SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

INFRA-2011-2.3.5 – Second Implementation Phase of the European High Performance Computing (HPC) service PRACE

PRACE-2IP

PRACE Second Implementation Project

Grant Agreement Number: RI-283493

D8.1.1
Community Codes Development Proposal

Final

Version: 1.0
Author(s): Claudio Gheller, Constantia Alexandrou, Miguel Avillez, Ivan Girotto, Sylvie Joussaume, Leonidas Linardakis, Ben Moore , William Sawyer, Thomas Schulthess
Date: 30.09.2011

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: RI-283493	
	Project Title: PRACE Second Implementation Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D8.1.1	
	Deliverable Nature: Report	
	Deliverable Level: PU	Contractual Date of Delivery: 30 / 09 / 2011
		Actual Date of Delivery: 30 / 09 / 2011
	EC Project Officer: Bernhard Fabianek	

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	Title: Community Codes Development Proposal	
	ID: D8.1.1	
	Version: 1.0	Status: Final
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2007	
	File(s): D8.1.1.docx	
Authorship	Written by:	Claudio Gheller, ETH-CSCS
	Contributors:	Constantia Alexandrou, CASTORC Miguel Avillez, UC-LCA Ivan Girotto, ICHEC Sylvie Joussaume, INSU/CNRS Leonidas Linardakis MPI-Hamburg Ben Moore University of Zurich William Sawyer, ETH-CSCS Thomas Schulthess, ETH-CSCS
	Reviewed by:	Peter Michielse, NCF; Dietmar Erwin, JUELICH
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	01/08/2011	Initial skeleton ready	
0.9	08/09/2011	All communities inputs collected	
1.0	12/09/2011	Draft for approval	
1.1	26/09/2011	Final version	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Scientific Communities, Scientific Applications, Astrophysics, Climate, Material Science, Particle Physics

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

The following references refers to the first two sections:

[1] http://www.prace-project.eu
[2] http://www.hp2c.ch/
[3] http://www.deisa.eu/
[4] http://www.deisa.eu/science/deci
[5] http://www.virgo.dur.ac.uk/
[6] http://www.hector.ac.uk/cse/distributedcse/reports/prmat/).
[7] http://www.astrosim.net/

Further references are listed in Section 3 and subsections.

List of Acronyms and Abbreviations

AMR	Adaptive Mesh Refinement
BSC	Barcelona Supercomputing Centre (Spain)
CEA	Commissariat à l'Energie Atomique (represented in PRACE by GENCI, France)
CFD	Computational Fluid Dynamics
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy)
CINES	Centre Informatique National de l'Enseignement Supérieur (represented in PRACE by GENCI, France)
CPU	Central Processing Unit
CSCS	The Swiss National Supercomputing Centre (represented in PRACE by ETHZ, Switzerland)
CUDA	Compute Unified Device Architecture (NVIDIA)
DEISA	Distributed European Infrastructure for Supercomputing Applications (EU project by leading national HPC centres)
DEMOCRITOS	Modeling Center for Research in Atomistic Simulations
EC	European Commission
EC-EARTH	European Earth System Model
ECHAM	European Centre Hamburg Model (global climate model)
EESI	European Exascale Software Initiative
ENES	European Network for Earth System Modelling
EPCC	Edinburg Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
EPFL	École polytechnique fédérale de Lausanne
EPSRC	The Engineering and Physical Sciences Research Council (United Kingdom)
ESFRI	European Strategy Forum on Research Infrastructures; created roadmap for pan-European Research Infrastructure
ETHZ	Eidgenössische Technische Hochschule Zuerich, ETH Zurich (Switzerland)
ETSF	European Theoretical Spectroscopy Facility
EU	European Union
FFT	Fast Fourier Transform
FZJ	Forschungszentrum Jülich (Germany)
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
GCS	Gauss Centre for Supercomputing (Germany)
GÉANT	Collaboration between National Research and Education Networks to build a multi-gigabit pan-European network, managed by DANTE; GÉANT2 is the follow-up as of 2004
GENCI	Grand Equipement National de Calcul Intensif (France)
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
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
GRAPE	Special purpose supercomputer for fast calculation of long-range forces

HI	Neutral Atomic Hydrogen
HII	Ionised Hydrogen
HITS	Heidelberg Institute for Theoretical Studies (Germany)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
ICON	Icosahedral Nonhydrostatic model, developed by DWD and MPI-M
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (represented in PRACE by GENCI, France)
IEEE	Institute of Electrical and Electronic Engineers
IESP	International Exascale Project
I/O	Input/Output
ISC	International Supercomputing Conference; European equivalent to the US-based SCxx conference. Held annually in Germany.
JSC	Jülich Supercomputing Centre (FZJ, Germany)
KB	Kilo (= 2^{10} ~10^3) Bytes (= 8 bits), also KByte
kpc	kiloparsec, approx. 3.08568025 × 10^{19} meters
LBE	Lattice Boltzmann Equation
LINPACK	Software library for Linear Algebra
LRZ	Leibniz Supercomputing Centre (Garching, Germany)
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. double precision) per second, also MF/s
MHz	Mega (= 10^6) Hertz, frequency =10^6 periods or clock cycles per second
MIC	Many Integrated Core architecture (Intel)
MoU	Memorandum of Understanding
Mpc	Megaparsec, approx. 3.08568025 × 10^{22} meters
MPI	Message Passing Interface
MPI-M	Max Planck Institute for Meteorology (Hamburg, Germany)
MPI-A	Max Planck Institute for Astrophysics (Garching, Germany)
MPP	Massively Parallel Processing (or Processor)
NCF	Netherlands Computing Facilities (Netherlands)
NFS	Network File System
NUMA	Non-Uniform Memory Access or Architecture
OpenCL	Open Computing Language
OpenGL	Open Graphic Library
Open MP	Open Multi-Processing
OS	Operating System
PetaFlop/s	Peta (= 10^{15}) Floating point operations (usually in 64-bit, i.e. DP) per second, also PF/s
PGAS	Partitioned Global Address Space
PM	Project Month
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
Psi-k	Electronic Structure Calculation of Solids and Surfaces.
PSNC	Poznan Supercomputing and Networking Centre (Poland)
QCD	Quantum Chromodynamics
QR	a procedure in linear algebra to factorise a rectangular matrix into an orthogonal (Q) and an upper triangular (R) matrix

RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
SARA	Stichting Academisch Rekencentrum Amsterdam (Netherlands)
SGEMM	Single precision General Matrix Multiply, subroutine in the BLAS
SHMEM	Share Memory access library (Cray)
SIMD	Single Instruction Multiple Data
SM	Streaming Multiprocessor, also Subnet Manager
SMP	Symmetric Multi-Processing
SNIC	Swedish National Infrastructure for Computing (Sweden)
SPH	Smoothed Particle Hydrodynamics
SSD	Solid State Disk or Drive
STFC	Science and Technology Facilities Council (represented in PRACE by EPSRC, United Kingdom)
STRATOS	PRACE advisory group for STRAtegic TechnOlogieS
TB	Tera (= 2^{40} ~ 10^{12}) Bytes (= 8 bits), also TByte
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

Executive Summary

Science and its challenges are primary targets of the PRACE effort. Work Package 8 (hereafter WP8) attempts to foster these by establishing and consolidating long-term collaborations between HPC centres and scientific communities. These collaborations in turn allow scientific numerical tools to be incrementally updated to innovative computational solutions. The premise applied in WP8 is that the responsibility of porting and scaling an application onto novel supercomputer architectures should remain with the application developers, who have the best grasp of methods and algorithms used in a particular scientific domain. Supercomputer centres with their profound expertise of computing architectures and programming models, have to recast their service activities in order to support, guide and enable scientific program developers and researchers in refactoring codes and re-engineering algorithms, thereby influencing the development process at its root. The resulting codes will fit both Tier-1 and Tier-0 allocation schemas, and it will be the particular requirements of the users of these community codes that determine whether applying for Tier-1 or Tier-0 resources.

