participants have reported many positive outcomes for the businesses involved in the activity including tangible measures of the Return on Investment: new staff people hired, increasing sales and new business, lower costs, continued HPC Access, in-house HPC systems deployment, integration into the HPC ecosystem and engagement with other national and European industry engagement programmes.

1 Introduction

Computational simulations have proved to be a promising way of finding answers to research problems from a wide range of scientific fields. However, such complex problems often have so high demands regarding the needed computation time that these cannot be met by conventional computer systems. Instead, supercomputers are the method of choice in today's simulations.

PRACE offers a wide range of different Tier-0 and Tier-1 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 optimization 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 for Tier-0 systems. PA C and the new PA D is managed by Task 7.1.A "Petascaling and Optimization Support for Preparatory Access Projects". 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 C and PA D in PRACE-4IP are listed in 2.1. The new Preparatory Access scheme type D connects the national Tier-1 systems with the Tier-0 infrastructure by allowing applicants to start the optimization work on this intermediate level before targeting the Tier-0 platforms directly. This new type is described in Section 2.2. The review process, the assignment of PRACE experts to the projects and the monitoring of the support work are detailed in Section 2.3, Section 2.4 and Section 2.5 respectively. The contents of Sections 2.3 - 2.5 can already be found in deliverable D7.1 of PRACE-4IP [2]. They are repeated here for completeness and the benefit of the reader. Section 2.6 describes the relation and hand over between the PRACE-3IP, the PRACE-4IP and PRACE-5IP project regarding PA. Section 2.7 gives an overview about the Preparatory Access projects covered in PRACE-4IP, which were not reported in former deliverables. The announcement of the call is described in Section 2.8. Finally, the work done within the projects along with the outcome of the optimization work is presented in Sections 2.9 - 2.10.

The second part of this deliverable is the report on the SME HPC Adoption Programme in Europe (SHAPE), which is a pan-European programme 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 continued under PRACE-4IP.

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

It can be challenging for SMEs to adopt HPC. They may have no in-house expertise, no access to hardware, or be unable to commit resources to a potentially risky endeavour. This is where SHAPE comes in, by making it easier for SMEs to make use of high performance computing in their business - be it to improve product quality, reduce time to delivery or provide innovative new services to their customers. Successful applicants to the SHAPE programme get support effort from a PRACE HPC expert and access to machine time at a PRACE centre. In collaboration with the SME, the PRACE partner helps them try out their ideas for utilising HPC to enhance their business.

The initial SHAPE pilot [4][5] 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 return on investment for the SHAPE work [6]. Following this pilot the PRACE Council decided to operate the SHAPE programme as a permanent service. An overview of the SHAPE programme is provided in section 3.1, and an overview of the status of the second and third call projects is in section 3.2.

The SHAPE second call was launched November 2014 and closed in January 2015, and 11 SMEs were approved to receive assistance from SHAPE. Most of the second call projects finished last year, so a request for follow-up was sent to those SMEs to evaluate the impact of working with SHAPE on their business, as reported in section 3.3. Summaries of the progress of the other ongoing second call projects are in section 3.4. The third call for SHAPE was launched in November 2015, closed in January 2016, with its eight approved projects also reported on in section 3.4. The fourth SHAPE call opened June 2016 and closed September 2016 with four SME projects now underway, discussed in section 3.5. Finally, section 3.6 looks at the plans and recommendations for SHAPE going forward.

The deliverable closes with a summary in Section 4 highlighting the outcomes of Task 7.1.A and Task 7.1.B and a summarising conclusion in Section 5.

2 T7.1.A Petascaling & Optimisation Support for Preparatory Access Projects – Preparatory Access Calls

Access to PRACE Tier-0 systems is managed through PRACE regular calls, which are issued twice a year. 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 optimize their applications on the envisaged Tier-0 system prior to applying for a regular production project. This is the purpose of the Preparatory Access 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. 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 optimization by the user. Users get also a small number of core hours; the allocation period is 6 months.
- Type C is also designed for code development and optimization 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 to support the optimization. As well as access to the Tier-0 systems, the applicants also apply for 1 to 6 PMs of supporting work to be performed by PRACE experts.

- Type D allows PRACE users to start a code adaption 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 to test the scalability improvements. The work is supported by PRACE experts similar to Type C. The maximum duration of Type D projects is 12 months.

The following Tier-0 systems were available for PA during the reporting period:

- CURIE, BULL Bullx cluster at GENCI-CEA, France
- HAZEL HEN, Cray XC40, replacing HORNET, Cray XC40 at GCS-HLRS, Germany
- JUQUEEN, IBM Blue Gene/Q at GCS-JSC, Germany
- MARCONI, replacing FERMI ,Lenovo NextScale at CINECA, Italy
- MARENOSTRUM, IBM System X iDataplex at BSC, Spain
- PIZ DAINT, Cray XC30 System at CSCS, Switzland
- SUPERMUC, IBM System X iDataplex at GCS-LRZ, Germany

For the Type D access scheme the following Tier-1 systems were available for the first PA Type cut-off:

- ABEL, MEGWARE MiriQuid at UiO, Norway
- ARCHER, Cray XC30 at EPCC, United Kingdom
- ARIS, IBM NextScale at GRNET, Greece
- BEM, Haswell based Cluster at WCSS, Poland
- CARTESIUS, Bull Bullx B720/B710 at SURFsara, Netherlands
- CY-TERA, IBM System X iDataplex at CaSToRC, Cyprus
- GALILEO, Lenovo NextScale , at CINECA, Italy
- SALOMON, SGI ICE-X at VSB-TUO at IT4I, Czech Republic
- SISU, Cray XC40 at CSC, Finland

2.1 Cut-off statistics

In PRACE-4IP, nine cut-offs for PA took place resulting in nine new projects so far. Two projects, which started in February 2017 from the December 2016 cut-off are to be taken over by PRACE-5IP. Similarly, one project was taken over from the last cut-off of the extension phase of PRACE-3IP, as written in D7.1 [2]. In total, eight projects were supported by PRACE-4IP T7.1.A. Two of these were already reported in the last deliverable. One project is ongoing.

In the June 2015, September 2015 and June 2016 cut-offs no new proposals had been accepted. In June 2015, only one single proposal applied for PA C but this project was already supported by SHAPE and used PA C to gain the necessary computing time. In September 2015 and June 2016, no new proposals applied for PA C.

The PA D cut-off took place for the first time in February 2017, which was a special cut-off to introduce PA D. For simplicity this special cut-off is merged together with the usual March 2017 cut-off in the following graphs. One PA D proposal was submitted for the special cut-off and three additional PA D proposals were submitted for the usual March cut-off. Accepted projects from the March 2017 cut-off will be supported by PRACE-5IP.

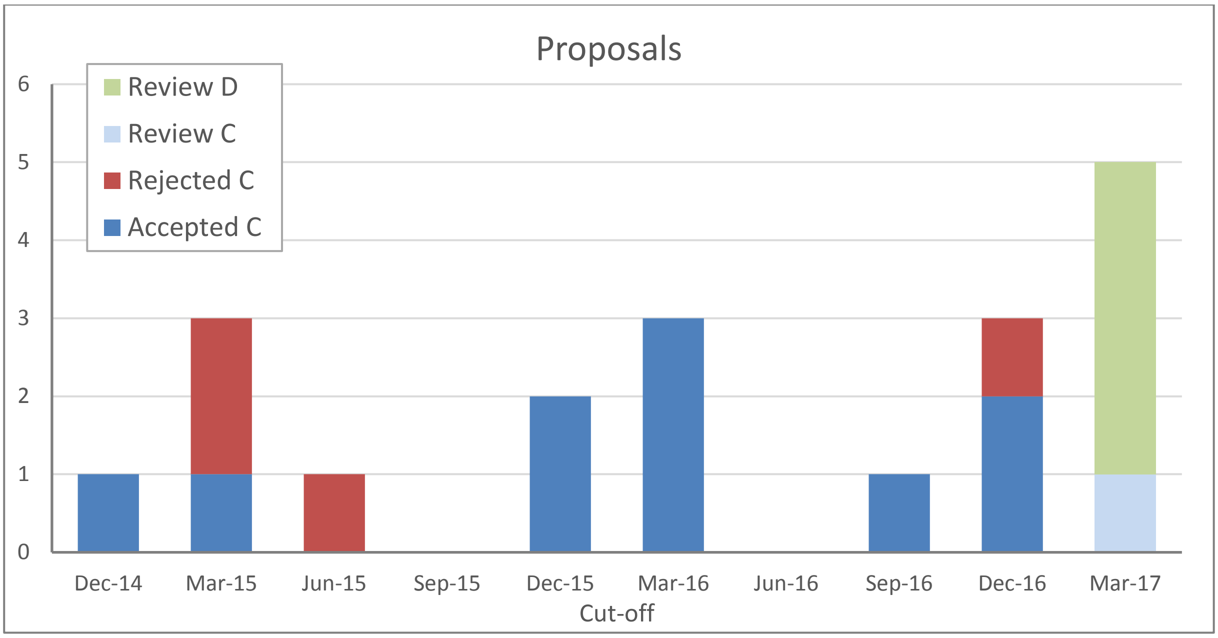

Figure 1: Number of proposals for PA type C and type D per cut-off.

Figure 1 presents the number of proposals that have been accepted, rejected or are currently under review respectively for each cut-off. In total 9 out of 13 proposals were accepted during the PRACE-4IP cut-off phase beginning in March 2015 until the cut-off in December 2016. The reviewing process of the March 2017 cut-off is currently in progress and therefore the final status is not yet available for the report.

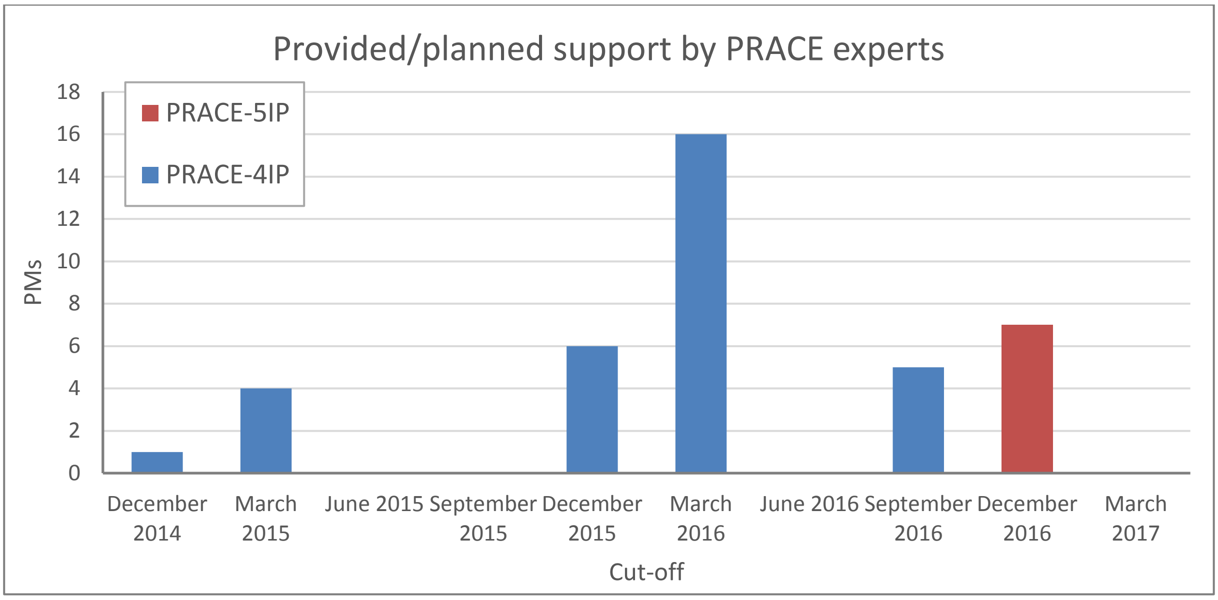

Figure 2: Amount of PMs assigned to PA projects per cut-off

Figure 2 gives an overview of the number of PMs from the PRACE project assigned to the projects per cut-off. In total 32 PMs were made available to these projects within PRACE-4IP.

Finally, Figure 3 provides an overview of the scientific fields, which are covered by the supported projects in PRACE-4IP.

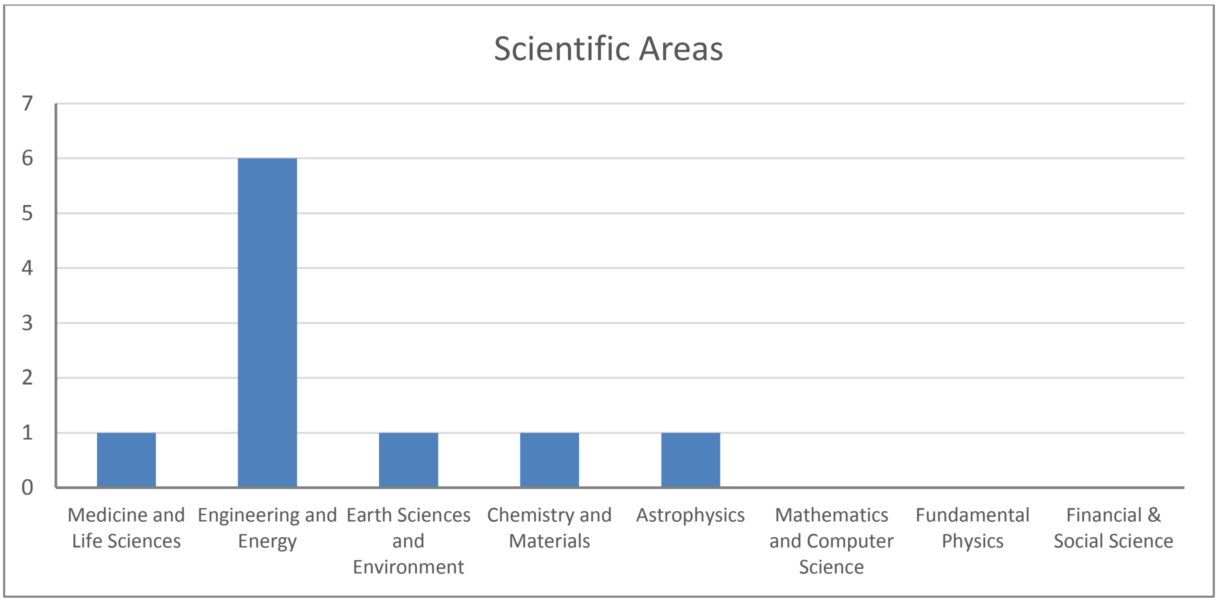

Figure 3: Number of projects per scientific field.

2.2 Preparatory Access Type D

Within the scope of PRACE-4IP a new type of preparatory access scheme was designed and finally established with a first cut-off in February 2017.

Preparatory Access Type D offers users the opportunity to start the optimisation work on a PRACE Tier-1 system to ultimately reach PRACE Tier-0 system scalability.

This access scheme is supposed to bridge the gap between the local university or site-specific computing resources and the PRACE Tier-0 platforms. The intermediate access to a Tier-1 system introduces an intermediate scaling level without having to face greater restrictions in the beginning of the optimisation work. It also helps to find the relevant bottlenecks, which have to be solved to reach Tier-0 scalability.

In addition to computing time on a Tier-1 and finally on a Tier-0 system, applicants will receive support by PRACE HPC experts of up to six person-months of effort. The expert will help to find relevant bottlenecks in the application and take part in the optimization work.

The maximum duration of Type D projects is twelve months.

As shown in the beginning of section 2, several Tier-1 systems, covering multiple architectures, are available for users of PA D. In contrast to PA A to C the exact system is not directly selected by the applicant. The system is selected based on preferences of the applicant and suggestions of the reviewers and the PRACE expert. This process helps new applicants and scientists who are not yet familiar with the different HPC architectures.

The final access to a Tier-0 system is used to perform at least a scalability run to validate the improvement of the code. To grant this access, a separate PA A proposal must be submitted during the runtime of the PA D project, which will then be connected to the ongoing project. The particular Tier-0 system is selected by the applicant with the help of the PRACE expert.

In total four proposals were submitted in the first PA D cut-off and are currently under review.

2.3 Review Process

The organization of the review procedure, the assignment of PRACE collaborators and the supervision of the PA C projects are managed by task 7.1.A. 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 work package 7 in order to assess their optimization 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 four 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.

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.A 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.

Based on the provided information from the reviewers the Board of Directors has the final decision on whether proposals are approved or rejected. The outcome is communicated to the applicant through PRACE aisbl. Approved proposals receive the contact data of the assigned PRACE collaborators, rejected projects are provided with further advice on how to address the shortcomings.

2.4 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.

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.

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 C or PA D project.

This procedure has proven to be extremely successful. No proposals had to be rejected in the past reporting period due to a lack of available support.

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 7.1A 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.

2.5 Monitoring of projects

Task 7.1.A includes the supervision of the Type C and the upcoming Type D projects. This is challenging as the projects' durations (six months for PA C and twelve months for PA D) and the intervals of the cut-offs (3 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 takes place in task 7.1.A every two months to discuss the status of running projects, to advice on how to proceed with new projects and to manage the finalization and reporting of finished projects.

In addition, the T7.1.A 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.A 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.

2.6 Hand-over between the different PRACE IP projects

The support for Preparatory Access Type C projects has been and is part of all PRACE projects (PRACE-1IP, -2IP, -3IP, -4IP, -5IP). For the hand-over between the projects, the tasks decided to treat the regarding projects in the following way:

The project out of the December 2014 cut-off started at February 16th, 2015. It was completely supported by PRACE-4IP and already reported within the last PRACE-4IP deliverable.