Note that this approach complements that of WP7, "Enabling Petascale Applications: Use of Tier-0 systems", in both PRACE-1IP and 2IP [1] projects, where experts from HPC centres port or "petascale" existing codes to current high-end HPC systems, in the framework of the Access Projects. WP8 addresses, instead, scientific communities and their needs toward the medium-to-long term future for their HPC numerical tools.

This document describes the work accomplished during the first month of the PRACE-2IP project in WP8, which aimed at identifying and selecting communities in different domains of science (hereafter denoted as "scientific communities"). This process has involved both PRACE-2IP partners and representatives of the scientific communities. In a first stage, the criteria for selection of the communities have been defined. This has been followed by the identification of the communities and their representatives that have been finally involved in the discussion, and have provided all the elements necessary to outline the work plan. Material Science, Climate, Astrophysics and Particle Physics are the scientific domains resulting from this initial selection phase. A set of additional disciplines has been identified that may be added later.

1. Introduction

Today, all successful development projects for sustained petaflops applications have in common that major portions of application codes had to be rewritten (software refactoring) and algorithms had to be re-engineered for applications to run efficiently and productively on novel architectures. This principle is not limited to the high-end of supercomputing, but applies to all tiers of the HPC ecosystem. The compute nodes architecture, in fact, follows current technology trends towards more parallelism and more customisation. Compute nodes will integrate many cores, potentially hundreds, some of which will serve as computational accelerators (e.g. GPU, MIC). These technology trends pose major challenges to which software needs to respond swiftly and effectively. However, the effort involved in software refactoring can be substantial and often surpasses the abilities of individual research groups. Furthermore, such effort often requires a deep knowledge of the algorithms and of the codes and, in many cases, even the understanding of the physics beneath the numerical algorithms. Therefore, a successful practice is to place these activities in community projects where scientific code design and implementation is undertaken and thus long-term software development for high-end computing can be sustained, exploiting the know-how of HPC

experts to support, ease and consolidate the effort toward the most innovative technological solutions.

This document describes the community selection accomplished in the first month (PM1) of the PRACE-2IP project. An early start-up was necessary, since a relevant number of different tasks had to be accomplished by the end of PM1:

- Set-up of the partnership (get in touch with the partners, define the contacts, begin and coordinate the work etc.)
- Assess the scientific communities selection criteria
- Select the communities
- Contact and involve the communities' representatives and share the vision, the objectives and the working strategy.
- Design the deliverable and coordinate the work on it
- Collect the inputs from the communities
- Write the deliverable.
- Prepare, organise and carry on the activities related to the PRACE-2IP Kick Off meeting.

The kick-off conference call was held two months before the official PRACE-2IP start-up (July 4th, 2011).

The present document is organised as follows:

Section 2 describes the methodological approach adopted to choose the scientific domains, introducing the selection criteria and describing the procedure followed to define the list of selected communities.

These communities are described in details in Section 3, together with their scientific challenges, an overview of codes and applications they use in their research work and the improvements required to make these instruments suitable to research topics that will be solved on the next generations of supercomputing architectures.

Section 4 aims at describing the organisation of WP8, in terms of hierarchies and workflow, necessary to properly coordinate the action of PRACE partners and the scientific communities.

Finally, Section 5 summarises the outcomes of the WP8 accomplished work thus far and summarises the next steps, leading to choose codes and algorithms to work on, by defining appropriate application performance model selection-based criteria.

2. Methodological Approach

PRACE is founded on the idea that scientific communities, with their high-end research challenges, are the main drivers for software development. Software must be re-designed in order to meet the scientific requirements and adapted to effectively exploit emerging supercomputing architectures. In order to reach such an ambitious target, a deep synergy between HPC experts and application developers from the communities must be fulfilled and a strong commitment between the scientific counterparts has to be achieved. A similar community software development platform has already been implemented in Switzerland by ETH Zurich, through CSCS, the workpackage coordinator. Such project, called *Swiss Platform for High-Performance and High-Productivity Computing* - HP2C [2], is concentrating on developing petascale applications in more than a dozen research areas, involving all major Swiss universities. The experience acquired in the HP2C framework has

been leveraged in order to have a prompt and effective start and set up of WP8, necessary to successfully carry out the ambitious workplan proposed by WP8 in the short time frame (two years) available. This, first of all, required a fast selection procedure, in order to enable the partners to start implementation work as soon as possible.

2.1 The selection criteria

The first step of the WP8 workplan was the identification and selection of the scientific communities. Here, the scientific community is defined as a structured group of research teams acting jointly in a specific scientific domain (astrophysics, chemistry…) with relationships and interactions among them, sharing opinions and using the same scientific methodologies to advance science in that specific domain. Scientific communities involved in computational sciences are generally interested in adopting similar computational approaches and methodologies, and using high-end HPC resources to advance science through simulation methods. These communities, in general, use specific categories of application codes and similar mathematical models.

In each different scientific domain, a large number of communities are present and active, and a selection is obviously needed in order to set up a realistic workplan, compatible with the finite amount of resources of the workpackage. Hence, a selection procedure had to be defined.

Such selection procedure has been designed by the PRACE partners, exploiting their different experience and competencies, in order to be

- fair and balanced (in order to distribute the effort over a broad spectrum of scientific themes);
- light and flexible (in order to prevent procedural bottlenecks and to extend, if/when possible, to a larger number of disciplines);
- consistent with the available resources (in order to avoid under or over commitment in the selected activities, achieving the best possible outcome out of the available PMs).

The procedure has been accomplished at PM1, beginning two months before the official start-up of the project. However, it can be easily iterated along the two-year duration of the project, if new communities propose collaborations with PRACE, and sufficient resources are available for their inclusion.

The foundation of the selection procedure lies on the following four criteria, defined in the PRACE-2IP description of work and further discussed, shared and agreed upon by the partners during the first WP8 meeting:

1. the candidate community must have high impact on science and/or society;
2. the candidate community must rely on and leverage high performance computing;
3. WP8 can have a high impact on the candidate community;
4. the candidate community must be willing to actively invest in software refactoring and algorithm re-engineering.

While the first three items can be expected, the last one represents a real breakthrough, since it requires an unfunded contribution of the community to the project, implicitly expressing the acknowledgment of the relevance and the impact of the PRACE action on the progress of its scientific activity. This also turned out to be the most severe selection criterion, distinguishing between those communities really willing to collaborate and all the others, thus minimizing possible risks of failure by involving only highly-motivated communities.

2.2 The selection methodology

The community selection was accomplished through a number of steps by:

1. Screening of the current European scientific HPC scenario, by means of the experience acquired in the PRACE-1IP and DEISA projects;
2. Identifying a number of scientific domains to search for communities suitable to the project;
3. Identifying candidate scientific communities in the selected domains and contacting these communities;
4. Harvesting information from the communities and selecting a few of them according to the defined criteria.

The selection procedure is an open process, the inclusion of further communities being possible "on-demand", namely by performing the fourth step for those communities interested to in working with PRACE.

Step 1

A screening of the current scenario of HPC scientific applications has been outlined and adopted for a first selection of scientific domains, harnessing the many years of experience achieved in the DEISA [3] and PRACE-1IP projects.

First, The DEISA Extreme Computing Initiative (DECI [4]) has been analysed during 2008-2011. The DECI program aimed at enabling a number of "grand challenge" applications in all areas of science and technology, running on the DEISA HPC grid. The supported projects represent a meaningful sample of the current trends and needs in the European High Performance Computing panorama. About 150 projects relied on DEISA resources. Also, the corresponding action promoted by PRACE-1IP project during 2011, with two "Access" calls for large-scale projects, has been considered. The number of projects per scientific area resulted to be:

Table 1: HPC large scale projects per scientific area.

Discipline	Number of projects	%
Life Science/Comp. Chemistry	30	17%
Material Science	35	20%
Climate	9	5%
Geophysics	6	3%
Engineering/CFD	27	16%
Astrophysics	22	13%
Plasma/Particle physics	38	22%
Other	6	4%
Total	173	

The PRACE-1IP "preparatory access" calls were also examined. Here, resources are assigned to prepare for Access projects, i.e. to test, benchmark and improve codes in order to be suitable for Tier-0 HPC systems. In this case we found:

Table 2: HPC development projects per scientific area.

Discipline	Number of projects	%
Life Science/Comp. Chemistry	9	16%
Material Science	7	12%
Climate	6	10%
Geophysics	2	3%
Engineering/CFD	10	17%
Astrophysics	9	16%
Plasma/Particle physics	6	10%
Other	9	16%
Total	58	

The outcomes of two relevant deliverables of the PRACE-1IP project have been analysed as well in the following documents.

D7.2.1, "Interim Report on Collaboration with Communities"

That deliverable reports initial steps of the community codes take-up action, performed by WP7.2 task of PRACE-1IP. The document is focused on the procedure designed and adopted to select a number of community codes to be enabled on Tier-0 systems. The WP7.2 methodology focuses primarily on applications and the way the related algorithms can be improved for the high-end computing systems currently available in PRACE. This is different from WP8 approach that focuses on the scientific communities and those scientific problems relying on the next generations of HPC architectures. However, it is interesting to notice that the communities addressed by WP7.2 are material and life science, computational chemistry, astrophysics and plasma physics, earth science and engineering, confirming the indications presented by the statistics above.

D7.4.1, "Applications and user requirements for Tier-0 systems"

This document reports on the surveys carried out on the PRACE-1IP partners' HPC systems, the applications running on them, and of current/potential users of the PRACE infrastructure. The user survey pointed out that a large fraction of the users could be identified in the fields of computational chemistry and material science, followed by computational fluid dynamics, life science and astrophysics. Interestingly, most of the users were also code developers (less than half of them declaring to be end users), proving the interest of the scientific community also in the numeric and code implementation aspects.

Step 2

From the previous analysis, it is clear that life science and computational chemistry, material science, earth science, engineering, astrophysics, plasma and particle physics are domains where HPC has a great impact. Among these domains, climate and astrophysics have a clear need for further code development on HPC systems, especially for innovative architectures and software solutions. Material science and plasma/particle physics can rely on more mature codes, but they prove to be major HPC resources consumers, requiring larger and larger computational resources. Therefore, their codes must be ready to exploit the new hardware solutions effectively. The need for code improvement is also clear in life science, and in particular in molecular dynamics. The same can be emphasised for the engineering and CFD

area. Other disciplines (mathematics, computer science etc.) are mainly focused on code development and experimentation, but are not high consumers of HPC resources, hence they can be considered irrelevant for PRACE.

Step 3

After Step 2, six scientific domains were ready to be explored in order to verify the presence of communities suitable and relevant for PRACE and preferably recognised also by the European Commission (and, possibly, funded in the FP7 framework). Each partner, according to its expertise, interests and contacts, contributed to identifying such communities as well as contacts to get in touch with them.

Step 4

The partners collected information and feedbacks from different communities. This information was used to direct further investigation toward specific communities and to drop some others, according to the adopted selection criteria. For instance, in the astrophysics domain, three communities expressed interest for the project: the AstroSim network [7], the Virgo consortium [5] the PFARM group [6]. After a first analysis, the PFARM group was dropped, being too focused on a single application and addressing the need of a too narrow community. Then, also the Virgo consortium was also excluded, since, although an outstanding community with extraordinary results in the field of numerical cosmology, it did not express the required interest in actively contribute to the PRACE effort.

All the outcomes of the performed investigation were finally presented and discussed in a conference call held on July 25th, 2011, when the final list of communities was defined (as will be presented in Section 3). In the same meeting it was also decided that this list was not conclusive, but further communities could be added later. A few of these communities were already “in the loop”, but, due to the time constraint for starting the work, they did not manage to formally express their involvement on time.

A final remark needs to be made to stress that, in contrast to that stated in the PRACE-2IP Description of Work, the Scientific Steering Committee has not been involved in the initial selection procedure. This exception was mainly due to the need to complete the selection in a short time, in order to start the work with communities as soon as possible. However, the Scientific Steering Committee is being consulted for a final approval of the accomplished communities selection, and will be actively involved in the next steps of the WP8 workplan, in order to effectively advice on the selection of the most promising between the codes proposed by the communities.

3. Selected Scientific Domains and Communities

The selected communities belong to four different scientific domains, all strongly relying on HPC for their scientific production:

- Astrophysics, with clearly identifiable major development teams in Europe and focused scientific, numerical and computational problems.
- Climate: structured and organised domain, with heterogeneous computational problems, ranging from HPC to data processing to data storage and accessibility. Its impact on the workpackage appears to be potentially extremely relevant
- Material Science: extremely large domain with a large number of research topics. Specific areas can be identified where HPC has an extraordinary impact.
- Particle Physics: clear scientific and numerical targets; well-defined stakeholders in the community; it needs HPC as a primary tool for research.

In the following sections, we will describe the different scientific areas and communities providing:

1. A general description of the selected scientific area and of the need for HPC in the field;
2. Highlights of scientific challenges in each field that can be addressed only exploiting a new generation of HPC architectures and related highly optimised numerical codes;
3. An overview of the main numerical approaches and codes adopted by the community. Some of them, according to the indications of the community experts, will be the subject of a redesign and refactoring activity.
4. A description of the communities selected in the various domains
5. The expected outcomes of the work proposed for WP8 and their impact.

Notice that most of the material presented below was provided directly by members of the different selected scientific communities. Therefore, all the information represents the expression of the need of the communities and of the numerical instruments they wish to adopt and exploit in their research activity, reflecting their actual expectations and requirements. For their contributions, we acknowledge in particular Sylvie Joussaume (INSU/CNRS) and Leonidas Linardakis (MPI-Hamburg) from the climate community, Ben Moore, Romain Teyssier (University of Zurich) and Miguel Avillez (University of Coimbra) for astrophysics, Constantia Alexandrou (University of Cyprus) for particle physics and Thomas Schulthess (ETH-Zurich) for material science.

3.1 Astrophysics

In the past decade a combination of satellite- and ground-based observations have mapped and quantified the observed universe to an unprecedented precision. Observers are acquiring accurate multi-wavelength data for over a million galaxies at large scales as well as for the interstellar medium and stellar systems in the Milky Way and nearby galaxies (e.g., Magellanic Clouds, M33 and M31) at small scales. They have high-resolution spectral information, color maps, element abundances and kinematical data for individual galaxies. The observational insight provided by these instruments often needs theoretical explanations, which in turn rely on the solutions of systems of complex equations coupling together macro and microscopic processes (e.g., fluid dynamics, turbulence, atomic and molecular...). Simple analytical solutions are impossible at three-dimensions and therefore, theorists have to solve these equations numerically using the most adequate numerical methods. However, this is not an easy task due to a set of constraints that act together to render difficulties to the numerical solution of these astrophysical problems.

In fact, the success of such efforts requires a compromise between resolution, dimensionality and adequacy of the physical processes for the resolution, evolution time and length scales and convergence of the solutions associated to the included physics. For instance dimensionality reflects itself in the processes evolution - two-dimensional and three-dimensional turbulence being rather different. Three-dimensional turbulence is characterised by the dissipation of energy through vortex stretching, a property that is not present in two-dimensional turbulence. Resolution is of most importance in order to capture the appropriate physics – if the adopted resolution is lower than the length scales for the relevant physical processes (e.g., cooling, atomic and molecular processes) the simulation does not catch these processes - the effects of their feedback into the system go unnoticed. Thus, in order to be able to resolve the smallest physical structures and have enough resolution for the relevant physical processes feedback to be captured special purpose numerical techniques that increase resolution (e.g., adaptive mesh refinement) on-the-spot and on-the-fly had to be developed and implemented. However, due to their complexity only a few codes have this capability, although these techniques are available in fluid dynamics since 1984 [1] and imported to astrophysics fluid dynamics by [2].