The two accepted projects from the cut-off in December 2016 were completely taken over by PRACE-5IP as these projects started by the end of February 2017.

All projects from the cut-off in March 2017 will be taken over by PRACE-5IP.

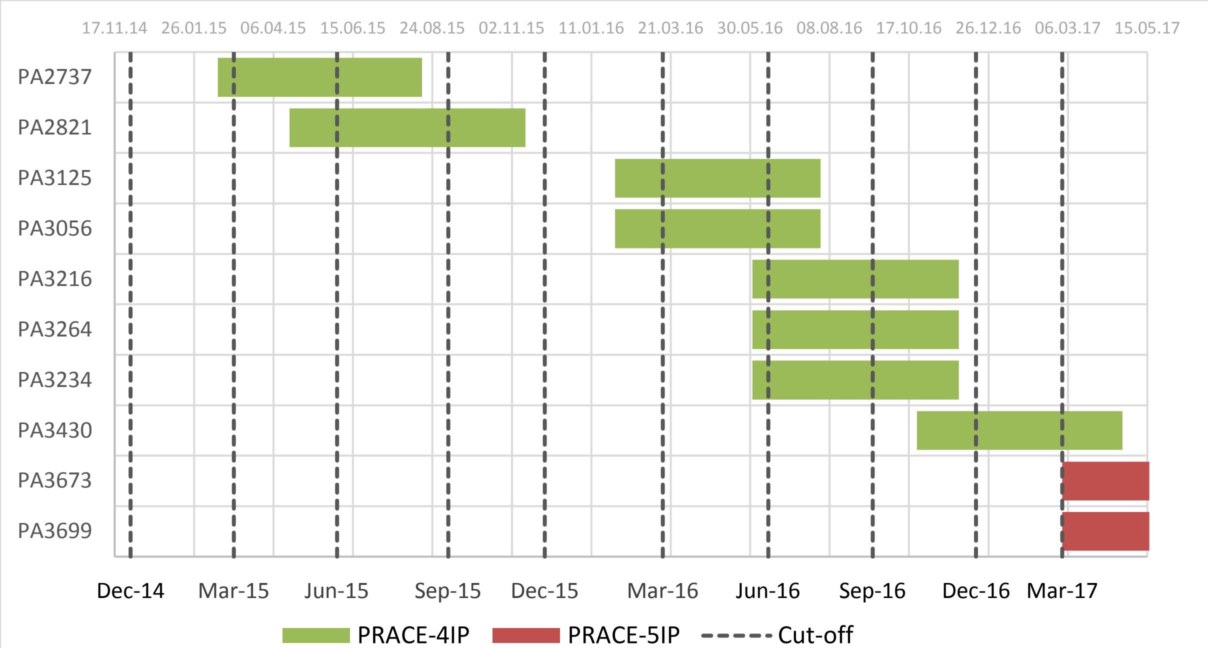

Figure 4: Timeline of the PA projects.

The timeline of these projects is shown in the Gantt chart in Figure 4. The chart shows the time span of each project. Projects, which are supported by PRACE-4IP, including the first project which were already accepted within the extension phase of PRACE-3IP, are shown in green. The projects which were handed over and are supported by PRACE-5IP are shown in red.

The slightly different starting dates of the projects per Cut-off is the result from the decisions made by the hosting members, which determine the exact start of the projects at their local site. Additionally, PIs can set the starting date of their projects within a limited time frame.

2.7 PRACE Preparatory Access projects covered by this report

Projects which were completed before March 2016 were already reported in deliverable D7.1 [2]. One project from the September 2016 cut-off started at the end of October 2016 and will run until end of April. The final outcome of the deliverable will be reported as part of the first T7.1A deliverable in PRACE-5IP. All projects which were taken over by PRACE-5IP will also be reported in the upcoming deliverable. In total five projects are reported in the current deliverable. Table 1 lists the corresponding projects.

Table 1: Projects, which are reported in this deliverable.

Cut-off December 2015	
Title	**Optimization of Hybrid Molecular Dynamics-Self Consistent Field OCCAM CODE**
Project leader	Antonio De Nicola
PRACE expert	Chandan Basu

Cut-off December 2015	
PRACE facility	MARENOSTRUM, FERMI, HAZEL HEN
PA number	2010PA3125
Project's start	01-Feb-2016
Project's end	31-Jul-2016
Title	**HOVE Higher-Order finite-Volume unstructured code Enhancement for compressible turbulent flows.**
Project leader	Claudia Adduce
PRACE expert	Thomas Ponweiser
PRACE facility	SUPERMUC, HAZEL HEN
PA number	2010PA3056
Project's start	01-Feb-2016
Project's end	31-Jul-2016

Cut-off March 2016	
Title	**Parallel mesh partitioning using space filling curves**
Project leader	Guillaume Houzeaux
PRACE expert	Mariano Vázquez
PRACE facility	MARENOSTRUM
PA number	2010PA3216
Project's start	01-Jun-2016
Project's end	30-Nov-2016
Title	**Simulation of the Ocean Cleanup Array**
Project leader	Monica de Mier Torrecilla
PRACE expert	Beatriz Eguzkitza
PRACE facility	MARENOSTRUM
PA number	2010PA3264
Project's start	01-Jun-2016
Project's end	30-Nov-2016
Title	**Parallel curved mesh subdivision for flow simulation on curved topographies**
Project leader	Xevi Roca
PRACE expert	Abel Gargallo-Peiró
PRACE facility	MARENOSTRUM
PA number	2010PA3234
Project's start	01-Jun-2016

Cut-off March 2016	
Project's end	30-Nov-2016

All remaining projects, starting with the project from the cut-off in September 2016, have to be reported in the next deliverable as those projects are still ongoing. Table 2 lists the corresponding projects.

Table 2: Currently running PA projects.

Cut-off September 2016	
Title	**Optimization of REDItools package for investigating RNA editing in thousands of human deep sequencing experiments.**
Project leader	Ernesto Picardi
PRACE expert	Andrew Emerson
PRACE facility	MARCONI
PA number	2010PA3430
Project's start	24-Oct-2016
Project's end	23-Apr-2017

Cut-off March 2017	
Title	**OPmized mulTI-fluid plasMA Solver (OPTIMAS)**
Project leader	Andrea Lani
PRACE expert	Chandan Basu & Thomas Boenisch
PRACE facility	HAZELHEN, JUQUEEN
PA number	2010PA3673
Project’s start	01-Mar-2017
Project’s end	31-Aug-2017
Title	**Water droplets and turbulence interaction inside warm cloud -- clear air interface**
Project leader	Daniela Tordella
PRACE expert	Andrew Emerson
PRACE facility	MARCONI, SUPERMUC
PA number	2010PA3699
Project’s start	01-Mar-2017
Project’s end	31-Aug-2017

The evaluation of the March 2017 proposals is currently in progress and they are therefore not listed here.

2.8 Dissemination

New PA Cut-offs are normally announced on the PRACE website [1].

After no new proposals where received in the Cut-offs in September 2015 and June 2016, PRACE sites were asked to distribute an email to their users to advertise preparatory access and especially the possibility of the dedicated support.

The new PA Type D was announced by a press release through the PRACE website and in the PRACE social media channels.

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 [1] and are also referenced by this deliverable.

2.9 Cut-off December 2015

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

The application evaluation form ensures a consistent and coherent presentation of all projects, which were managed in the context of PA. Within the preparation of PA D the evaluation form was updated to have a consistent evaluation for all types of PA. The projects started in December 2015 were still evaluated by using the older layout of the form. Additionally the white papers created by these projects are referenced so that the interested reader can seek further information.

2.9.1 Optimization of Hybrid Molecular Dynamics-Self Consistent Field OCCAM CODE, 2010PA3125

Code overview:

Beside ordinary coarse-grained and atomistic models, the OCCAM code is able to perform hybrid Particle-Field Theoretic Molecular Dynamics simulations. The recent hybrid Particle-Field technique combines Molecular Dynamics (MD) and self-consistent field theory (SCF). The main feature of the hybrid MD-SCF method is that the evaluation of the nonbonded forces between particle pairs is replaced by an evaluation of an external potential dependent on the local particle density. Using an efficient parallelization scheme it is possible to reach, by using coarse-grained models with chemical details, large time and length scales not accessible to the standard MD method [7][8].

Main objectives:

The main objective is to rewrite the input/output (I/O) routines of the parallel version of the OCCAM code. The original I/O implementation in the OCCAM code has several drawbacks. The code needs a large number of input/output files – as many as there are MPI ranks. For a large number of MPI ranks, it is difficult to manage such a large number of files.

Accomplished work:

The current I/O implementation in the OCCAM code, and pre-processor code IOPC used to build input files, is a strong limitation for parallel execution with a large number of MPI ranks. Moreover, the local operating system can impose a restriction on the number of files that be accessed simultaneously. In addition, the serial nature of the I/O in the current OCCAM code is another limitation. In the following code, a typical reading operation is reported.

```
do i = 1, loop
    read(MPI_id, *) x(i), y(i)
end do
```

The writing operation is done in a similar way. For each MPI ranks, several reading/writing operations, as reported in the example, are performed. All these small I/O operations, done at a prefixed step interval (for example the frequency writing of an MD trajectory), are not optimal for a parallel application, with the result that the I/O becomes a bottleneck for the whole code.

In order to reduce as much as possible the bottleneck due to the serial I/O, a new I/O scheme has been developed. In particular, a library called fpioLib, based on a parallel I/O has been written and implemented in OCCAM. In the new I/O scheme, each MPI rank reads information from the same input file, and each MPI rank writes own outputs in the same output file. In this way, the number of input/output files is reduced from N=number MPI ranks to a single input and output file. In this way we solved the problem related to managing a large number of files. Moreover, in order to reduce the I/O events for each MPI rank a large buffer read/write function has been implemented in the fpioLib. The set-up of this function can be done directly by a specific function included in the fpioLib Library.

Main results:

In order to test and measure the performances of the new parallel I/O implementation, tests were performed on two different systems. The description of the systems is reported in the subsection System Composition. The maximum number of processors used in the test is 4096.

The I/O routines of OCCAM have been rewritten to implement the parallel I/O, and to make a suitable interface with the fpioLib Library. Approximatively, 5 % of the OCCAM code has been rewritten, while the fpioLib Library is a completely new part of the code.

System Composition:

The systems simulated for the tests are realistic systems of surfactant in water solution. A coarse-grained (CG) model was used with a mapping very close to the atomistic (4:1), in which every effective CG bead corresponds to 4 heavy atoms. More details about the CG models can be found in the references [9][10].

System 1: 3,220,208 particles. There are 3 different molecule types and 12 different particle types. The box size of the system is: 67.2 x 62.2 x 79.9 nm3, the number of mesh points used in the simulation is 995328 (96x96x108). The density field, used for the calculation of intermolecular interactions, is updated every 100 steps. In Figure 5, the performances measured as 1E6 step/day, are reported for the old and new parallel I/O implementation.

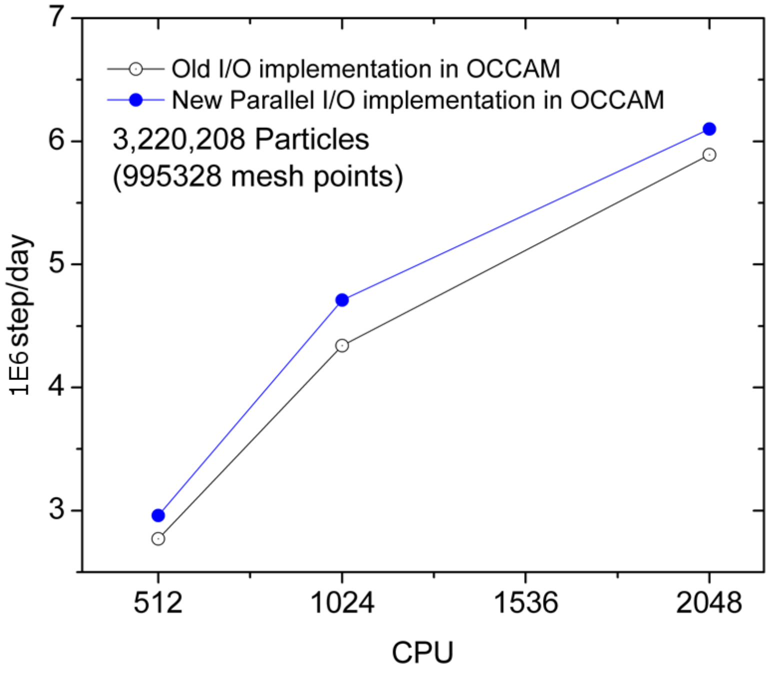

Figure 5: Step/day for the system 1 as function of no. of CPU.

System 2: 12,880,832 particles. Here there are 3 different molecule types and 12 different particle types. The box size of the system is: 134.4 x 134.4 x 79.9 nm3, the number of mesh points used in the simulation is 3981312 (192x192x108). The density field is updated every 100 steps. In Figure 6, the performances measured as 1E6 step/day, are reported for the old and new parallel I/O implementation.

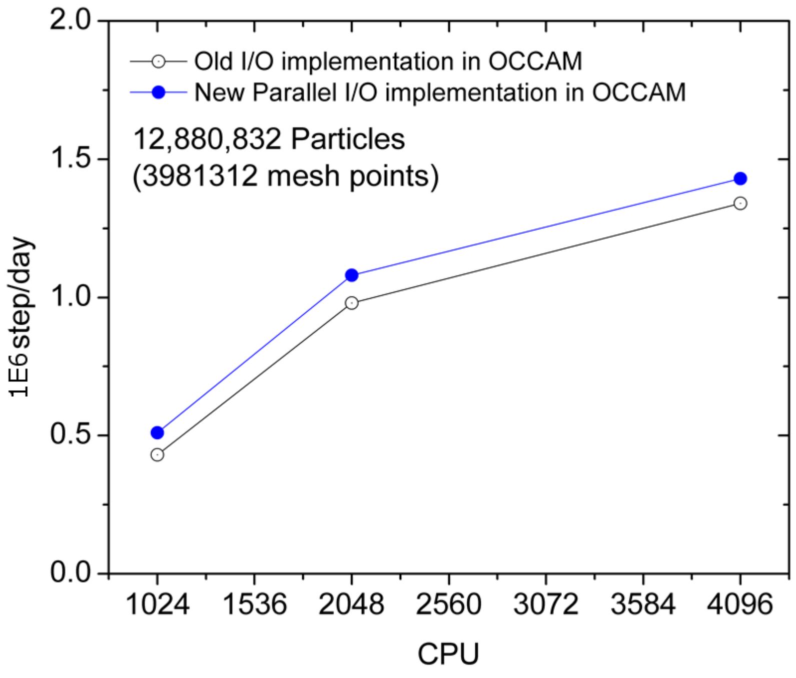

Figure 6: Step/day for the system 2 as function of no. of CPU.

As can be seen from Figure 5 and Figure 6, the performances with the new I/O scheme give an increase of about 6-18 %. Moreover, the number of files is significantly reduced to one each for input and output. Due to the efficiency of the fpioLib Library, now it is easy to gain access to perform a simulation on a larger number of CPUs, and also a reduction in the time to rebuild the whole trajectory in the post-production.

Computational Resources: CPU Intel Haswell Intel Xeon E5-2670 @2.6GHz 64-Bit, 16 cores/node with 64 GB RAM (4GB/core). Network communication InfiniBand QDR 40 Gb/s with a latency 0.7 microseconds.

The project also published a white paper, which can be found online under [11].

2.9.2 HOVE Higher-Order finite-Volume unstructured code Enhancement for compressible turbulent flows., 2010PA3056

Code overview:

UCNS3D Compressible CFD code, using 2D and 3D hybrid unstructured meshes and solving the compressible inviscid, viscous Navier Stokes Equations and the Reynolds Averaged Navier Stokes equations.

Main objectives:

The primary tasks were to pinpoint the sections in the code where most of the time is spent, implement calls to BLAS functions using the MKL libraries instead of using the matmul intrinsic Fortran functions, rearrange the memory access pattern of some of the matrices to speed-up performance, and implement non-blocking MPI communications.

Accomplished work:

The BLAS routines (GEMM, DOT) from Intel MKL where used to replace the matmul functions, and at the same time for two matmul calls the operations were reduced to one matrix multiplication, since an inverse transposed matrix that remained unchanged during the computations could be prestored. The non-blocking MPI isend and irecv replaced the MPI_SENDRECV, and the matrix indices where rearranged to obtain better memory access patterns. It should be noted that the greatest improvement was seen from the implementation of the BLAS functions, and prestoring the unchanged matrix.

Main results:

For the three sets of applications (ILES(TGV), LES, RANS) it was noticed that this implementation offered speedups for the very high-order schemes of spatial accuracy greater than 3. Specifically, the time taken per iteration was 2-6 times less than the original implementation as seen in the figures below.