The last decade has also seen a large growth in computational resources leading to massive parallel computations and sophistication on astrophysical software by including more and more physical processes, up to date numerical methods, and pursuing numerical optimisation and computational acceleration by the use of dedicated boards (GRAPE boards [3]) or the use of graphics processor units (GPUs). With these developments numerical simulations play a central role in understanding the Universe and the observational phenomena. However, theorists have not yet succeeded in making numerically a complete picture of the formation of disk galaxies from first principles, nor were they able to trace the evolution of turbulent systems in galaxies since injection scales down to the dissipation ones (a fundamental problem even in turbulence research). Similar difficulties in providing a complete picture exist on other astrophysics scales. Whilst astronomers are cataloging the properties of numerous extra-solar planetary systems, theorists have yet to form a single realistic solar system via computational techniques. Despite these difficulties, theorists are not so far behind and the solution of many ongoing problems will be at reach with the increasing access to computational resources and support to the community.

3.1.1 Scientific Challenges

Despite the computational and numerical difficulties that theorists face in providing realistic solutions to different problems, there are ongoing efforts by the community to address within the next few years several fundamental problems that are still lacking an adequate answer.

On the largest scales, comparable with the size of the universe, how cosmic structures form and evolve is a highly non-linear problem involving a huge range of time, mass and length scales and a host of important physical processes. The cosmological parameters that describe the big bang model and our expanding universe have been measured, although only about one percent of the universe has been physically identified and understood. Thus although the initial conditions for structure formation are known, it is not known what the dominant components of matter and energy actually are. Numerical simulations are an essential tool for constraining models and candidates for the nature of dark matter and dark energy. How these fluctuations in the dark matter and baryonic components form galaxies, stars and eventually planets, involves complicated non-linear processes including gravity, hydrodynamics, radiation and electromagnetism. Making accurate theoretical predictions requires the skills of researchers working on a broad range of topics combined with grand challenge supercomputing programs. This is an area in theoretical cosmology that is driven primarily by large scale supercomputer simulations.

Figure 1. A multi-billion particle simulation of the dark matter distribution surrounding a galaxy. The simulation took several million cpu hours using the PKDGRAV code on the Mare-Nostrum supercomputer [4].

Understanding the origin of the Milky Way has been a key goal of numerical cosmology and progress has been steady over the years. Constraints on the dark energy and cosmological parameters require large scale simulations where the locations of galaxies and their masses are accurately determined. That requires about 100 particles per dark matter halo. In order to resolve the formation of individual galaxies and their internal structure requires about 100 million particles per galaxy.

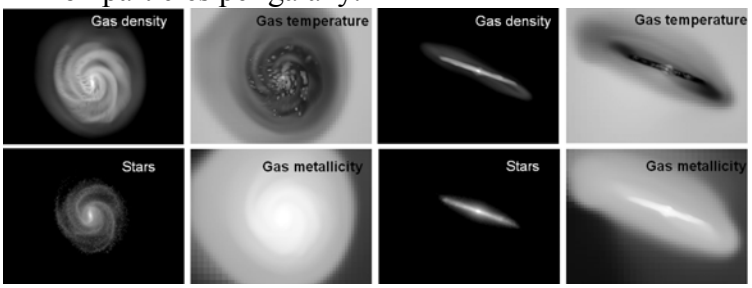

Figure 2: Face on and edge views of a "Milky Way" galaxy formation simulation [5]

Another open issue is how the cosmic magnetic fields pervading nearly all galaxies and clusters of galaxies originated? It is generally accepted that seed fields, whose strength is of order 10^{-20} Gauss, easily spring up in the era preceding galaxy formation. Several mechanisms have been proposed to amplify these seed magnetic fields to a coherent structure with the microgauss strengths of the currently observed galactic magnetic fields. It is more reasonable for the first and largest amplification to occur before the Galaxy forms, and the matter embedded in the field is gravitationally trapped. Primordial mechanisms have not yet been seriously developed, and this preliminary amplification of the magnetic fields is still open.

Star formation is central to a broad range of astrophysical problems, essential for explaining not only the distribution and types of stars we see around us today but also providing the concrete underpinnings for understanding the history and nature of the very first stars in the

early universe. How interstellar molecular clouds collapse and form new stars is a key question in astrophysics. Massive stars play an important role in the evolution of galaxies by injecting large amounts of enriched material and energy back into the interstellar medium and powering spectacular phenomena such as HII regions, stellar winds and supernovae.

Another challenging and computational demanding problem is represented by the birth and evolution of planets. Many simulations show the development of spiral structures, similar to those seen in some galaxies including the Milky Way. So far, different groups have reached conflicting conclusions about the role of the gravitational instability. Some results suggest that cooling in discs is efficient enough to allow giant protoplanets, precursors to Jupiter-like planets, to form due to the gravitational instability in the spiral structures. Other results lead to the opposite conclusion. Thus, high-resolution modelling of plane formation around solar-type stars may resolve the uncertainty.

Some of the most puzzling problems related to the physics of the interstellar medium (ISM) concerns the existence of a large fraction of extraplanar gas that is not contained in classical HII regions, but rather in a diffuse layer (named DIG for diffuse ionised gas), extending up to a few kpc on either side of the midplane of the Milky Way and external galaxies. It is ubiquitous in the Galaxy and it occupies approximately 20% of the volume within a 2 kpc-thick layer above and below Galactic plane.

Finally, the overall problem regarding cosmological/structure formation and Galactic/ISM modelling is the fact that plasmas keep a record of its history. In fact, most of the large scale simulations carried out by the different communities relies on collisional ionisation equilibrium, which implies that the number of recombinations equals the number of ionisations. However, this is not quite true. This means that the ionisation structure, cooling functions and emission spectra are quite different between collisional ionisation equilibrium and non-equilibrium conditions as can be seen in Figures 4 and 5. Furthermore, the usage of cooling using temperature grids does not work, since the plasma evolution depends on its own history and hence regions with the same initial conditions, end up having different cooling and emission paths (see review by, e.g., Avillez & Breitschwerdt [6]). Therefore, grid and sub-grid physics simulations must be self-consistent and couple together the dynamics with the atomic/molecular ionisation structure of the plasma in a self-consistent fashion.

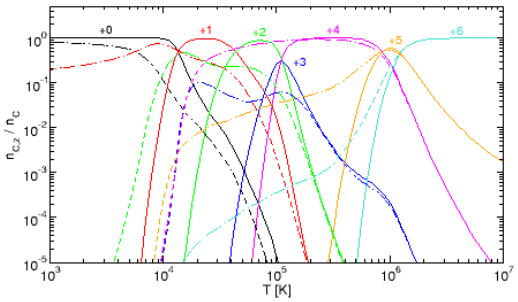

Figure 3: Ionisation structure of Carbon calculated in collisional ionisation equilibrium (CIE) and non-equilibrium ionisation (NEI) conditions (dashed lines)

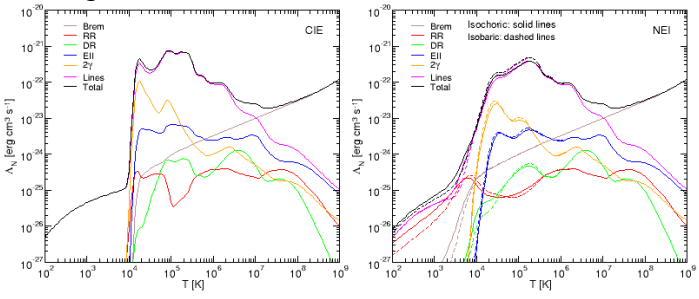

Figure 4: Normalised CIE (left panel) and NEI (right panel) cooling function calculated with solar abundances. Total cooling (dashed line), Bremsstrahlung (brown), radiative (red) and dielectronic (green) recombinations, electron impact ionisation (blue), two-photon (orange; 2γ) and line excitation (excluding 2γ; magenta) emissions.

3.1.2 Numerical Approaches and Community Codes

Modelling different parts of the Universe involves the adoption of a spatial scale and of a medium the system is represented by. In general, astrophysical structures and objects (clusters of galaxies, a galaxy, interstellar medium, star clusters, black holes, stars, planets), are represented as a fluid whose behavior is described numerically according to two different approaches: (1) the fluid is divided into a collection of fluid elements, whose history (evolutionary path) is traced (Lagrangean approach) or (2) the fluid's history is traced as a whole from a reference frame (Eulerian approach) on the cell of a computational mesh. Each of these approaches is used according to the problem at hand and the dominant physics: gravitationally dominated systems are better solved using a Lagrangean approach where the fluid is represented by an N-body system, with each particle enclosing a large mass (in the case of galaxies formation or cosmological simulations each particle can enclose millions of solar masses). The Eulerian approach becomes more effective for galaxy-scales, interstellar medium, stellar evolution and turbulence simulations. An hybrid scheme coupling these two approaches has been used to track the dynamics of individual systems (by particles) embedded in a fluid, e.g., evolution of newly formed stars (represented by particles) within a turbulent molecular cloud (represented by a fluid).

The calculation of the gravitational forces, driving the dynamics of most of astrophysical systems, is a crucial step for both approaches, both for the correctness and accuracy of the results and for the performances that can be achieved, in particular, on parallel supercomputers. In fact, due do their long-range nature, gravitational forces are intrinsically challenging to parallelise using a standard domain-decomposition approach. Many different methods have been adopted to face such challenge, but any solution to be effective and efficient has to be tuned for the target computing architecture.

One of the most recent numerical challenges is represented by the inclusion of the magnetic fields into numerical models. In magneto-hydrodynamics (MHD) the efficiency of the methods still lags its hydrodynamical counterparts. Since almost all problems involving magnetic fields require a 3D treatment, the development of accurate MHD schemes involving adaptive meshes has been a primary objective in the last decade and comprise developments for better accurate, monopole-free, schemes with low resistivity and be also able to cope with various non-ideal and collisionless effects like non-isotropic pressure and heat conduction. These developments allow a further increase in modelling MHD turbulence, dynamos and particle acceleration.

These methods have been adopted in a few public community codes widely used by the community. They can be classified similarly to the above numerical approaches: particle-based (e.g., Gadget, Gasoline, Hydra, Pkdgrav, etc.) and grid-based (Enzo, Athena, Zeus, Piernik, EAF-PAMR, Ramses, Nirvana, Flash, Pluto, and Pencil, among others) software. We briefly overview those that are developed by European groups and that may be subject of the refactoring and take-up action of WP8.

The GADGET code [7] in its three subsequent versions, represents one of the most successful applications for cosmology. Gadget, developed at MPI-A Garching (Germany) adopts a Lagrangian approach and combines smoothed particle hydrodynamics with a hierarchical Tree+PM algorithm for gravitational forces. It is highly optimised for parallel computing and it was used to run some of the largest cosmological simulations ever, like the “Millennium Run”, a pure dark matter simulation of about 10 billion particles.

The “Pkdgrav” N-Body code and its fluid dynamics counterpart “Gasoline” [8], represents further examples of TreeSPH codes developed with HPC in mind. A recent version of the code, called Changa [9], has been implemented adopting a combination of some of the up-to-

date approaches as CHARM++ and CUDA, addressing the effective exploitation of multiprocessors – multicore – multiGPU systems. The code is developed between the University of Zurich (Switzerland) and the University of Washington (USA).

RAMSES [10], originally developed in Saclay (France) now at the University of Zurich (Switzerland), is a cosmological AMR code that overcomes the limits posed by the usage of a tree structure by adopting an original refinement strategy, dynamically creating (or destroying) refinements on a cell-by-cell basis and assigning to each refined cell the information about the neighboring and parent cells. In this way the tree is eliminated and work and data can be distributed effectively between processors. This allows greater flexibility to match complicated flow geometries. Gravity is calculated by a novel approach based on the multigrid technique, supporting domains with irregular boundaries.

The ultimate tool for cosmological simulation is represented by the AREPO code [11], developed at MPI-A and HITS (Germany). AREPO overcomes the disadvantages of both the Lagrangian and of the Eulerian approaches relying on a moving unstructured mesh defined by the Voronoi tessellation of a set of discrete points. The mesh is used to solve the hyperbolic conservation laws of ideal hydrodynamics with a finite volume approach, based on a second-order unsplit Godunov scheme with an exact Riemann solver. The mesh-generating points can in principle be moved arbitrarily. The current version is parallelised for distributed memory computers. Notice that currently AREPO is not an open project.

The EAF-PAMR code [12] is a three-dimensional adaptive mesh refinement code uses the MHD extension of the Piecewise Parabolic Method [13] [14] and the grid refinement procedure discussed in [15] assuring that B is divergence-free during the refinement process and simulation time. It includes local self-gravity and heat conduction. The code further includes an atomic and molecular plasma emission module (E(A+M)PEC; http://www.lca.uevora.pt; [16]) to follow in a time-dependent fashion the ionisation structure of the ten most abundant elements in nature and determine on the spot the cooling function and emission spectra (the latter is then passed through instrument response to provide the observations by different telescopes like ROSAT, XMM-Newton, Chandra, FUSE, etc).

PIERNIK ([17] [18]) is a multi-fluid, grid-based MHD code based on the Relaxing TVD scheme. The original scheme has been extended by an addition of dynamically independent, but interacting fluids: dust and a diffusive cosmic ray gas, described within the fluid approximation, with an option to add other fluids in an easy way. As all the codes discussed above PIERNIK is parallelised by means of the MPI library.

The PLUTO code [19] developed at the Turin Astronomical Observatory is a modular Godunov-type code intended mainly for astrophysical applications and high Mach number flows in multiple spatial dimensions. The code solves the equations describing Newtonian, relativistic, MHD, or relativistic MHD fluids. It uses the Chombo adaptive mesh refinement library to manage structured adaptive grids. The Chombo package provides a set of tools for implementing finite difference methods for the solution of partial differential equations on block-structured adaptively refined rectangular grids. Both elliptic and time-dependent modules are included. Support for parallel platforms and standardised self-describing file formats are included. Just like ENZO, PLUTO has the main disadvantage to replicate the AMR tree on the processors.

NIRVANA [20] is a general-purpose C code for astrophysical research which numerically integrates the three-dimensional equations of time-dependent, compressible magneto-hydrodynamics. NIRVANA provides adaptive mesh refinement techniques to handle multi-scale problems and implements a Poisson solver to treat self-gravity. It further includes thermal conduction and tracer particles and is capable of dealing with ideal and non-ideal

magnetised systems. The MHD equations are solved within a semi-discrete finite-volume framework. Two Godunov-type solvers each combined with upwind constrained transport (CT) techniques for a divergence-free evolution of the magnetic field are implemented.

3.1.3 Community involvement, expected outcomes and their impact

The community

As discussed above, Astrophysics covers a broad spectrum of different research themes, and a number of numerical solutions are available for facing the various challenges. Our goal is to support the largest possible variety of scientific topics, compatible and consistent with the goals of WP8. These requirements perfectly fit the AstroSim Community [21]. The AstroSim programme is a five-year project, funded by the EC in the FP7 framework. It aims to bring together European computational astrophysicists working on a broad range of topics, from the stability of the solar system to the formation of stars and galaxies, strengthening the existing European activities in computational astrophysics, avoiding fragmentation as this field grows in strength, and to exchange expertise through an active program of conferences, workshops, training schools and exchange visits. The project is finishing in September 2011, and PRACE can provide a way to keep the community alive even after the end of the programme, capitalizing on the previous effort. Fourteen European countries and related national research organisation contributed to AstroSim. A Steering Committee, led by Prof. Ben Moore (University of Zurich), managed the project.

The scientific community code development effort in WP8 will be driven by eminent members of the astrophysical community, part of the AstroSim project, such as Ben Moore, Romain Teyssier and Joachim Stadel (University of Zurich - Switzerland), Stefano Borgani (INAF – Italy), Miguel de Avillez, Emanuele Spitoni, and Nuno Carvalho (University of Evora - Portugal).