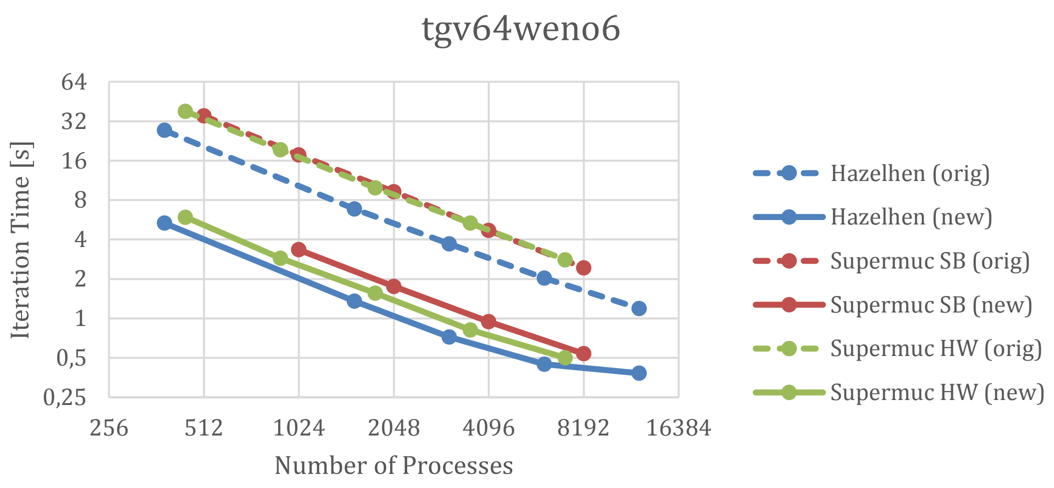

Figure 7: UCNS3D TGV iteration time comparison

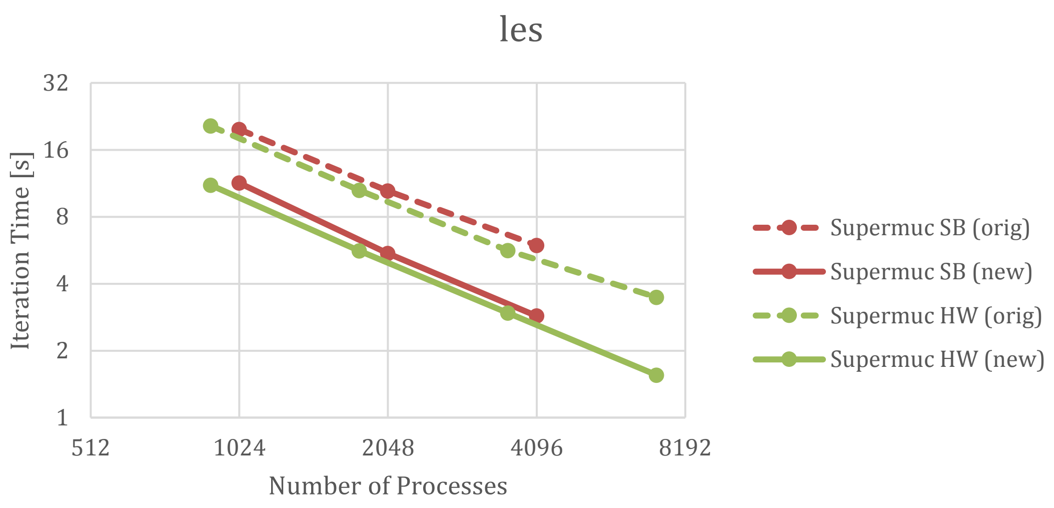

Figure 8: UCNS3D LES iteration time comparison

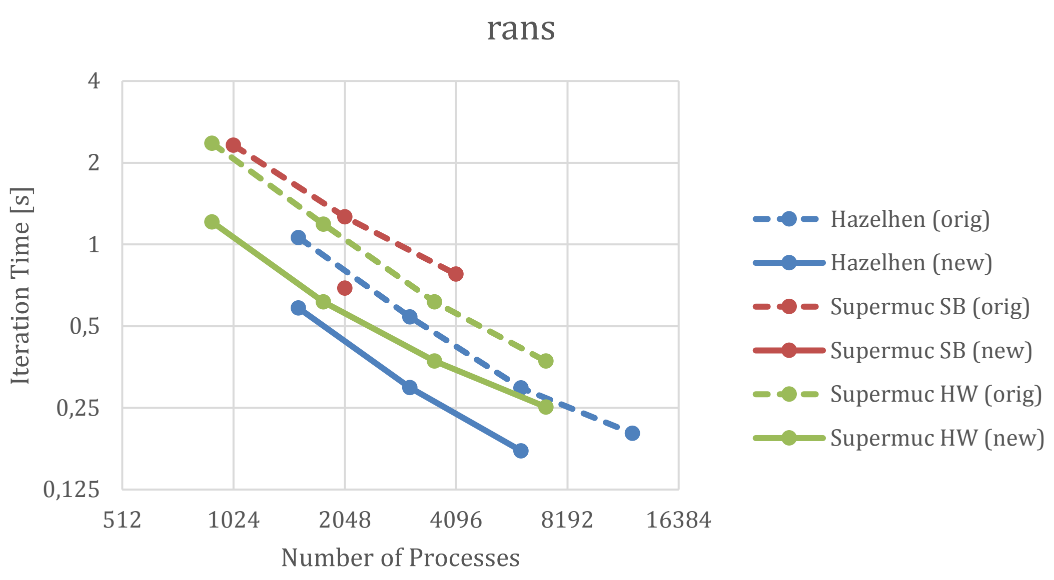

Figure 9: UCNS3D RANS iteration time comparison

The project also published a white paper, which can be found online under [12].

2.10 Cut-off March 2016

As written in Section 2.9 the layout for the final evaluation report was slightly changed starting with the projects of the March 2016 cut-off. The new layout is directly comparable between all types of PA and will be used for the upcoming reports.

2.10.1 Parallel mesh partitioning using space filling curves, 2010PA3216

Overview

The project aim is implementing a parallel mesh partitioner based on space filling curves in Alya, a high performance computational mechanics code.

Traditional mesh partitioning (METIS [13], SCOTCH [14]) is generally a sequential computational kernel. Although semi-parallel versions exist (PARMETIS [15], PTSCOTCH [14]), their performances in terms of load balancing and parallel efficiency remain low. In addition, one has to use a much more reduced range of MPI tasks to perform the partitioning in order to achieve a reasonable load balance. This poses therefore implementation and memory issues. All these drawbacks make the runtime re-load balancing (e.g. required for a mesh adaptation process) a bottleneck in the simulation process.

A strategy based on space filling curves can be used to implement a fully parallel mesh partitioner, which overcomes the mentioned drawbacks. This makes the partitioning and load balancing a fully parallel task that can be interleaved on the simulation process. The proposed project aims at implementing such a strategy in the Alya code.

Scalability results:

The aim of the tests carried out on MareNostrum III supercomputer was to compare the graph based partitioning implemented in the standard library METIS and an in-house implementation based on space filling curves (SFC). Both the cost of the partition process and the quality of the partition obtained are compared. The largest case was a 17 million element mesh representing the upper respiratory system, which we run on up to 1024 CPUs. The time saved in the mesh partitioning was over 50x (as compared with sequential metis) and, more importantly, with a better partition quality in terms of load balance. The theoretical imbalance was reduced by 15% on average and the computing time of the matrix assembly process (the target of this study) was reduced by 25% on average.

These results exceeded expectations and one could conclude that SFC based mesh partitioning is a robust solution for high performance computing.

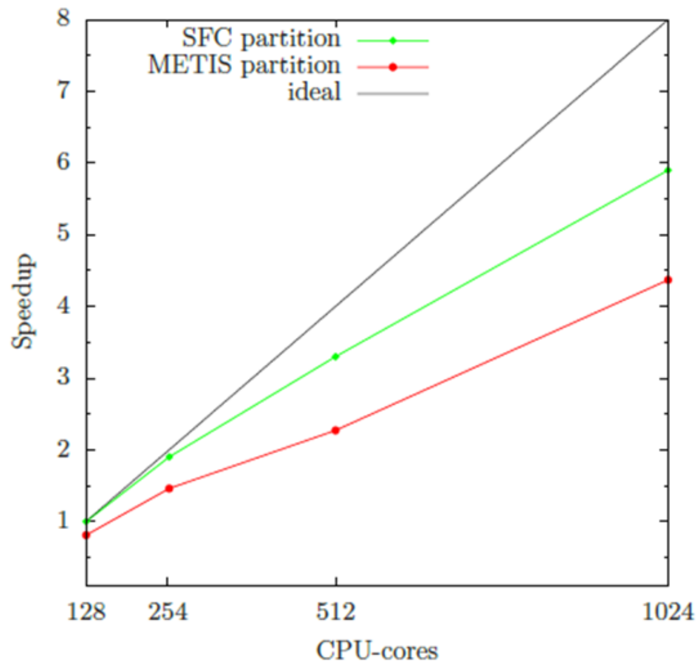

Figure 10: Strong speedup of the matrix assembly for the respiratory system case. The mesh is of 17M nodes. Results are normalized by the time achieved using SFC partitioning

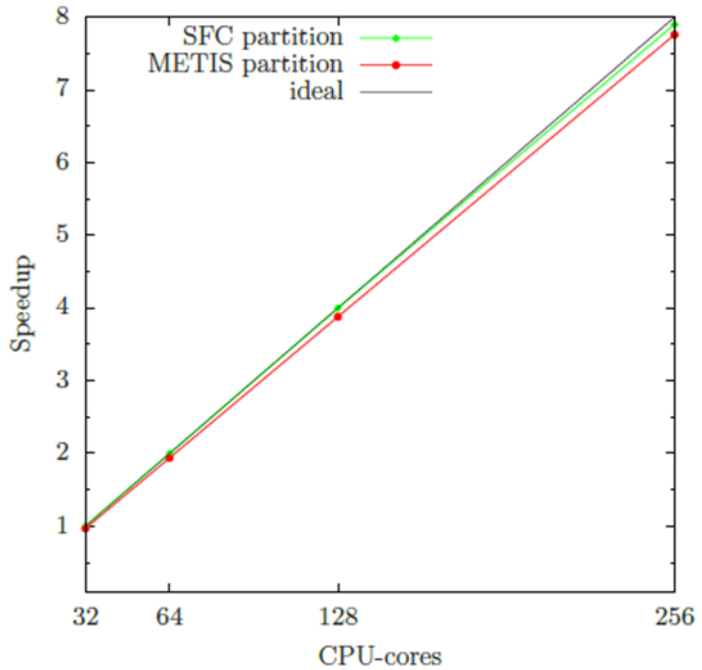

Figure 11: Strong speedup of the matrix assembly for the fusion reactor case. The mesh is of 100K nodes. Results are normalized by the time achieved using SFC partitioning.

Accomplished work:

The partitioning strategy implemented is based on a Hilbert space filling curve strategy. First a bounding box of the geometrical domain of the simulation is considered. This bounding box is divided into bins, each one with an assigned weight proportional to the number of mesh elements contained in it. Then the bins are aggregated on as many subsets as subdomains desired for the partition, such that the weight of each subset is as similar as possible. In order to obtain good properties in terms of locality, the order followed to aggregate the bins is defined by means of a Hilbert space filling curve. The strength of this implementation is that this is done in parallel, each process accounts for a part of the initial bounding box, but the process distribution and the local SFC defined by each one are generated in such a way that the partition obtained is independent of the number of processes engaged on it.

Results of the enabling process:

The result is a new partitioning algorithm implemented in Alya. It provides a robust solution for parallel mesh partitioning which avoids any memory bottle neck, and the partition can be done by the same number of parallel process that perform the simulation, avoiding resources waste.

Main results:

The conclusion of this project is that SFC is a very competitive approach for parallel mesh partitioning. Its integration in Alya overcomes some bottlenecks of standard libraries, such as METIS, to run at the largest scale. This project enhances the capabilities of the Alya code in terms of scalability, enabling the possibility of run-time load balancing, required for some specific algorithms like adaptive mesh refinement, or to attain the heterogeneity of modern architectures.

The implementation of the SFC partitioning algorithm is fully parallel and avoids any bottleneck in terms of memory. Speedup of 50x comparing with the METIS sequential algorithm at partitioning a mesh of the respiratory system with 17M elements using 512 CPU-cores has been obtained, without sacrificing the quality of the partition which in fact provides a better load balancing (15%) and acceleration of the explicit part of the time-step (25%).

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

2.10.2 Parallel curved mesh subdivision for flow simulation on curved topographies, 2010PA3234

Overview

Standard commercial linear mesh generators (e.g. ANSYS [17]) are targeted to approximate geometries. To this end, they feature boundary layer meshing, and local adaptivity. However, they are not ready to generate meshes suitable for petascale runs, mainly due to memory constraints. Furthermore, they only provide piecewise linear (straight-sided tetrahedra) or trilinear (ruled-sided hexahedra) approximations of the computational domain and therefore, subdivided meshes do not target to approximate curved boundaries. In this project a parallel uniform curved mesh subdivision method has been implemented in Alya, a HPC code developed at Barcelona Supercomputing Centre.

The scientific case has been prompted by the interest of combining the research results of two of BSC's ongoing projects, both in the Horizon 2020 framework of the European commission funded by the programs Energy oriented Centre of Excellence for computer applications (EoCoE, contract H2020-EINFRA-2015-1) and Marie Sklodowska-Curie Actions (contract HiPerMeGaFlowS), respectively. The goal of this combination is to obtain mesh convergence results for the turbulent flow on real curved topographies, with applications to wind farm design.

To this end, the chosen approach starts by generating a coarse linear mesh. Then the former coarse mesh is curved to match the boundaries of the computational domain. Finally, the curved mesh is given to the improved Alya code that now reads the curved mesh, partitions it, and sends the subdomain meshes to the slaves. The result is a finer linear mesh obtained in parallel with improved geometric accuracy. Then, a steady state flow solution is obtained on the finer linear mesh.

The advantage of the proposed and implemented strategy is that the subdivision is performed and stored in parallel and therefore, there are no memory constraints. Furthermore, the finer linear mesh targets the curvature information described by the first curved mesh. Finally, if the original numbering is conserved, then the post processing can be performed on the coarse curved mesh.

Scalability results:

The speedup of the initial Alya code was almost linear for a sufficient number of elements per subdomain, say 20,000. The same performance results are obtained, however in the present proposal it is possible to solve larger problems where the inherent curvature of the topography is now preserved. This was not possible in the previous version of the code where after successive refinement, the geometric accuracy of the topography representation was not improved.

Accomplished work:

1. It is assumed that a coarse linear mesh has been obtained and that the mesh features the required resolution for the current problem (CFD) such as boundary layer elements near the wall.
2. The elements of the former coarse mesh are curved to match the boundaries of the computational domain. In this manner, the inherent curvature of the domain can be reproduced with more accuracy. To this end, we use a curved meshing method developed by the researchers at the BSC. The method uses piece-wise polynomial representations of the elements that can be curved to match the domain boundaries by minimizing the mesh distortion.
3. The coarse curved mesh is given to Alya, which is based on a master-slave strategy. The master reads the curved mesh, partitions it, and sends the subdomain meshes to the slaves. After this task, the original coarse curved mesh no longer exists, neither does any array defined on it. The slaves then refine their corresponding subdomains in parallel, by performing the subdivision of the curved elements in finer linear elements. During the process it is necessary to ensure that the boundary nodes of the new elements match with the curved elements of the neighbouring subdomains. The result is a finer linear mesh obtained in parallel with improved geometric accuracy. For standard engineering applications, the computational mesh has between 10M and 100M of elements whereas, the curved mesh that has to be partitioned and distributed has only between 10K and 100K of elements.

Results of the enabling process:

Starting with a coarse curved mesh composed by 60K hexahedra, and after four consecutive subdivisions, a finer distributed linear mesh composed by 16K x 8 x 8 x 8 x 8 = 65M elements can be obtained in seconds and stored completely in memory. Now, the resulting refined linear mesh targets the curvature of the topography while in the previous version of the code this was not possible. Using the finer mesh, it is possible to converge to a steady flow state with an improved accuracy of the topography geometry.

Main results:

A subdivision strategy performed and stored in parallel, thus avoiding memory constraints, has been proposed and implemented. Furthermore, the finer linear mesh targets the curvature information described by the initial curved mesh. The mesh can be used to obtain the steady state of a flow solution of interest in simulating the generated power of a wind farm. Finally, if the original numbering is conserved, then the post-processing can be performed on the coarse initial mesh.

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

2.10.3 Simulation of the Ocean Cleanup Array, 2010PA3264

Overview:

In this project a very strongly coupled multi-physics problem has been simulated which represents a huge challenge for efficiency and accuracy.

The trajectories of the plastic particles in an ocean are simulated with a Lagrangian transport method immersed in a two-phase flow solver. The position of the interface is tracking with the Level Set method. A long floating barrier to capture the particles is the other ingredient of the problem. The simulation has been done with the Alya code, which has demonstrated its high potential to face complex problems.

A wave generator has been implemented in order to evaluate the action of waves on the particle trajectories and the forces of the boom and skirt. This issue also adds more complexity in the Level set solution. Finally, accurate viscous solutions require refined meshes close to the skirt. In order to simulate flows at high Reynolds numbers the use of turbulence models is mandatory.

All of these ingredients result in a very complex simulation. The fluid-structure interaction has some peculiar difficulties like different physical descriptions, different mesh sizes, different numerical treatments, different time scales and physical and numerical instabilities.

The coupling with the particles simulation also implies problems related with the load imbalance and synchronization. In a distributed memory context, particles may be concentrated in very few MPI tasks and one cannot predict where they are going. One solution could be to use a library called DLB (dynamic load balancing) to load balance particles at runtime, and a multi-code strategy to achieve coarse grain asynchronism.