For PRACE-2IP (University of Evora, which is project partner through UC-LCA) will lead the work and several other partners, like CSCS and EPCC will contribute.

Outcomes and Impact

It is clear from the above discussion that, although many different numerical and computational solutions are available, there are problems (i.e. physical processes) that are relevant for and common to most disciplines (e.g. calculation of the gravitational forces or of the magnetic field) and that are particularly challenging, especially when large scale applications are considered and HPC is decisive.

In WP8 we will address some of these challenges, supporting the scientific codes developers in the adoption of the most effective solutions in view of the next generation HPC architectures. More specifically:

1. Optimisation, acceleration and testing of specific modules (comprising the latest numerical techniques) to be used by the astrophysical community allowing their integration directly into ongoing code development by the different research groups throughout Europe. The modules under study include: (i) adaptive mesh refinement library, (ii) plasma emission on-the-stop calculation modules, (iii) advection schemes for grid-based software, and (iv) tree-structure for particle-based calculations.
2. Implementation of hybrid numerical schemes exploiting multi-node, multi-cores hybrid architectures (i.e. with accelerators like GPUs or MICs) adopting well established programming paradigms relying on MPI, OpenMP, CUDA, OpenCL, and their effective combination and synergy. This allows the community to better decide which approaches and schemes are better suitable for the particle and grid-based calculations.
3. Analysis of the possible interplay between codes acting on different scales and/or

describing different physical processes in order to couple them, possibly better exploiting heterogeneous systems, where some algorithms run efficiently on specific computing devices while some others are most suitable for a different architecture (e.g. hydrodynamics can overwork GPUs, while gravity can exploit shared memory multicore nodes best).

The impact that the WP8 work may have on the astrophysics community could be extraordinary. Since the focus is not on a specific code, but on the refactoring and optimisation of methods common to most of the application codes, the proposed solutions can be adopted and exploited by a large number of developers in order to exploit in the most suitable way future HPC systems for the current high-end scientific goals.

3.1.4 Relevant Bibliography

[1] Berger M.J., Oliger J., 1984, Journal Computational Physics, 53, 484
[2] Berger M.J., Colella P., 1989, Journal Computational Physics, 82, 64
[3] http://www.astrograpе.org/
[4] Stadel, J., D. Potter, B. Moore, J. Diemand, P. Madau, M. Zemp, M. Kuhlen, and V. Quilis 2009. Monthly Notices of the Royal Astronomical Society 398, L21-L25
[5] Agertz, O., R. Teyssier, and B. Moore 2011. The formation of disc galaxies in a LCDM universe. Monthly Notices of the Royal Astronomical Society 410, 1391-1408
[6] Avillez, M.A. & Breitschwerdt, D. 2004, Astronomy & Astrophysics, 425, 899
[7] Springel V., 2005, MNRAS, 364, 1105
[8] http://hpcc.astro.washington.edu/faculty/trq/brandon/pkdgrav.html
[9] http://www-hpcc.astro.washington.edu/tools/changa.html
[10] http://irfu.cea.fr/Projets/Site_ramses/RAMSES.html
[11] Springel V., 2009, Mon. Not. of the R. Astron. Soc., in press
[12] Avillez, M.A. & Breitschwerdt, D. 2005, Astronomy & Astrophysics, 436, 585
[13] Dai, W., & Woodward, P.R. 1994, Journal Computational Physics,
[14] Dai, W., & Woodward, P.R. 1998, Astrophysical Journal, 494, 317
[15] D.S. Balsara, J. Comput. Phys., Vol. 174(2), pp. 614-648 (2001).
[16] Avillez, M.A., Spitoni, E., & Breitschwerdt, D. 2011, in "Advances in Computational Astrophysics: Methods, Tools and Outcomes", ASP Conf. Series., In press.
[17] http://piernik.astri.umk.pl
[18] Hanasz, M., Kowalik, K., Woltanski, D., & Pawlaszek, R. 2010, in “Torun Exoplanets”, EAS Publications Series, 42, 275
[19] http://plutocode.ph.unito.it
[20] http://nirvana-code.aip.de
[21] http://www.astrosim.net/

3.2 Climate

Climate change is one of the global challenges of our time: it has been, and is still extensively researched, and the overwhelming majority of climate scientists agree that the observed modern day global warming is unprecedented and is caused in its largest part by humans [1]. Still, from a more quantitative point of view, the magnitude, feedbacks, and manifestations of climate change are subject to deep uncertainty.

While climate change creates great risks for humanity, these risks can be minimised if the uncertainty can be reduced, or at least quantified. Climate modelling currently represents the single best method to predict the future, and allows humankind to adapt the way we live. Based on their cousins, the numerical weather prediction models, for their atmospheric component, climate models include the long-term interactions with ocean, land, chemistry, … as well as human and non-human influences (so-called forcing terms, e.g. the increase in greenhouse gases). The central limiting factors to the validity and efficacy of a climate model are, on the one hand, its completeness in terms of physical processes, and, in the second hand, its resolution, both of which are ultimately limited by the computing resources available. With increasing resolution, physical phenomena are being incorporated with more significant accuracy, such as the dynamics of clouds, the transport of heat, moisture, and of biogeochemical components such as ozone and other greenhouse gases, which in turn affect radiative processes and the global heat budget of the system [2].

This complexity explains why climate modelling is now one of the most compute-intensive areas in science, and one where new and emerging high-performance computing technologies could have a decisive effect in improving the answers to important climate questions, as for example: how will climate change effect communities and individuals on a local scale? who will benefit? who will be disadvantaged? how can we adapt to the changes?

Climate science thus meets the criteria defined in section 2: there is a clear impact on society, a long history of relevance to supercomputing, and, finally, a documented willingness of the developers to adapt and modify their codes to make use of new technological developments. Further, the spatial and physical interconnectedness of the problem bring particular challenges to model scalability.

3.2.1 Scientific Challenges

Climate models have gone through a revolution in the last decade, enabled on the one hand through increases in computing power, and, on the other hand, through improvements in algorithms and inclusion of new physical processes. A central achievement has been the steady increase in both spatial and time resolution for climate simulations. Such improvements, clearly allow for demonstrably better forecasts [3]. An example of the benefits of higher resolution can be seen in Fig. 3. Another important one has been the introduction of more complexity with the inclusion of biogeochemistry processes, such as the carbon cycle, included in atmosphere-ocean-land components, turning climate models into Earth system models (ESMs).

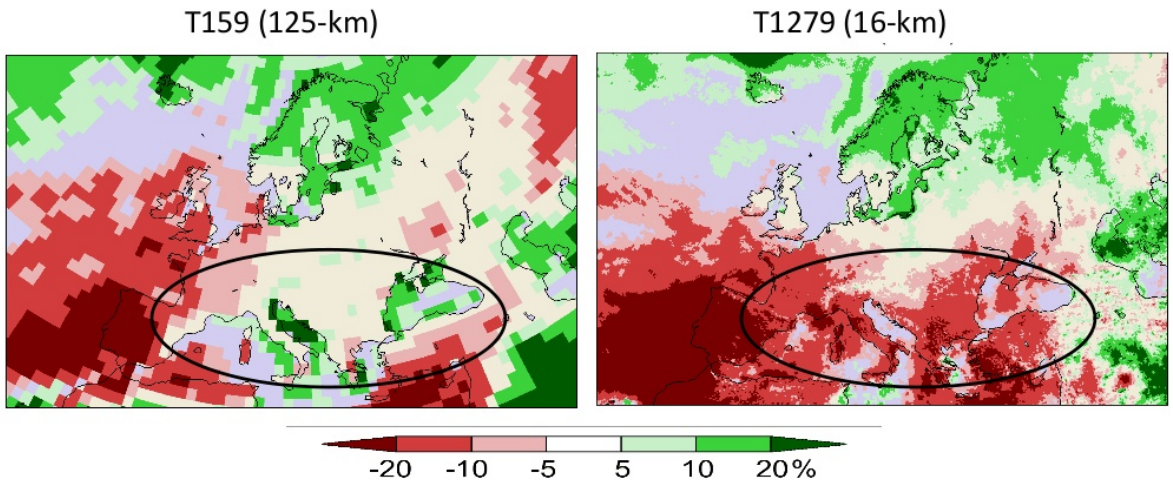

Figure 5. Most of our experience indicates that climate change will lead to a reduction of precipitation over the Mediterranean region. However, a low-resolution simulation (left: 125km) with the ECMWF IFS model over the 21st century seems to indicate a slight increase in the annual mean precipitation over the indicated portion of the Mediterranean basin. A much higher resolution simulation (right: 16km) clearly predicts a sharp (5-20%) decrease over the century, in line with most other models. [3]