Scalability results:

In this work different numerical algorithms have been implemented to simulate the problem in hand as it will be described.

Regarding the performance of the simulation, the main issue to be improved is the efficiency. As explained earlier, the computational load is imbalanced and this drawback is attempted to be solved. Load imbalance implies losing efficiency. If efficiency is defined as the ratio between the summation of the times of each CPU and the maximum time multiplied by number of CPUs this could be seen as the load balance, defined as the ratio between the average time and maximum time. The average time is defined as the summation of the times of each CPU divided by the number of CPUs.

The target of the present development is to keep the efficiency and scalability of the code while making it able to solve multi-physics problems through multi-code coupling strategies. This means that multiple instances of the code will run in the same MPI environment, using the best configuration possible for each physical problem involved in the corresponding instance of the code.

Accomplished work:

In this project the difficulties related to the simulation of a multi-physics problem have been tackled. In particular, the interaction between a structure and bifluid with the level set method to capture the interface between air and water has been solved. A particle transport immersed in this fluid is also solved.

Each component of the global problem has its own challenges to be overcome. The level set method used in this project for tracking the moving interface has the main drawback in a poor mass conservation. The multiphase flow presents density and viscosity discontinuities in the interface which leads to numerical difficulties. A pressure enrichment strategy is used to improve the solution as well as a more accurate integration rule with cut elements in the interface.

These techniques have been implemented and validated in the Alya code to be used in the simulation of this project. The code based on finite element formulation has been developed in Barcelona Supercomputing Centre and is written in the FORTRAN programing language.

Particles are transported solving the Newton's second law considering a series of forces: drag, gravity, buoyancy and friction. The time integration is based on the Newmark's method. Only a one-way coupling between fluid and particles is considered such that the particles are transported by the fluid but have negligible effects on the fluid dynamics.

The parallelisation of the fluid and particle solvers is implemented with two levels of parallelism, MPI and OpenMP.

In the distributed memory context in which the Alya code is executed in this work the main difficulty to be overcome is the imbalance problem due to the particles being concentrated in a very few messages in the message passaging interface.

This can be solved using multi-code coupling between fluid and particle solver since independent partitions could be used for the fluid and particle computational domains. Another very useful strategy is a dynamic load balance applied at the shared memory level to fully exploit the available resources when MPI processes of the two solvers are idle.

The coupling with the structure is solved in this project via two-code coupling. This means that in one time step, the fluid mechanics problem is solved, the forces in the interface surfaces between fluid and structure calculated and sent to the solid mechanics code, which, after deformation by the effect of these forces, sends back the new position of the interface surface and the time step is advanced. The resources between solid and fluid are independent. This is another critical point for which dynamic load balance is a very useful tool.

With all this context it has been studied, which is the optimum configuration to solve the global coupled problem such that the maximum efficiency of the computational resources is achieved.

Figure 12 represents the global problem and in Figure 13 and Figure 14 different options analysed in this work are shown. ALE and N-S stand for Alefor and Nastin and are the modules to solve the fluid flow problem.

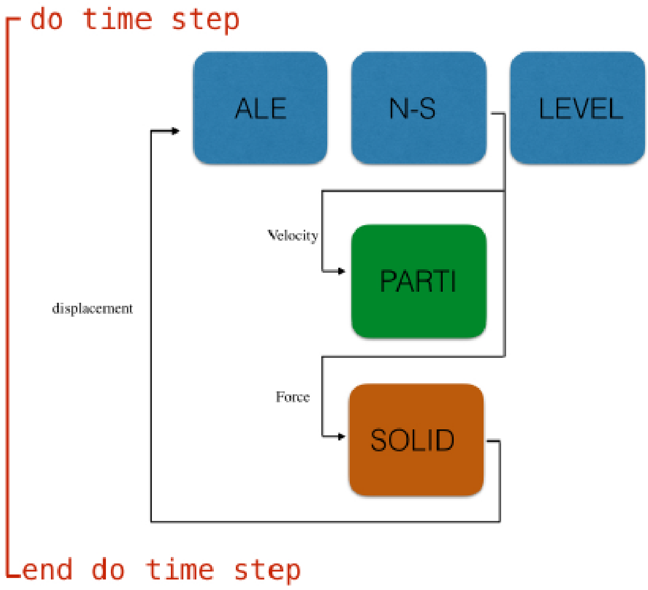

Figure 12: Global problem

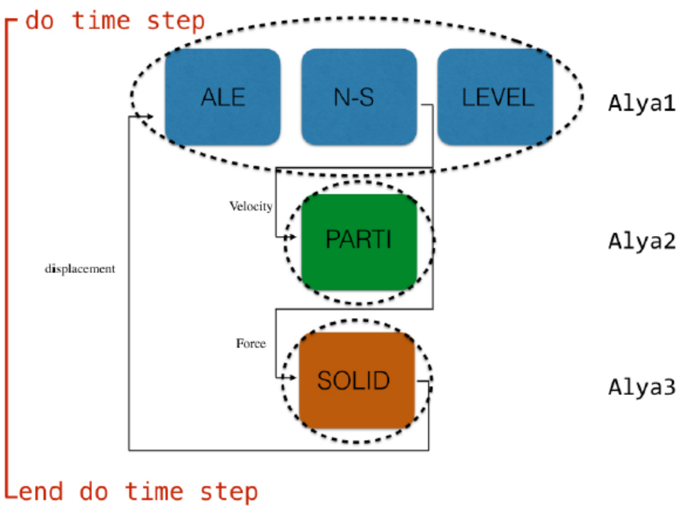

Figure 13: Three different instances of the Alya code to solve the problem

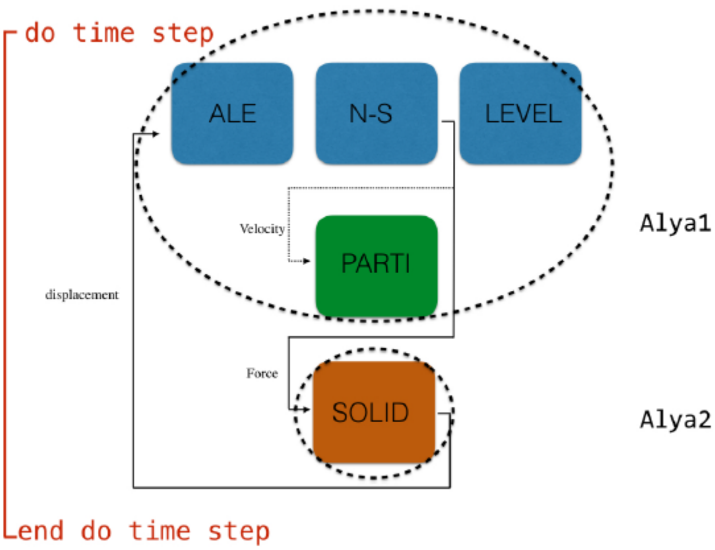

Figure 14: Two instances of Alya code.

Results of the enabling process:

As it has been explained in this report, to carry out this complex multi-physic simulation implies a lot of effort. The fluid-structure interaction has some peculiar difficulties like different physical descriptions, different mesh sizes, different numerical treatments, different time scales and physical and numerical instabilities. A fine mesh is mandatory when complex physics has to be captured. The simulation of particle transport leads to imbalance problems as explained before. All of these requirements mean that a lot of computational resources are necessary to solve these kind of problems.

The hybrid paradigm is very useful in terms of computational performance in the particle transport scenario. As discussed earlier, issues in load balancing can appear when particles are considered. For this case, a dynamic load balance library is a good solution since it can vary automatically, smartly and dynamically with the MPI/OpenMP distribution.

The computational cost to face with the problem in hand is huge and sometimes beyond a reasonable scope. Therefore, through this project the best strategy to optimize this effort and to be able to solve the complex problem has been studied. Improvement of the algorithms and implementing a different parallel environment have been crucial points to achieve this objective.

An enrichment pressure method to improve the solution of the bi-fluid has been implemented. As it is well known, for two different density fluids the interpolation errors in the pressure give rise to spurious velocities. Also, viscosity discontinuities can lead to discontinuous velocity gradients. Enrichment methods add degrees of freedom at elements cut by interface in order to reduce interpolation errors. In this work one pressure degree of freedom per cut element is added. In order to capture discontinuities and take advantage of these enrichment functions, the integration rules have been modified in elements cut by the interface.

Another new ingredient implemented in the Alya code to solve the problem has been the wave generator. This consists of a sector of the domain in which an analytical solution for the level set function is imposed and coupled with the numerical solution in order to generate a specific free-surface profile. To validate the free surface model, a 3-dimensional dam break problem has been carried out.

For the particle transport, the friction force has been implemented to consider the friction in the skirt.

A general difficulty to solve this complex multi-physics problem is calibration of all the parameters and details necessary for the simulation, not only at the physical level (time step, relaxation parameters, etc.) but at the computational level e.g. how are the resources divided? All of these questions are very difficult and analysed in this work.

Main results:

As conclusions of this project the following points can be outlined:

- The usability of the assigned PRACE system has been very good. The Mare Nostrum machine has powerful resources with a simple and quick system queue. The support and all the necessary information related to the Mare Nostrum account has been very clear and the problems or questions arising have been solved immediately and very efficiently.
- The results of this project have a great impact in two directions. On the one hand, the problem faced with this simulation is one of the major disasters of the humanity. An accurate simulation of the proposed solution could be a determinant for the success of the project. On the other hand, general problems presented in industry and science have to face up to this type of strongly coupled multi-physics problems. This work presents

some strategies to be more efficient solving these kind of problems. Simulations involving multi-physics, complex geometries and fine meshes capable of capturing accurate results need computational efficiencies like the ones proposed in this project.

The project also published a white paper, which can be found online under [19].

3 T7.1.B SHAPE

In this section, the progress in task 7.1B SHAPE will be discussed.

In the period since the previous deliverable (D7.1 [2]), the three unfinished second call projects have continued, along with the third call projects started last year. Follow-up requests have been sent to the SHAPE second call projects that concluded last year, to see how the work performed with the assistance of SHAPE has affected the SMEs' business – this is reported in section 3.3. Summary reports for the ongoing second and third call projects are reported in section 3.4, along with an overview of the lessons learned with regards to the implementation of the SHAPE programme. In addition, a fourth SHAPE call has been held, and the successful applicants' projects have now begun, as discussed in section 3.5. Finally, the future of SHAPE is discussed in section 3.6.

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.

It can be challenging for SMEs to adopt HPC. They may have no in-house expertise, no access to hardware, or be unable to commit resources to a potentially risky endeavour. This is where SHAPE comes in, by making it easier for SMEs to make use of high-performance computing in their business - be it to improve product quality, reduce time to delivery or provide innovative new services to their customers.

Successful applicants to the SHAPE programme get effort from a PRACE HPC expert and access to machine time 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 assisted over 30 SMEs (see the project website http://www.prace-ri.eu/hpc-access/shape-programme/ for examples), and the 5th call for applications will open in March 2017 so there will be more SMEs in the pipeline able to benefit from SHAPE.

The pilot programme was run under PRACE-3IP, and ten SMEs were given assistance via SHAPE. For the second call, eleven projects were approved, and eight more in the third call. Finally four further projects have been approved in the fourth call. Figure 15 shows the broad range of industry domains that have been represented in SHAPE so far, and Figure 16 the countries from which the participating SMEs originate.

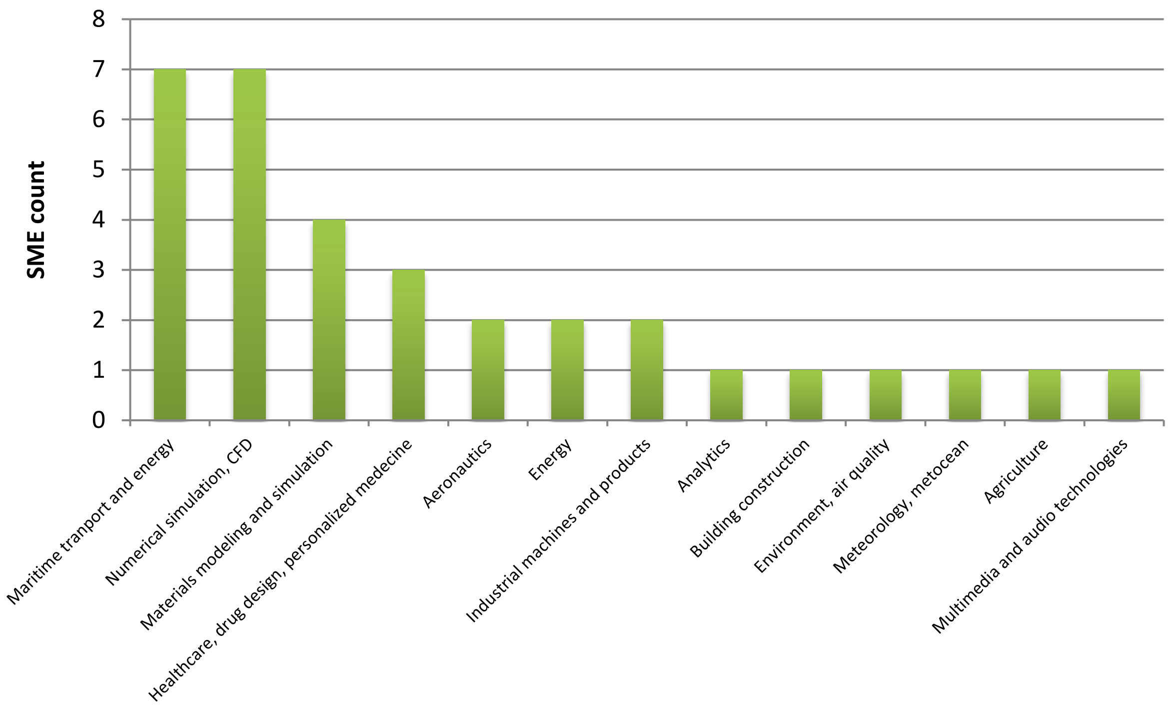

Figure 15: SHAPE: SME industrial domains

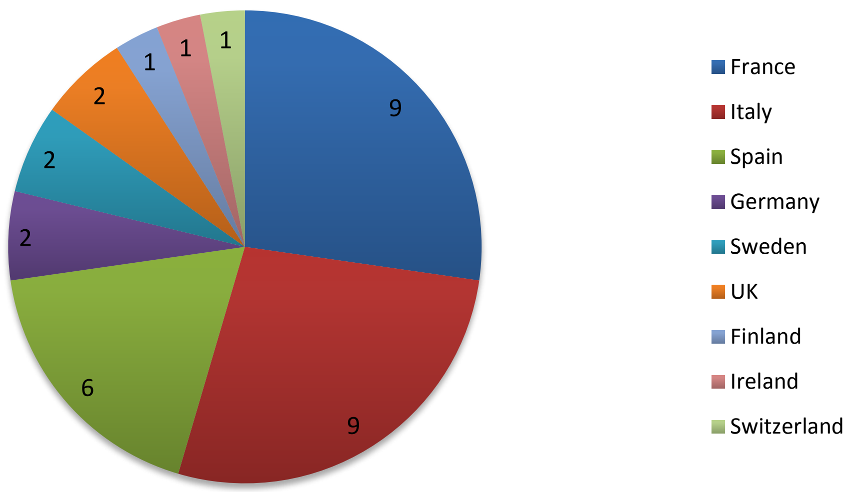

Figure 16: SHAPE: SMEs country of origin

The process for SHAPE is as follows:

- SHAPE Call 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 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 MB/TB 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 have been used 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: 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. On conclusion 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, 6-12 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.

The composition of the review panel has been changing slightly with each call, for the fourth call it was:

- Two members representing 4IP-WP7,
- Two members from the PRACE Board of Directors,
- One member from the Industry Advisory Committee,
- One member from the PRACE Peer Review Team,
- One member representing the preparatory access team

Given the current two-stage process of applying for SHAPE (apply to SHAPE, then if successful apply for machine time, usually via Preparatory Access), this final panel member is included to give a preliminary opinion on the technical suitability of the application from the PA point of view.

The criteria used in the review were:

- **Coherence with the goals of SHAPE** – the application should match the aim of SHAPE, that is, to assist SMEs with no or little experience of HPC to get involved with HPC, via both expert assistance and machine time from PRACE.
- **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 or recruit dedicated staff. The solution implemented should be part of a business plan to further engage in HPC in the long term.
- **Technical Adequacy** – the applications are expected to be realistic and achievable in the timeframe and resources available to the programme.
- **Commitment from the SME** – the level of commitment of the SMEs to co-invest with PRACE in achieving the project goals. The effort should be at least approximately equally split between the company and PRACE.
- **Social and Economic Impact** – the impact of the proposed project on society as a whole.

The applications are reviewed and ranked according to these criteria, then the final recommendations are put forward to the Board of Directors.

3.2 SHAPE status: second and third call projects

Here the status of the second and third call projects is reported. As mentioned earlier, where it has been some time (6-12 months) since the conclusion of a project, the participating SMEs have been asked to provide a follow-up report to assess the impact of the SHAPE work on their business. For the remaining ongoing projects, each were asked to provide a summary of their project for this document. On conclusion of the work, each project is expected to produce a white paper for review and publication on the PRACE website.

As of March 2017, the status of the second and third call projects is as described in Table 3 and Table 4 respectively. Note that the actual technical work for all of the second and third call projects has either finished, or is being finalised.

Table 3: SHAPE status - second call projects

SME	PRACE	White Paper	Reporting
Cybeletech	CINES	Published	Follow-up, section 3.3.1
OPTIMAD Engineering	CINECA	Published	Follow-up, section 3.3.2
Open Ocean	IDRIS	Published	Follow-up, section 3.3.3
Hydros Innovation	CINECA	Published	Follow-up, section 3.3.4
Vortex Bladeless SL	BSC	Published	Follow-up, section 3.3.5
Design Methods	CINECA	Published	Follow-up, section 3.3.6
IngenieurbüroTobias Loose	HLRS	Published	Follow-up, section 3.3.7
Ergolines SRL	ITU	Published	Follow-up, section 3.3.8
WB-Sails	CSC	Received, review Apr 2017	Summary, section 3.4.1
Principia	CINES	N/A (internal tech report due April 2017, see summary for details)	Summary, section 3.4.2
Algo’tech	INRIA	Due April 2017	Summary, section 3.4.3

Table 4: SHAPE status - third call projects

SME	PRACE	White Paper	Reporting
ACOBIOM	CINES	Received, review Apr 2017	Summary, section 3.4.4
Airinnova AB	KTH	Due April 2017	Summary, section 3.4.5
Creo Dynamics AB	KTH	Due April 2017	Summary, section 3.4.6
AmpliSIM	IDRIS	Due May 2017	Summary, section 3.4.7
ANEMOS SRL	CINECA	Received, review Apr 2017	Summary, section 3.4.8
BAC Engineering Consultancy Group	BSC	Due April 2017	Summary, section 3.4.9
FDD Engitec SL	BSC	Due April 2017	Summary, section 3.4.10
Pharmacelera	RISC	Due April 2017	Summary, section 3.4.11

The amount of PRACE effort allotted to the second and third call projects was 45.25 and 30.75 person-months respectively.

3.3 SHAPE Second call: Follow up for completed projects

This section presents the follow up reports from the SMEs which concluded their SHAPE work in 2015-2016. The key aim of this exercise is to ascertain how taking part in the programme has affected the SMEs: what are they now doing with HPC following the SHAPE activity, what is their return on investment, and so on.

A report template was distributed to the SMEs and PRACE partners, and included the following Key Performance Indicators (KPIs) affected by the SHAPE work to be considered (but not limited to):

- New customer acquisition
- Turnover
- Return on Investment (RoI)
- New product offers
- New service offers
- Access to new markets
- Enhanced software features
- Staff recruitment

The responses are presented below, and summarised in section 3.3.9

3.3.1 Cybeletech (France): Numerical simulations for plant breeding optimization

PRACE Partner: CINES

Project Duration: 1st May 2015 - 24th Dec 2015

Project overview

CybeleTech is a young SME whose aim is developing the use of numerical technologies in agriculture. The core products of CybeleTech are based on numerical simulations of plant growth through dedicated biophysical models. This new technology can bring added values at different stages of the agriculture and food chain: in varieties breeding, simulating plant growth can help reducing the amount of field trials by approximately 50% and consequently reduce the time needed to produce a new cultivar. Plant growth simulations are also used to optimize cropping practices, helping farmers to save resources and to maximize yields.

Another application is the forecast of yields and production at large scale in order to better anticipate storage and market variations. CybeleTech provides a range of software as a service dedicated to each step of the food chain described above. One important aspect of the technology of plant growth modelling is its genericity and adaptability to various species. CybeleTech services are already available for the major field crops in Europe (corn, wheat and rapeseed). The company has the ambition to broaden its field of expertise in the near future to market gardening (tomatoes and cucumber mainly), arboriculture and to forest management.

CybeleTech works with various academic partners including INRA for soil and plant databases, Ecole Centrale Paris for plant growth modelling continuous development and CEA for sensors usage. HPC technologies are nowadays becoming essential for CybeleTech in order to be able to process the accumulating amount of soil data, climatic data and plant data together with numerical simulations.

Breeding a new variety is a long process that requires a decade and thousands of experimental trials in fields so as to select the most robust and efficient traits. Some steps in the process of plant selection could be conducted in-silico to reduce the duration and development cost of a new variety.

The plant growth model used in numerical simulations must be calibrated with plant phenotypes data. To define the optimal experimental protocol to be followed for calibrating the model, the model is run with a genetic algorithm.

In this project, the optimisation of the plant growth model enabled the reduction of its computing time by a factor of five. The performance of the whole application was highly improved by implementing a master-slave approach to the optimisation of the evolutionary algorithm. The performance is a function of the number of protocols and realisations considered.

HPC usage

The SHAPE project has been a valuable opportunity for CybeleTech to gain important skills in HPC technologies thanks to the partnership with experts from CINES, who assisted in the optimization and parallelization of the algorithm. CybeleTech engineers gained skills in the following areas:

- Sequential code optimization with profiling tools
- Efficient MPI implementation of parallel algorithms

Such in-house skills are of high importance to CybeleTech as they have been valued in other projects such as relational database modelling, image processing, data assimilation to plant growth models…

In this regard, CybeleTech continues to work on HPC applications through initiatives accessible to SMEs (Mont Blanc project…) and will carry on using HPC technologies for specific research and development applications in the near future.

Business impact of the SHAPE project

Thanks to the SHAPE project, CybeleTech can provide new breeding methodologies to seed companies. This is offered by CybeleTech as a new service for seed companies that aim at reducing experimental R&D costs in plant breeding trials and validation. The SHAPE project demonstrated that by switching to the new methodology, seed companies are able to save as much as 50% on field trials costs and consequently reduce the time for breeding a new variety. This is a strong commercial argument for making the case for numerical technologies in general to seed companies. However, moving to the new methodologies as suggested by the results of the SHAPE project is a big change of paradigm for seed companies and is a long process that has not demonstrated a quantifiable return for the moment. CybeleTech is in continuous commercial discussion with big seed companies on this topic and the output of the SHAPE project is a valuable result that helps to convince their partners in this regard.

3.3.2 OPTIMAD Engineering (Italy): RAPHI: aerodynamic analysis, design and optimization

PRACE Partner: CINECA

Project Duration: 1st July 2015 – 31st Dec 2015

Project overview

Within RAPHI, a feasibility study of porting COPPA, a kinematic CFD code for rarefied flows, to the Intel Phi architecture has been carried out. The aim was to quantify the necessary investments and to assess potential benefits of adopting this architecture for this computing intensive application. To this scope, the entire code has been compiled on the Intel Phi architecture and the computing core has been refactored to efficiently exploit the Phi architecture. Different performance metrics, indicating scalability, floating point operations and energy consumption, have been monitored and compared to the ones obtain when using the host architecture (Intel Xeon 2630).

HPC usage

OPTIMAD has a deep commitment to HPC since all codes developed by the company use a classical MPI paradigm to express parallelism. Scalability has been demonstrated up to O(104) cores on some of Europe's Tier-0 and Tier-1 systems. For production the company buys cycles at well-established HPC centres, and in particular there is a long-term collaboration with CINECA. The entire technical staff has a profound knowledge in MPI based parallel computing and during the SHAPE project this knowledge has been extended to different paradigms suitable for the Intel Phi architecture, thanks to the specialized workshops organized by CINECA and the interaction during the collaboration.

The SHAPE project RAPHI has been used to quantify hypothetical business benefits (energy consumption, total-cost-of-ownership) in porting KOPPA to the Intel Phi architecture (Knights Corner at the time) with respect to the Intel Xeon processors used in most Tier-1 systems.

Unfortunately, the computed metrics did not indicate a significant business benefit in adopting this technology. However, the company is aware that the employed processor (Knights Corner) was only a transition point within Intel's roadmap for next generation HPC.

Business impact of the SHAPE project

Although the results of the feasibility study are not very encouraging, the SHAPE project has been valuable to Optimad since the project accelerated and catalysed this study significantly.

Without the collaboration with the HPC centre, the company would have had to face noticeable burdens to obtain the same result. The principal benefits are

1. the access to a working and supported system with the Intel Phi architecture;
2. support in porting of the code to Intel Phi and profiling of the code;
3. support in implementing and analysing of performance metrics.

Beyond the result of the feasibility study that has been carried out during SHAPE, the company gained valuable knowledge on the usage of a new (heterogeneous) HPC architecture. While the quantification of the business impact is difficult in this stage, concrete follow up actions have been taken in order to push the Technology Readiness Level of the company's products with respect to forthcoming HPC architectures:

1. Start implementing support for hybrid parallelization strategies in the bitpit library (http://optimad.github.io/bitpit/) developed by Optimad;
2. Collaboration with the INRIA Storm team to deploy efficiently industrial codes on heterogeneous systems.

3.3.3 Open Ocean (France): High Performance Processing Chain - HPPC

PRACE Partner: IDRIS

Project Duration: 16th September 2015 – 29th February 2016

Project overview

Open Ocean is a French SME company, which develops innovative on-line solutions to help plan and manage offshore developments. It has conceived an oceanographic data study tool which computes and formats data (pre-processing and processing) and which provides relevant oceanographic information to industrial marine companies (post-processing) through a web interface. However, the "time-to-solution" of this post-processing step is prohibitively long and hence not compatible with industrial use.

Therefore, the goal of this SHAPE project was to improve post-processing by optimising a parallel Python program developed by Open Ocean that processes and computes statistics (e.g. wind speed) on big datasets. Open Ocean and IDRIS engineers worked together to optimise this program by using resources available in a national supercomputing centre: high performance parallel machine and parallel file system (GPFS, 100 GB/S bandwidth).

HPC usage

The Open Ocean computing infrastructure is composed of a dedicated head node server, which can send tasks to two dedicated compute nodes. The storage is provided through an NFS "share" server.

Open Ocean is renting this infrastructure from the French infrastructure provider online.net. This solution was chosen because it offers sufficient flexibility; the infrastructure can be rented per month for a reasonable cost. However, a relatively good knowledge of system administration is needed to install and maintain the different software layers and Open Ocean had only basic knowledge of parallelization.

At the completion of the SHAPE project, the complete workflow was parallelized; it was also strongly linked to a proprietary solution provided with the scheduler installed on the cluster. Open Ocean's knowledge of parallelization today is more generic and can be used with different software and in different contexts. The Open Ocean workflow is also more portable now.

A parallelization strategy recommendation was also performed, namely with regards to the order of the tasks, which greatly reduced the elapsed time.

Business impact of the SHAPE project

As mentioned earlier, this project provided Open Ocean with a better understanding of the concept of parallelization and today the new optimised workflow is being used.

They also plan to change their current proprietary scheduler for an open source scheduler. This will be low cost and easier to maintain.

A precise description of the workflow has been written for this project. This has helped Open Ocean very much when working with other partners on different R&D projects. However, it is difficult to refer to the web interface.

Some of the recommendations for optimisation arising from this SHAPE work were not always applicable to the SME's infrastructure, and therefore could not be implemented. Also, running the code on a super computer required significant recoding. As such, one conclusion arising from this work would be to recommend revisiting the workflow optimisation in the context of the SME's internal infrastructure.

3.3.4 Hydros Innovation (Switzerland): Automatic Optimal Hull Design by Means of VPP Applications on HPC Platforms

PRACE Partner: CINECA

Project Duration: 1st August 2015 – 1st March 2016

Project overview

Hydros is a Swiss Engineering & Research company founded in 2007 with several patented designs in the field of marine and sailing yachting. Hydros‘ most famous projects are:

- l’Hydroptère: two world speed records in 2009 and absolute sailing speed record (51.36kts);
- Hydros C-Class: first yacht built in TPT, 2nd place at Little Cup 2013;
- HY-X: first hybrid « Fly and Float » motorboat.

In recent years Hydros is exploring new market segments such as yacht and super-yacht hull design. The usage of HPC resources and open-source software can be a valuable tool for massive shape design optimizations as usually performed by Hydros. The main scope of the proposed SHAPE project was therefore to evaluate the feasibility of automatic optimal hull design on an HPC infrastructure and the impact of such a workflow on the day-to-day work of Hydros personnel.

The ability to provide a robust black box for CFD evaluation is a big gain for the complete workflow, nevertheless the methods and processes put in place during this study are still work in progress. Before using it for industrial purposes, the market confidence in the method should be improved showing that the cost of automated optimization is in accordance with the gain in performances.

HPC usage

The company moved to Dubai in January 2017 after being acquired by another company. The new challenges are mainly the commercialization of the foil yachts (http://www.foiler.com) and development of other concepts and yachts. For these reasons, the HPC systems usage has not been pursued at this stage.

Business impact of the SHAPE project

Not available due to the takeover of the company as discussed above.

3.3.5 Vortex Bladeless SL (Spain): Parallel multi-code coupling for Fluid-Structure Interaction in Wind Energy Generation

PRACE Partner: BSC

Project Duration: 1st April 2015 – 1st October 2015

Project overview

Vortex-Bladeless is a Spanish SME whose objective is to develop a new concept of a wind turbine without blades, called Vortex or vorticity wind turbine. This design represents a new paradigm in wind energy and aims to eliminate or reduce many of the existing problems in conventional generators. Due to the significant innovation in the project concept, its scope is different from conventional wind turbines. It is particularly suitable for offshore configuration and could be exploited in conventional wind farms and in environments usually closed to existing ones due to the presence of high intensity winds.

The objective of this SHAPE project was to develop the necessary tools to simulate Fluid-Structure Interaction (FSI) problems and to reproduce the experimental results for scaled models of the Vortex-Bladeless device. In order to do this the Alya code, developed at the Barcelona Supercomputing Centre, was adapted to perform the Fluid-Structure Interaction (FSI) problem simulation.

HPC usage

Following the SHAPE project, the SME has continued the collaboration with BSC in the framework of the H2020 SME instrument program within the project "New Concept of Affordable Wind Energy Generators without Blades – VORTEX". In this new phase, the HPC resources are being used to investigate the reach of the technology and to help in the design and development of the Vortex device. The most demanding calculations are performed in the MareNostrum supercomputer but also numerical test are being carried out by the SME on a system of their own (typical runs of 96 CPUs).

Business impact of the SHAPE project

Revenue from sales is yet to be obtained by Vortex Bladeless as the product is still at the design and development stage. However, thanks to the relationship between BSC and the project, through the H2020 SME INSTRUMENT program and the SHAPE work, very interesting results are being obtained for the future business:

1. Computational validation of wind tunnel experiments is now possible.
2. Studies can be performed on the geometry of the device, looking for the best aerodynamic performance and the influence of some detrimental effects.
3. At the present time and in this particular area, Vortex is focusing on finding the answer to what is the largest device (maximum nominal power) that could be built without entering in the critical Reynolds area (high aerodynamic performance zone).

Despite not having sales revenue yet, great advances in the development of the technology have enabled capital investments in the company via two new significant partners. This has provided Vortex with enough financial resources to recruit the technical staff they need to carry out their planned tasks.

Great support for the project has also been received from all over the world (awards, donations, altruistic collaborations …), with more than tens of thousands emails from people interested in the project and the future products.

3.3.6 Design Methods (Italy): Coupled sail and appendage design method for multihull based on numerical optimization

PRACE Partner: CINECA

Project Duration: 1st April 2015 – 1st October 2015

Project overview

Design Methods is an engineering firm with fifteen years of experience in the aerospace field. The mission of Design Methods is to provide multidisciplinary engineering consulting and design services to industries and design teams supporting them with highly specialized competences on aerodynamic design, CAE analysis, software development, CAD modelling, numerical optimization environment and customized design tools development. They operate in aerospace, automotive and marine fields. The main scope of the project was to evaluate the feasibility of a numerical optimization workflow for sailing boats sail plans and appendages.

HPC usage and Business impact of the SHAPE project

Despite the wishes of Ubaldo Cella, owner of Design Methods, he has been unable to exploit the results of the SHAPE project yet, due to other ongoing projects and completion of his doctorate. However, there are plans to utilise the SHAPE work in the future.

3.3.7 Ingenieurbüro Tobias Loose (Germany): HPC Welding

PRACE Partner: HLRS

Project Duration: 1st August 2015 – 31st March 2016

Project overview

As an integral part of the PRACE SHAPE project “HPC Welding” the parallel solvers of LS-DYNA were used by Ingenieurbüro Tobias Loose (ITL) to perform a welding analysis on the Cray XC40 “Hazel Hen” at the High Performance Computing Centre Stuttgart (HLRS).

A variety of test cases relevant for industrial applications have been set up with DynaWeld, a welding and heat treatment pre-processor for LS-DYNA, and run on different numbers of compute cores. The explicit mechanical solver was tested on up to 4080 cores with significant scaling. As far as the project partners are aware, it was the first time that a welding simulation with the LS-DYNA explicit solver was executed on 4080 cores.

HPC usage

The project clarified how HPC can be used for welding analysis, which welding processes, welding tasks, modelling methods and analysis types are applicable to HPC and how much effort is necessary. The overall effort for welding analysis on HPC is now much better known with the help of this SHAPE project, leading to the ability of gaining a more accurate cost estimate of welding consulting jobs. This is a commercial benefit for ITL. The project provides a good basis for further investigations in high performance computing for welding structure analysis.

After the end of the project, during the year 2016, the results were published at several conferences and in the journal “Schweiß- und Prüftechnik” (Vienna, 2016, pp. 204-206: “Höchstleistungsrechnung für die Schweißsimulation mit DynaWeld und den parallelisierten LS-DYNA Solvern”.

Further publications in other journals are under preparation.