In the context of WP8, we restrict ourselves to two key issues. Firstly, mature models need to be ported, validated, optimised and supported on the next generation of high-performance platforms, and run at high resolution to gain insight into key climate questions. Numerous technical challenges need to be overcome to achieve this, from the coupling of model components, to parallelising the model output, to the efficient execution of ensembles. This requires access to not only significant computing resources, but also post processing and data archiving facilities.

The second issue is the development of new algorithms and techniques for the models of the future. For example, the newest dynamical cores (atmospheric dynamics solvers, but also ocean dynamics solvers) use novel grids (Fig. 4) to span the Earth. The homogeneity of these grids generally results in improved parallel performance, at the price of additional programming complexity. Several of these emerging dynamical cores have already proved their viability [4], in full-scale models. Besides the basic development of these algorithms on HPC platforms, other challenges include the incorporation of (pre-existing) physical parameterisations to achieve the necessary scalability.

For both of these areas, a key challenge is then to port the codes to disruptive architectures, such as general-purpose graphical processing units (GPGPUs), e.g., [5]. This step requires extensive interaction between the algorithm developers and the experts in the programming paradigms for these platforms in a way that allows us to maintain a good pace of scientific development and does not lock us into single architectural solutions.

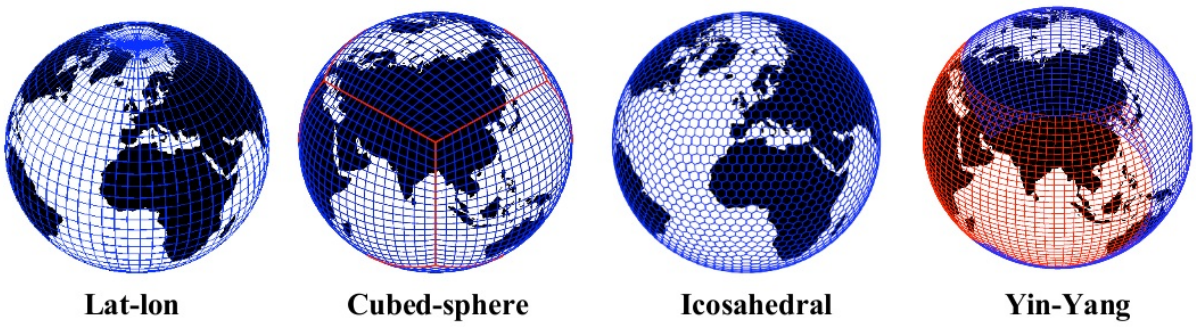

Figure 6. Classically most dynamical cores have used a latitude-longitude grid for calculation, which is computationally practical, but highly inhomogeneous. In particular, the meridians converge at the poles, which requires special techniques to avoid computational instabilities. Many new models are attempting to use grids with a more uniform coverage, e.g., gnomic (cubed-sphere), icosahedral (consisting of either triangular or penta-/hexagonal cells), or two overlapping, but carefully joined, regular grids ("Yin Yang"). [6]

3.2.2 Numerical Approaches and Codes

The European Network for Earth System modelling (ENES) gathers the European climate modelling community and organises its modelling infrastructure within the FP7 project IS-ENES (http://is.enes.org). Within the ENES initiative, the EC-EARTH [7] and HadGEM3-A (Unified Model) [8] models have been selected as the first models to benefit from a cooperation with PRACE-1IP (EC-Earth) and DEISA2 (HADGEM3-A) to prepare for and use Tier-0 machines. Both teams would need technical help with issues such as ensemble runs, code coupling, process mapping, and the management of large output data. But progress achieved this way will also benefit to other models using a spectral dynamical core: (i) the ECHAM-6 [9] model although the ECHAM developers at Max Planck Institute for Meteorology have now channelled development of the Icosahedral Non-hydrostatic, or ICON, [10] model, mentioned subsequently; (ii) the French "ARPEGE-Climat" model, sharing with IFS, ECHAM and EC-Earth most of the features of its spectral dynamical core. Other models, based on the grid-point approach for the dynamical core, such as the Hadley Centre's HadGEM3-A model [8] and the IPSL LMD-Z model [11], will be regularly and fully informed of the progress, so that they can benefit from them in their own concomitant developments.

Instead of concentrating on one particular model, a larger impact can be made by concentrating on community-developed libraries, which can be used in multiple models. The areas currently creating bottlenecks are known to be I/O and couplers between different model components. Of particular significance in this context is the ScalES [ScalES] initiative. One of its projects is to create a portable, parallel I/O library within the framework of Climate Data Interface (CDI). This will offer ECHAM6, and subsequently ICON, parallel I/O capability. The community effort for coupling model components is OASIS4, which is also being extended within another ScalES project.

Within the context of forward-looking models, several emerging dynamical cores (solvers of the atmospheric dynamics) use icosahedral grids to span the earth. A G8-funded ICOMEX concentrates on four of these cores: MPAS, NICAM, ICON, and DYNAMICO, summarised in [6], where the latter two are European developments. ICON is close to official release and has already been evaluated in an inter-dycore comparison [5]. The UK MetOffice, STFC and other UK academics are similarly working on the "GUNG HO!" project and IPSL on the DYNAMICO project, both are still in their first development phase.

Currently, several atmospheric models are being ported to GPGPUs, for example, the global models ASUCA (Tokyo University) and the Non-hydrostatic Icosahedral model (NIM,

National Oceanic & Atmospheric Administration), as well as regional models, e.g., the Weather Research and Forecasting model (WRF) and the Consortium for Small-scale Modelling *(COSMO)* model.

3.2.3 Community involvement, expected outcomes and their impact

In this section we describe proposed areas of work in collaboration between the ENES community and the PRACE-2IP partners. They correspond to important, common issues for the climate modelling community and have been discussed during a meeting organised by the IS-ENES infrastructure project which has among its objectives to organise the interface with PRACE. Further discussions with ENES, possibly through a workshop between ENES and PRACE-2IP partners, will allow to make the final choice between the priority issues described below, and provide the rationale for affecting resources.

(i) Coupler

The OASIS coupler is used today by more than 30 climate modelling groups around the world. In Europe in particular, OASIS is the coupling software used in 6 of the 7 Earth System Models (ESMs) involved in ENES. This strong community network around the coupler results from the permanent investment of CERFACS (and therefore its shareholder Météo-France) and CNRS, who are both devoting one person full time to OASIS, from numerous collaborations on specific developments and from temporary but important funding (PRISM, CICLE, METAFOR, IS-ENES projects).

The current OASIS3 version of the coupler offers a limited field-by-field parallelism of the coupling, which is becoming a bottleneck for high-resolution ESMs run on massively parallel platforms. To remove this bottleneck, two short-term solutions are currently under evaluation. The first one is based on a user-defined regridding functionality recently developed in IS-ENES and the second one is to interface the American Model Coupling Toolkit (MCT) with the OASIS3 communication library. These solutions both assume that the regridding weights are pre-computed offline and implement fully parallel regridding and exchanges of the coupling fields. If these solutions both prove to be valid, the most efficient one will be chosen. On the longer term, we are also evaluating the other coupler currently developed at CERFACS, Open-PALM and its new parallel online regridding library CWIPI. In fact, online calculation of the regridding weights will most probably become a clear requirement on the longer (~5 years) term, in particular for adaptive grids.