Business impact of the SHAPE project

Based on the experiences with this project ITL established the DynaWeld GmbH & Co. KG with an increase of manpower to continue in the business of manufacturing simulation welding and heat treatment. DynaWeld and ITL continue to work using HPC and are expecting orders of HPC related projects. The marketing campaign for this is running.

3.3.8 Ergolines SRL (Italy): HPC-based Design of a Novel Electromagnetic Stirrer for Steel Segment Casting

PRACE Partner: ITU

Project Duration: 12th May 2015 – 19th April 2016

Project overview

Ergolines is an Italian SME dedicated to the development of custom technologies for steel continuous casting, its core business being the design and engineering of electromagnetic stirrers (EMS). These are electrical machines dedicated to the contactless stirring of liquid metal in industrial processes through eddy current induction, leading to key process benefits. The project applied HPC-based numerical simulations to the analysis of the magneto hydrodynamics of liquid steel in an industrial furnace.

HPC usage

Ergolines is not using HPC resources at the moment.

Business impact of the SHAPE project

Participation in the SHAPE project was definitely beneficial for Ergolines, since it enabled the company to considerably speed up computational times in the performance analysis of EMS applied to industrial furnaces.

3.3.9 Summary of follow-up responses

The follow-up feedback from the SMEs has generally been very positive with regards to SHAPE. Unfortunately, this has mainly been provided in a qualitative rather than quantitative sense, i.e. measurable against the suggested Key Performance Indicators. The key points from the SMEs questioned are noted below:

KPI: New product offers

Ingenieurbüro Tobias Loose: are promoting and expecting orders for HPC related projects.

KPI: New service offers

Cybeletech can offer savings up to 50% on field trial costs of breeding new varieties – SHAPE results provide convincing argument to their business partners.

KPI: Enhanced software features

Ergolines SRL: considerable speed up of their performance analysis software.

KPI: Staff recruitment

Vortex Bladeless: have taken on more staff following advances in the development of the technology, partially enabled by their SHAPE work;

Ingenieurbüro Tobias Loose: have set up a new organisation based on their experiences with SHAPE, which has meant new staff employed.

Other:

Optimad: made much quicker evaluation of new technologies and gained valuable experience in the usage of heterogeneous architectures, which has fed into their ongoing HPC strategy;

Open Ocean: gained knowledge of parallelisation and enhanced their workflow, which has helped their relationships with partners;

Hydros Innovation: company taken over so HPC usage no longer a priority;

Design Methods: no exploitation yet but still plan to make use of the SHAPE results in the future.

3.4 SHAPE second and third call: Project Summaries

This section provides summaries of the three projects still ongoing from the second call of the SHAPE programme, and the eight projects from the third call. For each of these there is a brief overview describing the problem to be solved, the activity undertaken, how PRACE was involved, the benefit to the SMEs, and finally the lessons learned for the further development of the SHAPE programme itself. The lessons learned are discussed further in Section 3.4.12.

Note that each SHAPE project is also producing a technical white paper that will cover the activities and results of the projects in greater detail than presented here. The intention of this section is to give a flavour of the broad range of projects and the diversity of the subject areas, along with summarising the benefits of the SHAPE programme to the SMEs.

3.4.1 WB-Sails Ltd. (Finland): Modelling sail boat performance in HPC

Overview

Project partners:

- SME: WB-Sails Ltd.
- Technical partner: Next Limit Dynamics (NLD)
- SHAPE contact: CSC – IT Centre for Science Ltd.

In this project WB-Sails wanted to utilize HPC resources in CFD and simulations of sails and sailboats using a particle based CFD code XFlow from NLD. Coupled FSI (fluid-structure interaction) simulations using a FEM solver were also planned. The goal was to be able to perform more complex simulations in a shorter wall-clock time than possible with WB's in-house desktop systems.

Activity performed

XFlow was installed on CSC's Taito cluster system and numerous test runs have been performed, including scalability testing, implementing the XFlow GUI on the Taito-GPU system, and robustness testing of XFlow's various solvers, in particular for external aerodynamics as well as multiphase free surface problems.

During the testing it was found that XFlow does not support adaptive mesh refinement when MPI is used for parallelization. This leads to wasting computing resources and reduces the parallel speedup. Also, using more than 192 cores caused the software to crash.

The biggest problem was that CSC was not able to provide free core hours for the project due to its ownership policy. Several alternative funding scenarios were investigated to try and alleviate this, but none were successful.

Thus the more computing intensive FSI simulations had to be abandoned. However, a number of test runs with and without free surface were performed using different meshes and various core counts. Additionally, valuable experience and knowledge about using HPC resources was gathered.

A more detailed report can be found in the white paper "WB-Sails: Modelling sail-boat performance in HPC", under review at time of writing.

PRACE cooperation

CSC supported WB in the project by setting up the simulation software and environment at CSC. Substantial effort was used in helping WB to use CSC's HPC environment and especially the batch queueing system, GPUs for visualization and so on. WB considered CSC's support during the project outstanding.

Computing time from CSC's Taito cluster including GPUs was allocated for testing purposes.

Benefits for SME

Moving from a desktop to a HPC environment proved to be more complicated than expected. In the course of the project WB-Sails gained experience and skills in using a HPC system and now has clear understanding of the benefits. Also limitations of the software were discovered.

Although actual production runs could not be performed, even the test runs provided valuable data to WB.

The experience gained in the project allowed WB and CSC to jointly submit a proposal to another EU funded HPC collaboration project.

Lessons learned

Generally the SHAPE program works very well and offers SMEs an easy way to start using HPC resources. The support offered by PRACE is crucial in the transition from the local system, e.g. a desktop, to a HPC environment.

However, at least for CSC, it would be important that the computing resources needed in the SHAPE projects could be funded by EU. Due to the Finnish national policies they cannot provide resources to industry without compensation.

3.4.2 Principia (France) – HPC for Hydrodynamics database creation

Overview

Principia is a scientific engineering company that performs engineering studies for large industrial companies and develops and commercialises added-value numerical software solutions.

The main goal of this project is to optimize one Principia code, Diodore, on HPC infrastructure in order to improve the Deepline HPC product.

Principia and CINES are working together on this topic, with the aim of transferring optimization and parallelization skills so that Principia staff can reuse the methodology on their other HPC codes.

Activity performed

The beginning of the project was delayed due to security aspects and Principia's staff availability.

CINES profiled the code with Intel VTune and found hot spots on lines of code concerning a temporary file, mathematic functions (*exp(), sqrt()*) and BLAS calls for solving a linear system of equations.

CINES compiled the sequential FORTRAN code with additional runtime checking and fixed uninitialized variables and other potential bugs that could lead to errors if their default value was not zero.

There were numerical instabilities when running the test cases (e.g. plate in water). The set up involved a fake MPI execution with two processes. One rank is the unchanged binary code and the other one contains the modified code. At the end of the simulation the arrays are compared by the zero rank process, but despite various efforts these were not equivalent.

NOTE: Due to these issues with the work a white paper will not be produced, rather, an internal technical report will be provided to PRACE, due April 2017.

Benefits for SME

The SME now has some improved checks and processes to validate its code.

Lessons learned

CINES learnt that error propagation is hard to follow and subsequently to correct in an industrial code.

3.4.3 Algo'Tech (France) – Electrical design enabled by high performance linear algebra

Overview

Project partners:

- Algo'Tech contact: Jacques Peré-Lapernes.
- INRIA technical partners: Matias Hastaran (INRIA Bordeaux-Sud Ouest) and Pierre Ramet (Bordeaux University and INRIA Bordeaux-Sud Ouest).

Manufacturers increasingly use electrical and electronic components in all kinds of products to provide new functions, improve performance and reduce the operational and development costs. The counterpart to this is the introduction of new constraints such as the electromagnetic vulnerability where the equipment functioning can be disturbed by electromagnetic effects. The perturbation can happen in various conditions:

- Inside an electrical harness between power supply wire and command wire;
- Between electrical harnesses;
- Or because of external sources (electrical equipment, engines, Wi-Fi etc.) or natural ones like lightning.

This raises concerns for machine design, robotic assembly lines, devices for commercial buildings and public facilities, embedded systems (aircraft, trains, boats, tramways, drones), and in general, all automatic devices and electrical systems. To reduce the time to market, manufacturers need tools able to:

- Detect detrimental electromagnetic effects as soon as possible: before prototyping, during the design phase;
- Avoid cable shielding when it is not necessary: lighter equipment, less manufacture operations (such as welding).

An electromagnetic simulator, integrated into an electrical CAD software suite, used during the design phase, on a simple PC connected to HPC, would allow the manufacturers to adapt the design to electromagnetic constraints before building the first prototype. Currently, they generally identify such problems during the testing phase, or even on the customer site during installation, with all the consequences that entails.

Activity performed

For conducting electromagnetic vulnerability studies, Algo'Tech has developed a numerical simulator, for which a major computational bottleneck was the solution of a large sparse system of linear equations. To overcome this problem, the aim of this SHAPE project was to use a state-of-the art parallel sparse direct solver, namely PaStiX [21], which has been developed at INRIA Bordeaux-Sud Ouest for several years, into the Algo'Tech simulator. The features of the PaStiX solver should reduce both the memory footprint and the time to solution on computers ranging from multicore desktop up to large-scale platforms.

In short, the Algo'Tech software provides an electrical model for the wire harness, which is translated into a sparse system of linear equations written in a matrix format, Ax =b, where the matrix size depends on:

- The number of wires in the electrical harness;
- The length of the wires combined to the frequency range selected for the simulation

The white paper will report on the computational performance benefits for the Algo'Tech simulator observed on a set of currently considered large calculations displayed in Table 5 for computations performed on SME affordable multi-core platforms.

Table 5: Number and length of wires and their associated matrices

Number of Wires	Length (meters)	Order of the square matrix A	Number of non-zero terms into A
200	1	40 209	164 568
	10	398 409	1 636 368
	20	796 608	3 272 519
	50	1 990 608	8 178 519
500	1	101 038	417 434
	10	1 001 038	4 150 634
	12	1 101 538	4 567 508
	20	2 220 076	8 301 268
1000	1	202 153	835 133
	10	2 220 300	8 351 330
1500	1	303 440	1 252 850
	10	3 344 000	12 530 000
2000	1	404 434	1 670 213
	10	4 443 400	16 702 130

PRACE cooperation

The PRACE technical partners from INRIA have closely interacted with the Algo'Tech contact to facilitate the integration of the PaStiX solver into the SME simulator. They also provided advice to optimize the performance of this new solver in order to reduce as much as possible the wall clock time of typical electromagnetic vulnerability studies.

Benefits for SME

Thanks to this experiment and its collaborative work with INRIA, Algo'Tech has developed a solution to have a simulator working seamlessly on a PC or HPC systems. Algo'Tech tested the full chain, designing through an electrical CAD suite with real wiring diagrams and testing through its simulator the electromagnetic vulnerability. The simulations have been run both on PC and HPC systems, with various numbers of cores. Also it has been tested if it is possible to do this transparently for the users.

In terms of a business plan, this project has allowed Algo'Tech to reach 2 main objectives:

- Consolidate Algo'Tech's business case and the feasibility of its preferred business model (SaaS);
- Understand better what kind of issues the company will face to propose this solution operationally on the market – meaning by that, how to be able to propose at a reasonable price an "on demand" HPC service.

Algo'Tech has the ambition to have 10% of its business driven directly by the simulation in mid-term future and considers that it will help them significantly on 15 to 20% of their current business (CAD software) by creating an important competitive advantage. Algo'Tech also expects to get business with large groups currently out of "arm's reach", thanks to this simulation capability.

Lessons learned

Algo'Tech used to develop their whole source code themselves. This SHAPE project has allowed them to envisage using other solutions, such as INRIA libraries for HPC. This combination of solutions helps to provide better services to Algo'Tech's customers, saving time and increasing design efficiency. HPC allows saving hours of runtime but this gain must not be lost waiting for days the schedule of the job when using external computing resources.

3.4.4 ACOBIOM (France): MaRS Matrix of RNA-Seq

Overview

ACOBIOM has more than 18 years of experience of working with computational experts and biological specialists. ACOBIOM's bioinformatic platform has been dedicated to genomic data management and it plays a leading role in analysing the results generated by its new high throughput platform and integrating all these data into a network of interacting routes. This provides the biotechnology industry and clinicians with accurate targets or correct information for curing patients.

Since the early days of the Human Genome Project, data management has been recognized as a key challenge for modern molecular biology research. Recent years have seen a dramatic increase in the amount of genomic and transcriptomic data produced by typical projects in this domain. The research program, in association with GENCI/CINES (Montpellier, France), focused on transcriptomic data and the RNAseq approach (study of RNA molecules). The RNAseq approach is used in a wide variety of applications. These include identifying disease-related genes, analysing the effects of drugs on tissues, and providing insight into disease pathways. The RNAseq is widely used to characterise gene expression patterns associated with tumour formation. Since the RNAseq provides absolute values and does not require any calibration with arbitrary standards, results can be compared at any time with other data, even that raised by independent laboratories. Once collected, these data can be digitalised and then easily and reliably compared in silico with the growing library of RNAseq databases generated for normal and pathological situations in other laboratories around the world (Human: ~27000 libraries; Average size of a library: 1.7GB. Total size: 120TB).

This project aims at the development of the Generic platform called MARS (Matrix of RNA-Seq) with an innovative method and software to analyse, integrate and contextualize large-scale biological data in the fields of Human Health. MARS is made timely by the exponential increase in the throughput of molecular (Omics) approaches to cover the unmet needs in the specific field of Health.

Activity Performed

Many challenges were encountered and solved during the project: the download of a very large amount of data, the installation and the settings of the various software in the cluster, the optimization of the compiler and the packages, and the constraints in jobs duration and jobs limitations.

Installation of each code of the toolchain (Trimmomatic, bowtie, picard, samtools and htseq) along with all the third party requirements (java, boost) was done and used through SLURM jobs. It represented a huge amount of data computed and required the use of high performance computing on the cluster Occigen (GENCI/CINES): 120 TB of compressed data downloaded and 1.2 million core hours consumed thanks to the extra allocation given by the CINES director.

PRACE cooperation

ACOBIOM will give a presentation of the MaRS project during the "Industrial Parallel Track 3" which will take place during PRACEdays17.

Benefits for SME

ACOBIOM has been able to reach the goal of this project that is, knowing how to adapt and monitor the work to be done on a remote supercomputing facility.

Lessons learned

ACOBIOM learnt how to follow, thanks to python scripts, the workflow of thousands of jobs crunching the 27000 genomic libraries, and how to re-submit the failed ones that are statistically inevitable.

3.4.5 Airinnova (Sweden): High level optimization in aerodynamic design

Overview

Airinnova AB is a start-up company providing and developing advanced computational technology solutions for cutting-edge aircraft preliminary design, computational aerodynamics, and multi-disciplinary optimization. The PDC-HPC Centre at SNIC-KTH volunteered as PRACE contact partner in their project focusing on aerodynamic shape optimization for a reference aircraft model, namely the common research model (CRM) aircraft released by NASA.

Shape optimization is an important task in aircraft design. The ultimate goal is to design a lighter, greener and quieter aeroplane by reducing drag, especially at high speeds.

In order to reach solutions that are not only close to the baseline shape, methods involving computationally expensive fluid dynamics analysis have to be applied to reveal many more candidate shapes. Accurate Computational Fluid Dynamics (CFD) runs, for example the Reynolds-average Navier-Stokes (RANS) solver with turbulence models, must be parallelized to reduce the total run-time.

Activity performed

During the project the performance using different proposed CFD solvers to carry out high-fidelity RANS aerodynamic shape optimization was analysed and benchmarked. The SME needed essential benchmarking help and assistance with porting (namely to deploy, run in an automatized way and monitor) the CFD solvers on a PRACE Tier-0 machine.

For a more detailed reporting of the results see the white paper ("Automation of High Fidelity CFD Analysis in Aerodynamic Design", currently under development) and also the final AIAA conference paper.

PRACE cooperation

PRACE helped with clarifying the scope of the project and documented it in a detailed, agreed work plan. PRACE also helped with the application for Tier-0 Preparatory Access resources. Besides that regular (bi-weekly) meetings were held to discuss the progress and any extraordinary circumstances. In August 2016, also a midterm report was created to clearly document the progress.

Beyond the scope agreed-upon in the work plan, the SHAPE project was also presented in the local Newsletter, a common abstract to prestigious 2017 AIAA AVIATION Forum was prepared and accepted, and additional computing time on the local PRACE Tier-1 resource Beskow was granted.

Benefits for Airinnova

PRACE tried to act as a sounding board and give new inspirations: for some parties getting different viewpoints and the possibility to discuss the project was already of great benefit.

PRACE could help with a fresh outside view in reformulating the scope of the project to make it easier to be understandable by a broader range of people.

The current SHAPE project is already a step towards making all the core processes that require special skills - like creating the mesh and execute the high-fidelity CFD analysis code – automatic, and thereby helping a wider field of engineers to design the aircraft in a more efficient and simpler way.

The SME has subsequently applied for a Preparatory Access Type-D project, which will follow-up on the work done during this SHAPE project and take advantage of the produced optimization results and the existing scripts. The final technical goal is a flexible integration of the whole automation process into a MDO (Multi-Disciplinary Optimization) design environment. The final research goal is still the wing shape optimization for drag reduction.

Lessons learned

The fact that two main representatives from the SME and PRACE come from a very similar field and were having the same mother tongue helped tremendously in speeding up the project and getting to terms right from the start. As such, the main challenge within this set-up was to put the plan on paper so everybody else would also understand the detailed scope and come with ideas. Fortunately, the partners were able to create a common open communication atmosphere, with everyone being engaged and happy to explain things all over, so that issue was avoided.

Concerning the general SHAPE process, it was not clear that the SHAPE application was not automatically coupled to a Preparatory Access Application. The SME should not have to write two applications for essentially the same thing, and a SHAPE project without computing time would not make much sense. Additional computing time from local KTH resources were used to bridge over the times without Tier-0 access for the SHAPE project.

3.4.6 Creo Dynamics (Sweden): Large scale aero-acoustic applications using open-source CFD

Overview

Creo Dynamics is an engineering company with core competence in fluid dynamics, acoustics and structural dynamics with sound experience in participating in national and international research programmes focussing on the development of new emerging technologies for the automotive or aerospace industry. During this SHAPE project simulation processes, based entirely on open-source software, were put together to demonstrate an efficient and robust workflow implementing real life Computational Fluid Dynamics (CFD) applications in parallel at large scale that can compete with current industry best practices at a considerably lower cost.

The PDC-HPC Centre at SNIC-KTH volunteered as PRACE contact partner in this promising project aiming at encouraging new users to adopt and invest in HPC and open-source software.

Activity performed

During the project a relevant test case of industrial relevance (the general truck model) was chosen.

Benchmarks were performed for the snappyHexMesh mesh generator part of OpenFOAM [22] for OpenFOAM 3.0+, 4.0 and 1606+.

The memory usage of snappyHexMesh during all steps of meshing was monitored and analysed to find critical bottlenecks and inefficiencies in the workflow, the overall meshing parallel scalability was analysed and a semi-automated process for parallel meshing, solving and post-processing of large CFD cases was demonstrated.

Effort was also put into memory consumption tests to locate, replicate and really identify and describe the leak problem in detail in every step of the calculation as part of task 4 in the workplan.

A complete description of the findings is given in the corresponding white paper ("Overcoming challenges in computational aero-acoustic processes using scale resolving open source CFD", currently under development).

PRACE cooperation

PRACE, together with the SME, agreed on a detailed work plan. PRACE also helped with the application for Tier-0 Preparatory Access resources. Besides that, regular (bi-weekly) meetings often over video were held to discuss the progress and any extraordinary circumstances. Regular communication was done over email.

Beyond the scope agreed-upon in the work plan, the SHAPE project was also presented in the local newsletter and additional computing time on the local PRACE Tier-1 resource Beskow was granted.

Benefits for Creo Dynamics

Within the PRACE SHAPE project, Creo Dynamics got access to world-class HPC resources at BSC and PDC-HPC where they can perform benchmarking tests, and share and discuss their experiences with other experts in the HPC area.

Creo Dynamics has its own small HPC resource. However, the project-driven timescales in industry usually have requirements to run larger simulations in a short time. The SHAPE project illustrates that PDC or other PRACE systems can provide additional computational resources for industrial users.

Moreover, sometimes some features cannot be exploited when only using small clusters. For example, it was thought that advantage could be taken of the distributed STereoLitography (STL) not only on the memory consumption but also on the total run-time. However, after some benchmarking tests, we found that the runtime performance of distributed STL meshing is too poor compared to the standard approach.

PRACE provided an expert in using the Allinea development debugging tool DDT to find the memory leak.

In the end the distributed STL worked, with the resulting mesh being valid and correct. The bugs found were rather documentation missing on limitations and requirements.

Lessons learned

"Open source is not free."

To choose the actual geometry of the test case took longer than expected, as it was desirable to avoid the hassle of signing No-Disclosure Agreements (everyone was ready to sign, but the additional complication seemed unnecessary), and therefore in the end only a general model could be used, rather than a concrete one that may have been more impressive. In hindsight, at least part of that discussion with the third party (in this case Scania) would have been useful to have happened earlier.

Further third party dependencies were introduced when finding a crucial error for our use-case within Allinea's DDT and performing additional testing for eventual bug reports towards OpenFOAM and the distributed STL support on bug/feature experiences.

At one time even a security issue was interfering with running the debugger.

As SHAPE were not a paying customer, this often led to low priority or long reaction times. OpenFOAM required payment to fix the error at an estimated worth of 5 working days.

Concerning the general SHAPE PRACE process, it was not clear that the SHAPE application was not automatically coupled to a Preparatory Access Application. The SME should not have to write two applications for essentially the same thing, and a SHAPE project without computing time would not make much sense.

The account creation for PRACE SHAPE users is not fully automatized yet. It was not clear whether PRACE SHAPE users should be put in the LDAP system and handled in the same way as PRACE DECI users. For the access to the Tier-0 system papers had to be signed on paper and sent via snail mail twice since the first set-up got lost with the account being blocked in the meantime.

3.4.7 AmpliSIM (France): DemocraSIM (DEMOCRatic Air quality SIMulation)

Overview

Project partners:

- Olivier OLDRINI, AmpliSIM, olivier@amplisim.com
- Sylvie PERDRIEL, AmpliSIM, sylvie@amplisim.com
- Isabelle DUPAYS, IDRIS, isabelle.dupays@idris.fr
- Pierre-François LAVALLÉE, IDRIS, pierre-francois.lavallee@idris.fr
- Sylvie THEROND, IDRIS, sylvie.therond@idris.fr

AmpliSIM is a French start-up based in Paris, created at the beginning of 2015. AmpliSIM provides an on-demand browser-based Numerical Simulation Service for business companies and public authorities, enabling urgent computing for operational air quality.

Their market is impact assessment for air quality and the management of industrial hazards. They provide private users or public authorities with the capability to perform seamless numerical simulations needed for the assessment of air quality regarding the impact of industrial plants.

The DemocraSIM project aims to bring air quality simulation to the general public. When the Fukushima nuclear disaster occurred, it was possible to display crowd-sourced data of radioactivity measurements on a single web map (see the map [23] and the history [24]).

Why are people able to do their own radioactivity measurements and share them, while they cannot simulate an event and get an idea of their own potential exposure?

The DemocraSIM SHAPE project will allow AmpliSIM to overcome technological barriers through collaboration with IDRIS engineers in the use of urgent computing, advanced visualisation and data analytics, thereby enabling them to launch an industrial solution on a private HPC cloud (such as the Fortissimo [25] platform). It will also allow IDRIS HPC experts to address these issues on “real cases” with potential re-use in academic cases of urgent computing.

Activity performed

The first part of the activity was the deployment of all the software and libraries that the project needed on the IDRIS supercomputers:

- Saturne (version 4.0.5 with hdf5 and MED), using the Intel Compiler and MPI Intel library,
- OpenFOAM (3.0.1), using the Intel Compiler and MPI Intel library,
- GDAL (2.0.1),
- QT (5.0.6),
- Proj4 (4.8.0),
- Mercurial (4.1).

Compiling and executing the FORTRAN/C++ codes with other compilers such as Intel instead of GNU highlighted some problems. These were corrected (e.g. adding BIND(C) for the subroutine declarations and not defining a specific subroutine name for a C++ call). This solution worked with all compilers.

Python profiling was also performed in order to resolve any possible bottlenecks in the application.

In order to test the simulation with a portal it was necessary to create a small test case with a multi-step job in 4 parts:

- Sequential data transfer on the supercomputer,
- Parallel software execution (OpenFoam, Saturne, …),
- Sequential data post-processing,
- Transfer of the data on the local machine.

Currently the project is in its final phase, integration with the portal.

PRACE cooperation

The collaboration with AmpliSIM has been very successful with monthly meetings held at IDRIS.

At the last face to face meeting of WP7 in Nice Sophia Antipolis, December 6th-7th, a presentation was given by Alexandra Fresneau from AmpliSIM.

The project will also apply for the next PRACE Preparatory Access type D in order to start a code optimisation process on a PRACE Tier-1 system.

Benefits for SME

Benefits for AmpliSIM modelling tools

The post processing of the numerical models binary outputs for visualization is an actual challenge for AmpliSIM's web service. The constraints are to be able to visualize (rapidly) large amounts of data. To overcome both the size of the data and the rapidity, AmpliSIM have developed their proprietary post-processing suite.

This suite, called bin2tile, relies on tiling process, like that used in a web geographical information system: geographical data are defined at different zoom level and sliced in tiles. Then these tiles are loaded for visualization according to the location and the zoom level being explored by the user. It enables AmpliSIM to visualize very large computational domains.

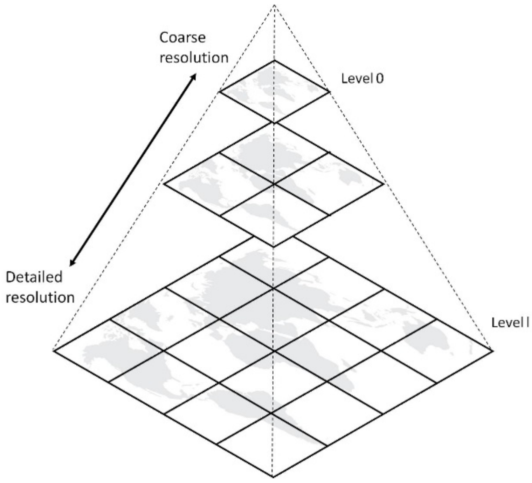

Figure 17: Example of tiling and pyramidal process for data visualization

Bin2tile uses parallel programming via MPI to enable AmpliSIM to rapidly produce images from binary model outputs.

Thanks to the IDRIS team, it was possible to optimize the bin2tile suite, including better compilation options, and solve parallel paradigm issues related to file access and directory creation.

AmpliSIM also received help in successfully increasing the portability of the bin2tile suite.

Benefits for AmpliSIM connection backend

The workflow of AmpliSIM service consists of several steps:

- Case setup for the numerical model,
- Uploading of input data on the IDRIS cluster,
- Launching of the simulation,
- Retrieval of post processed data from the IDRIS cluster.

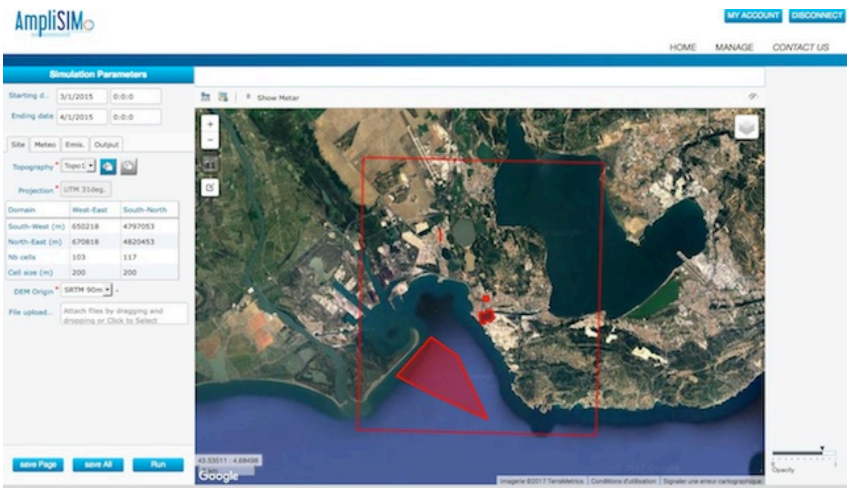

Figure 18: Example of case setup using AmpliSIM web service

Thanks to the DemocraSIM project and connection on IDRIS infrastructure, AmpliSIM could upgrade the connection backend software. This backend allows the web portal to drop input files on the cluster, launch the simulations and retrieve the outputs.

AmpliSIM connection backend (ACB) relied on ssh to connect to the IDRIS cluster, and as such AmpliSIM had to improve the ACB with regards to modularity and security to meet IDRIS standards.

Regarding modularity, the ACB launching scripts had to be generalised and the IBM LoadLeveler (LL) process was introduced. In addition to being able to hook launching steps thanks to LL steps properties, AmpliSIM is now able to connect to other LL type clusters.

IDRIS security constraints enforced us to strengthen SSH options used by the ACB.

Benefits for AmpliSIM web portal

AmpliSIM adapted their platform to meet the project constraints: in order to allow for scaling of different services provided by the platform, they had to gain modularity in the component. In order to do so, Docker containers were introduced to split the portal into different components that can then be scaled independently.

Security of the web services has hence to be defined in a manner that allows for a Single Sign On process: a user of the portal must not have to identify themselves multiple times. In order to do so, a Central Access Service (CAS) was introduced within the portal. The user, once signed in, has the ability to access the different web services provided by the portal without any additional authentication, while at the same time ensuring protected access to the data.

Follow-up

In the following months the AmpliSIM team will focus on:

- The capability to retrieve information regarding the calculation context of the model run (like processors layout, time taken for each step...) from the cluster,
- The implementation of a machine learning algorithm to implement predictors of simulation runtimes and job consistency checks.

As these tasks will need longer time than that remaining inside the SHAPE framework, AmpliSIM are considering applying to suitable candidates in the PRACE preparatory access programmes to continue the work.

3.4.8 Anemos srl (Italy): Modeling SUNSTAR Simulation of UNSteady Turbulent flows for the AeRospace industry

Overview

Anemos SRL is a recently founded highly innovative start-up company established to leverage on the longstanding and internationally recognized competence of the CFD Group of the Department of Mechanical and Aerospace Engineering (DIMA), to set up services appealing to the Aerospace Industry. CINECA was the PRACE centre supporting the SME, in collaboration with CNAM Paris.

The SUNSTAR project aimed at improving the robustness and testing the predictive accuracy of an in-house CFD software for the simulation and the analysis of turbulent high-speed flows relevant for the aerospace industry, exploiting the HPC premises provided by PRACE. The software kernel is a compressible parallel flow solver, incorporating state-of-the-art numerical methods and advanced features for the direct numerical simulation of supersonic flows.

Activity performed

The activities in the SHAPE project are aimed at improving the performance and the parallel scalability of an existing in-house solver, which incorporates state-of-the-art numerical methods and advanced features for the simulation of supersonic flows in complex configurations, including high-order shock-capturing capabilities.

For the purposes of the project (i.e. optimization and improvement of parallel efficiency, and code validation) the kernel version of the solver has been considered (i.e. no complex immersed boundaries, no shocks) so as to focus on improving the key elements.

PRACE cooperation

The optimization activities have been carried out in close contact with CINECA. In particular, Dr. Vittorio Ruggiero provided specialist assistance for code implementation and performance improvement.

One of the main tasks was to optimize the solver in order to enhance data vectorization, to exploit the novel Intel architectures, in particular Broadwell and Knights Landing, made available through the CINECA system MARCONI, equipped with 256 (AVX) and 512 (MIC) bit registers respectively.

Benefits for SME

The outcome of the project has been highly satisfactory for the SME, having led to:

- successful implementation of the baseline software platform on modern parallel computer architectures, made available by PRACE;
- improvement of the code vector performance;
- improvement of the parallel scalability;
- code validation on large-scale test cases.

The next step to be undertaken will be exploiting the implemented innovation to reach the market. The main idea for that is to establish a cloud-based, pay-per-use CFD service for companies operating in the aerospace industry. In summary, the pay-per-use model to be implemented will also have the potential to create invaluable business benefits in terms of:

- Substantial reduction of the simulation cost and time, allowed by smart integration of hardware, expertise, applications and visualization tools in a “one-stop-shop”;
- Greater efficiency in the advanced design stage thanks to the access to high-fidelity 3D, unsteady computations;

- The possibility to validate and improve simplified predictive tools, typically used for preliminary design.

Lessons learned

The SME expressed their full satisfaction with the activity.

From the PRACE centre point of view, two issues may be considered as feedback for future SHAPE calls:

- PRACE should elaborate a Terms and Conditions document which should be signed by the selected company detailing in particular how IP-related issues will be handled;
- The concerned PRACE partner should provide a specific amount of CPU-hours agreed upon with PRACE as in-kind contribution, available at the start of the project. It is felt the activities of a SHAPE project are not adequately supported by having to use the mechanism of applying to a further call for machine time, even on an ad-hoc basis, which has often timeframes incompatible with the project and gives the SME an extra bureaucratic burden.

3.4.9 BAC Engineering Consultancy Group (Spain): Numerical simulation of accidental fires with a spillage of oil in large buildings

Overview

BAC Engineering Consultancy Group is an engineering consultancy SME. It develops four main activities:

- Civil Engineering, which includes activities in road works, rail, airports, waterworks and river hydrology, drainage and water treatment, and environment;
- Structural engineering, focusing on the field of structures. Activities include new works, rehabilitation, large equipment, and building height;
- Quality control, specializing in areas such as quality control in building and civil engineering, inspections and testing, such as non-destructive testing, auscultation and control, naval inspection (thickness measurement) and inspection of tanks (Oil & Gas);
- R&D on prefabrication of concrete and metal structures, software development, studies and research prototypes in Fire Engineering and Wind Engineering.

The objective of the present project is to develop a fire engineering analysis (Performance-Based design) of a steel structure building that belongs to the ITER (International Thermonuclear Experimental Reactor) industrial complex in Cadarache (France), which is devoted to research in the field of Nuclear Fusion. Numerical simulations of different accidental fires that can take place inside of the building have been carried out.

This project was realised with the technical support of the Barcelona Supercomputing Centre.

Activity performed

The SME performed benchmarking and scalability tests of the Fire Dynamics Simulator (FDS) code on BSC's MareNostrum machine up to 2048 cores, and then the results were compared with its in-house cluster. The company verified that BAC's infrastructure has an important weakness between node connectivity and system configuration that limit the scalability of the code. The BSC team supported them to improve parts of the local installation, helping to optimize their current infrastructure and detect the bottlenecks for future clusters and installations.

PRACE cooperation

The Partnership for Advanced Computing in Europe (PRACE), via SHAPE, provided the expert support to install the code, adapt their scripts for running the system, helped in their local installation and fixing technical problems in their use of Tier-0 systems and also within their own company cluster. Also, PRACE Preparatory Access Type B provided the required computational time necessary to perform these initial simulations.

Benefits for SME

The numerical simulations have allowed understanding possible fire situations and their effects on the buildings in one of the most representative international projects and non-pollutant energy icon for the future generations.

The participation in the SHAPE program has made it possible to test the scalability of the code in an HPC environment and to develop the capacities inside the company to use HPC when required.

The high capacity of parallelization and calculation of HPC has allowed to the company to use fine spatial meshes and guarantee the correct resolution of the numerical simulation in a suitable time.

In summary, this SHAPE project has improved BAC's know-how on high performance computing and has strengthened its position for the next challenges.

Lessons learned

The experience in this collaboration in the SHAPE project framework has been satisfactory. It provided a good testing environment for the SME in order to test their software in thousands of cores, also to work with experts in compilation and porting of applications from BSC, so they were able to use the expertise in their own cluster and understand the limitations of the code from the scalability point of view.

One of the most confusing points has been to have making two different applications for PRACE resources, the first one to get the approval of the SHAPE project and the second one to get the real access for the calculation time.

3.4.10 FDD Engitec S.L. (Spain): Pressure drop simulation for a compressed gas closed system

Overview

The mission of FDD Engitec S.L. is the design and manufacturing of firefighting components and systems. Its core expertise is focused on the development of gaseous clean agent extinguishing systems. Its activities cover the design of components and systems, validation tests, manufacturing and certification. It is oriented to the global market, but mainly based in the European and North Africa markets.

The techniques the company employs for developing its technology is mainly mechanical engineering: elasticity and material resistance and fluid mechanics for compressed and non-compressed fluids. The team is currently focused on the design of a new firefighting system that will use inert compressible fluid.

The objective of the project is to calculate, by simulation, the pressure drop of a firefighting pressure regulated discharge valve for an inert gas agent.

This project was performed in collaboration with the Barcelona Supercomputing Center (BSC). The code used to perform the numerical simulations is Alya, a multi-physics code developed at BSC. The code has been adapted to solve this application problem.

Activity performed

The problem presented much more difficulties than those originally foreseen. These difficulties meant that solving the problem proposed completely would require a much larger amount of CPU time. Therefore, the focus was changed from the original project, to simulate a similar problem under the same conditions but to study the initial stages of the discharge, i.e., the first second instead of the first five minutes. This gave a good insight on what kind of resources would be necessary to solve the complete problem.

The problem complexity also required additional code developments, especially related to the boundary conditions. This fact led also to unforeseen additional programming effort. The final white paper will have a thorough discussion and a summary of the obtained results.

PRACE cooperation

The Partnership for Advanced Computing in Europe (PRACE) provided the expert support to adapt the code for this application and the machine time needed to perform the simulations.

Benefits for SME

The development of simulation techniques for discharging compressed gas across a valve installed over a closed system when the inlet condition changes with time is a big step forward in the design of these valves, as currently there is no commercial simulation tool that can manage it. During this SHAPE project the potential of the simulations and their complexity has been shown, as they require large HPC resources.

Lessons learned

The experience in this collaboration in the SHAPE project framework has been in general satisfactory. However, in the relationship with the SME it has been difficult to meet their expectations. The SME expected full solutions from the project that would allow them to take decisions on their product design. Unfortunately, with the resources available (the computational resources provided by Preparatory Access) it was not possible to achieve a full solution of their complex problem.

One of the most confusing points has been to have to make two different applications for PRACE resources, the first one to get the approval of the SHAPE project and the second one to get the real access for the calculation time.

3.4.11 Pharmacelera (Spain): HPC methodologies for PharmScreen

Overview

Pharmacelera is a company that develops hardware and software solutions for drug discovery. Among other products, Pharmacelera develops and sells PharmScreen, a revolutionary software tool for ligand-based drug design. PharmScreen scans compound databases consisting of hundreds of thousands of molecules and tries to identify potential hits by comparing molecules using a full 3D representation of its interaction fields.

In the course of the SHAPE project “HPC methodologies for PharmScreen”, Pharmacelera pursued the following four main goals:

- Defining a parallelisation strategy for PharmScreen using a combination of traditional CPUs, GPUs and highly parallel vector machines and setting up an appropriate testing methodology for such a HPC solution;

- Implementing a first parallel version of PharmScreen beyond OpenMP (which was already available);
- Exploring HPC alternatives for Pharmacelera's hardware platforms PharmaBox and PharmaBlade;
- Assess the performance and accuracy of PharmScreen by doing a sensitivity analysis of its different parameters.

Table 6: Project partners overview: HPC methodologies for PharmScreen

	SME	PRACE Partner
Company	Pharmacelera S.L.	RISC Software GmbH
Website	http://pharmacelera.com/	http://www.risc-software.at/en/
Contact / Support experts	Enric Gibert, CEO	Thomas Ponweiser Aleksey Kondratyev Thomas Steinreiter

Activity performed

For improving PharmScreen and taking advantage of HPC methodologies, the following main tasks have been performed during the project:

- PharmScreen performance profiling, bottleneck detection and definition of a parallelisation strategy;
- Code refactoring and parallelisation with OpenMP and OpenCL;
- Improvement of Testing;
- Porting of the application to run on the MareNostrum supercomputer;
- Scalability evaluation;
- Sensitivity analysis.

The scalability of the code has been evaluated in MareNostrum. Several applications have been tested with a variable number of threads and with different input set size in order to assess the influence of the initialization and data gathering stages. Figure 19 shows the speedups of PharmScreen showing a very nice scalability for the evaluated datasets. Other studies have been performed with the GPU implementation, which are shown in the white paper.

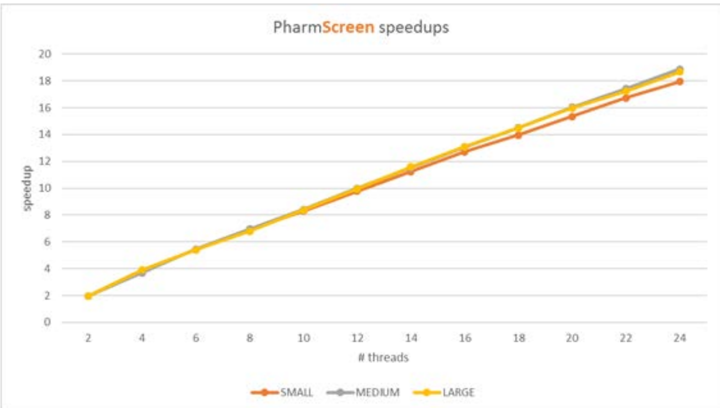

Figure 19: Speedups of PharmScreen execution for a variable number of threads

Moreover, the availability of the MareNostrum supercomputer allowed multiple sensitivity studies of the PharmScreen configuration parameters to be performed. Among them, the influence of grid spacing on the quality and simulation time was investigated.

Table 7: Grid spacing influence on quality

Grid spacing influence (Quality: From 0 to 1)						
	PreAlign grid spacing (Å)					
Final align grid spacing (Å)	0,5	1	1,5	2	3	4
0,5	0,815	0,838	0,839	0,853	0,837	0,816
1	0,826	0,8258	0,84	0,854	0,838	0,827
1,5	0,8255	0,824	0,843	0,8541	0,837	0,826

Table 8: Grid spacing influence on simulation time

Grid spacing influence (Simulation time h)						
	PreAlign grid spacing (Å)					
Final align grid spacing (Å)	0,5	1	1,5	2	3	4
0,5	5,783	3,162	2,911	2,915	2,914	2,914
1	3,714	1,137	0,900	0,862	0,743	0,738
1,5	3,373	0,821	0,732	0,569	0,498	0,492

Very interesting conclusions were drawn from this study, since it provided results that were not expected. Simulation time for different grid values was reduced as the spacing was higher, however, the quality of results did not improve after a certain distance reduction. The performed sensitivity studies have allowed Pharmacelera to optimize PharmScreen and increase its quality.

PRACE cooperation

Partners from RISC Software supported the SME with:

1. performance profiling and analysis;
2. porting computation kernels to OpenCL, enabling the usage GPU, Xeon Phi and other accelerator platforms;
3. porting the application to MareNostrum;
4. static code analysis and refactoring;
5. improving the testing methodology.

Computational resources for carrying out scalability and sensitivity analyses were provided via PRACE Preparatory Access Type B project 2010PA3391 by the Barcelona Supercomputing Center (BSC; MareNostrum supercomputer).

Benefits for SME

The SME has benefited from the HPC expertise of the PRACE partners in order to adapt its proprietary software to run in HPC infrastructures. Moreover, PRACE has provided them with computing hours that have allowed Pharmacelera to optimize the tool and make it more efficient.

Lessons learned

The collaboration between Pharmacelera and RISC Software is considered to have been unproblematic and fruitful by both parties. Initially Pharmacelera expressed preference for a support expert located near to their site. However in the end, fortunately neither spatial distance nor any language barriers affected the quality of the collaboration.

Key for the success of the project were regular Skype meetings as well as the responsiveness and commitment of both parties. In addition, the well-structured cloud-based document sharing and source version control provided by Pharmacelera were the basis for a frictionless collaboration. All in all, Pharmacelera is very satisfied with the project outcomes and both parties hope for the opportunity to collaborate again in the future.

3.4.12 Summary of lessons learned

In this section an attempt is made to summarise the "lessons learned" arising from the project reports. These will be fed back into the SHAPE process to help improve and reinforce the programme.

- National policies, with respect to allowing resource provision for industry users, need to be considered;
- On-demand HPC is important for some services offered by SMEs: this can be incompatible with the terms of usage for some HPC systems;
- In terms of communication, a common language (both literally and in terms of domain knowledge) between the PRACE partner and the SME can be very useful;
- Third-party software dependencies need to be considered carefully;
- There is an impact to using Open Source when there may be limited support available and issues arise;
- Expectation management with regards to the SMEs is important: what can actually be achieved in the time available and with the machine resources allocated needs to be carefully considered and made clear to the SME and in the workplan at the outset;
- Whilst geographical locality between the SME and the PRACE partner is often desirable, when it is not possible it does not appear detrimental to the project outcomes;
- Regular interactions (e.g. email, telcons, skype calls) and engagement with the SMEs are conducive to positive results;
- The double application approach of SHAPE (apply to SHAPE, and then apply for machine time) is confusing and cumbersome.

This final point was raised several times and was also noted in the previous deliverable, D7.1. As mentioned earlier, the inclusion of a preparatory access representative on the review panel was intended to ease this process by ensuring that the SHAPE applications would be suitable for preparatory access, but clearly this falls short of entirely removing the second application stage. The suggestion from CINECA that the PRACE partner working with the SME should be able to provide CPU hours to the SME at the start of the project is worth consideration. This is on the understanding that those hours are provided on an in-kind basis, and the equivalent hours would be made available from PRACE to the partner at a future date.

3.5 SHAPE: Fourth call

The fourth SHAPE call was launched In June 2016 and closed in September 2016. From seven applications, four were deemed appropriate for assistance from SHAPE:

- Artelnics, a Spanish SME looking at using HPC in their Neural Designer (BSC);
- Milano Multiphysics s.r.l.s, an Italian company hoping to use HPC acceleration to improve prediction of erosion and corrosion due to flowing liquid metals (CINECA);
- Renuda UK Ltd, looking at optimising 2D flow for faster, better steam turbine design (EPCC);
- Scienomics, a French SME looking at realistic modelling of complex, composite systems (IDRIS).

These projects are all underway with 17 person-months effort from PRACE partners allocated in total.

3.6 SHAPE: future

The SHAPE programme will continue in PRACE-5IP. Indeed, the next call will launch in April 2017. This call will close June 1st, as do some of the preparatory access calls so they are now aligned. As noted earlier, it is recommended that efforts should be invested in reviewing the SHAPE process, with the goal of a single-application process.

One issue raised by reviewers of SHAPE is what happens to the SME once they have ended their project. With this in mind, it is intended that a web resource be created with support and ideas for the SMEs to continue working with HPC, e.g. third-party service providers, other programmes and such-like.

Consideration should be given to creating stronger links with some of the other industry engagement programmes (such as Fortissimo [25]), to ease the SMEs transition from SHAPE to accessing commercial HPC. In addition, some thought should be given to any potential benefits for PRACE of creating a "second-round" SHAPE offer, to which SMEs who have already worked with SHAPE can apply if they are still not quite ready for accessing HPC via the more traditional methods.

As always, it is key to ensure that SHAPE is publicised effectively to ensure that SMEs can have the opportunity to participate. To this end, the ongoing engagement with the Industry Advisory Committee is important, but also new avenues are being considered. For example, following on from an exercise in identifying the publicity produced by SMEs who have already taken part in SHAPE, the expectations on new SMEs in this regard have been made clear at the outset of the projects, which should increase SHAPE awareness in their particular domains.

4 Summary

Two parallel sub-tasks on application enabling in Work Package 7 of PRACE-4IP have been described including final reports on the supported applications. These two activities have been organized into support projects formed on a basis of either scaling and optimisation support for Preparatory Access or SHAPE.

4.1 Preparatory Access

During PRACE-4IP Task 7.1.A successfully performed nine cut-offs for preparatory access including the associated review process and support for the approved projects.

In total eight Preparatory Access type C projects have been supported or are currently supported by T7.1.A in PRACE-4IP. Most of the projects reported in this deliverable plan to or have already produced a white paper. Approved white papers are published online on the PRACE RI web page [20]. Table 9 gives an overview of the status of the white papers for all projects. It shows the success of task 7.1.A as almost all finalized projects published their results or plan to publish it in the near future.

In addition to the already established PA C access scheme, the new Preparatory Access type D could be designed and finally launched in the context of PRACE-4IP. Together with the other PA schemes this provides a basis for the computational science community in Europe supporting and enabling them to port of further codes to PRACE Tier-0 systems.

Table 9: White paper status of the PA projects in PRACE-4IP.

Project ID	White paper	White paper status
2010PA2737	**WP211**: Numerical Simulation of complex turbulent Flows with Discontinuous Galerkin Method	Published online [26]
2010PA2821		No white paper produced
2010PA3125	**WP221**: Parallel I/O Implementation in Hybrid Molecular Dynamics-Self Consistent Field OCCAM Code using fpioLib Library	Published online [11]
2010PA3056	**WP222**: Optimising UCNS3D, a High-Order finite-Volume WENO Scheme Code for arbitrary unstructured Meshes	Published online [12]
2010PA3216	**WP223**: Enabling Space Filling Curves parallel mesh partitioning in Alya	Published online [16]
2010PA3234	**WP224**: Parallel curved mesh Subdivision for flow Simulation on curved Topographies	Published online [18]
2010PA3264	**WP225**: Particle Transport in a Fluid interacting with an immersed Body with Alya	Published online [19]
2010PA3430		Project will finish by the end of April 2017. White paper will subsequently be produced

4.2 SHAPE

The SHAPE programme continues; the fourth call projects are underway and the next call is planned to open in April 2017.

Feedback from the completed second call projects has been sought to evaluate the Return on Investment that participating in SHAPE has provided to the SMEs. Whilst much of this feedback was qualitative, the programme can be considered a success for most of the SMEs.

The three ongoing second call projects, and all eight third call projects, are either complete or will complete by May 2017, with the subsequent white papers available shortly after. Summaries of the activities of those projects have been described here.

Recommendations have been made on improvements to the SHAPE process, and also areas identified for future enhancements such as follow-on resources for the SMEs post-project.

5 Conclusion

Thanks to the Preparatory access Type C and the new preparatory accesses Type D, the PRACE infrastructures offered resources to European research community to assist in leveraging their applications. The processes in place allow discussion between the centre and the researchers to ensure that their request can be addressed.

All the applications that have been accepted by the Preparatory access C and D programmes have seen an increase in terms of performance or scalability of the application. This would suggest that this specific task must be continued in the upcoming PRACE phases in order to enable our research community to deal with the even larger supercomputers that will be available within the coming years. These tasks have enabled computing at the Petascale, in preparation for the road to Exascale.

Since its beginning in 2013, the SHAPE programme has provided assistance to 33 SMEs all over Europe. Furthermore, these successful applicants come from a broad range of industrial domains, highlighting the fact that there is a large cross-section of European small and medium businesses with the potential to be able to benefit from HPC.

SHAPE participants have reported many positive outcomes for the businesses involved in the activity including tangible measures of the Return on Investment: new staff people hired, increasing sales and new business, lower costs, continued HPC Access, in-house HPC systems deployment, integration into the HPC ecosystem and engagement with other national and European industry engagement programmes.

With regards the future of SHAPE, as mentioned in Section 3.6, one of the next challenges will be to ensure that all European SMEs can access the information and the opportunity to participate in the programme, and indeed to continue to engage with HPC following the programme. With this in mind the objective should be to ensure the range of SHAPE is far-reaching both geographically and in terms of the industry domains, to allow more SMEs to take advantage of the innovation possibilities opened by HPC.