In PRACE-2IP WP8 the focus will be on the optimisation of the best of the short-term solutions described above and on the development of a fully parallel on-line interpolation library suited to current and future climate modelling needs.

(ii) I/O

Several groups in Europe and elsewhere are currently aiming at increasing the resolution of their models as much as the machines in use, and the software developed, allow. One of the common obstacles for optimal performance of such models is the I/O. This has been recognised for a long time, and e.g. the IS-ENES project is trying to tackle this problem. The IO issue includes the development of high performance, asynchronous I/O servers within model simulations and the I/O within post-processing itself. In the first instance, we propose to deal with the ocean European platform NEMO widely used in Europe and for post-processing with the Climate Data Operator used by many European groups.

NEMO: NEMO has been interfaced with an IO server for 2 years. A new version of this IO server has been completely rewritten from scratch in order to solve known performance issues. This new I/O server has been tested with small configurations. It is now crucial to validate its use and its performance on high resolution configurations parallelised on about

1000 cores. These tests must also deal with parallelisation and compression functionalities offered by NetCDF4, the file format on which this I/O server is based. These tests have to be performed on several computers in order to evaluate the portability, the performance and the optimal use of the IO server according to the different architectures of the PRACE machines. Tests will also be needed in configurations coupled to other components with the OASIS coupler. Possible interface with other data formats should also be examined (e.g. using CDIs see below).

Post-processing: post-processing, which encompasses the scientific analysis and evaluation of the model results, on-line or after execution, often comprises routines that are not fully optimised for performance. One of the packages used for this process are the Climate Data Operators. The CDOs are based on a library called CDIs (climate data interfaces), which is handling the I/O of the codes. This library is also used by some of the most modern model codes around today, e.g. the ICON model. CDIs could also be used by other models to incorporate the functionality provided by the CDOs. So, improving the I/O of the CDIs has the potential to be crucial for a larger number of climate prediction centres in Europe.

For PRACE-2IP, we propose to examine the basic structure of the CDI and to identify those portions of the code that can be improved in their performance using e.g. one of the I/O methods developed in IS-ENES. Many institutions would thus benefit from a performance improvement of the CDIs for their post-processing, and could possibly benefit for their models, also.

(iii) Dynamical cores

Several groups in Europe and elsewhere are currently developing new dynamical cores for atmosphere and ocean models to be used for global and regional climate research as well as for numerical weather prediction. For physical and numerical reasons, these new models are based on unstructured, in particular icosahedral, grids.

Probably the most advanced of the models in Europe is the ICON model, which is developed jointly by the German Weather Service (DWD) and the Max-Planck-Institute for Meteorology (MPI-M). In a cooperation with ETH Zürich and CSCS parts of ICON have been ported to GPGPU architectures. Within the G8 project ICOMEX, ICON will be used as a test case for an abstract model description scheme for efficient use of memory bandwidth.

For PRACE-2IP WP8, we propose to examine the basic structure of ICON and to identify parts that can be modularised and abstracted in order to increase performance, portability and maintainability of this type of models. ICON is chosen because the code is basically operational but in this context serves as a sandbox to develop and test solutions that may benefit any other model of this type. Thus other endeavors within ENES, e.g. the UK project Gung-Ho and the French project DYNAMICO, currently in an evaluation phase, can build upon the results.

(iv) Fault-tolerance

First adaptations of NEMO ocean model algorithm to fault tolerant (FT-)MPI [12] are currently implemented at CERFACS. They will have to be tested on an environment able to reproduce failures. PRACE-2IP, could offer such experimental conditions on one of their future platforms.

3.2.4 Relevant Bibliography

[1] Pachauri, R.K. and A. Reisinger (Eds.). Core Writing Team. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. IPCC, Geneva. 2007

[2] Drake, J.B., P.W. Jones and G. R. Carr. Overview of the Software Design of the CCSM. Int. J. High Perf. Comp. Appl. 19 (3), 177—186, 2005.

[3] Dirmeyer, P., B. Cash, J. Kinter, T. Jung, L. Marx, M. Satoh, C. Stan, H. Tomita, P. Towers, N. Wedi, D. Achuthavarier, J. Adams, E. Altshuler, B. Huang, E. Jin, and J. Manganello. Climate Dynamics On-line. Jun. 2011. Springer-Verlag. http://www.springerlink.com/content/9256820kjx800015/

[4] Lauritzen, P., C. Jablonowski, M. Taylor and R. D. Nair. Rotated versions of the Jablonowski steady-state and baroclinic wave test cases: A dynamical core intercomparison. J. Adv. Model. Earth Systems. Vol. 2, Art. #15, 34 pp., doi:10.3894/JAMES.2010.2.15. 2010.

[5] Govett, M.W., J. Middlecoff and T. Henderson. Running the NIM Next-Generation Weather Model on GPUs. Cluster Computing and the Grid, IEEE International Symposium on, pp. 792-796, 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, 2010.

[6] Lauritzen, P. H. and D. L. Williamson. Dynamical core development. Atmosphere Model Working Group (AMWG) winter meeting, NCAR, Boulder, Colorado. February 12, 2010.

[7] Hazeleger, W., C. Severijns, T. Semmler, S. Ştefănescu, S. Yang, X. Wang, K. Wyser, E. Dutra, J. M. Baldasano, R. Bintanja, P. Bougeault, R. Caballero, A. M. L. Ekman, J. H. Christensen, B. van den Hurk, P. Jimenez, C. Jones, P. Kållberg, T. Koenigk, R. McGrath, P. Miranda, T. Van Noije, T. Palmer, J. A. Parodi, T. Schmith, F. Selten, T. Storelvmo, A. Sterl, H. Tapamo, M. Vancoppenolle, P. Viterbo, U. Willén. EC-Earth: A Seamless Earth-System Prediction Approach in Action. BAMS Issue 10, Vol. 91, pp. 1357–1363, Oct. 2010.

[8] D. N. Walters and M. J. Best and A. C. Bushell and D. Copsey and J. M. Edwards and P. D. Falloon and C. M. Harris and A. P. Lock and J. C. Manners and C. J. Morcrette and M. J. Roberts and R. A. Stratton and S. Webster and J. M. Wilkinson and M. R. Willett and I. A. Boutle and P. D. Earnshaw and P. G. Hill and C. MacLachlan and G. M. Martin and W. Moufouma-Okia and M. D. Palmer and J. C. Petch and G. G. Rooney and A. A. Scaife and K. D. Williams The Met Office Unified Model Global Atmosphere 3.0/3.1 and JULES Global Land 3.0/3.1 configurations Geosci. Model Devel. 2011 4 1213--1271

[9] Roeckner, E., G. Baeuml, L. Bonaventura, R. Brokopf, M. Esch, M. Giorgetta, S. Hagemann, I. Kirchner, L. Kornblueh, E. Manzini, A. Rhodin, U. Schlese, U. Schulzweida, A. Tompkins, The atmospheric general circulation model ECHAM5 - Part 1, MPI-Report 349, 127 pp, 2003.

[10] Giorgetta, M., T. Hundertmark, P. Korn, S. Reich, M. Restelli, Conservative space and time regularizations for the ICON model, Berichte zur Erdsystemforschung, 67, 32 pp., 2009.

[11] Marti, O. and 28 co-authors, Key features oft he IPSL ocean atmosphere model and ist sensituvity to atmospheric resolution, Climate Dynamics, 34, pp1-26, 200.

[12] Fagg, G., Gabriel, E., Bosilca, G., Angskun, T., Chen, Z., Pjesivac-Grbovic, J., London, K., Dongarra, J. Proceedings of ISC2004, Heidelberg, Germany, June 23, 2004.

3.3 Material Science

3.3.1 Introduction and Scientific Challenges

In contrast to traditional areas of physics such as cosmology or high-energy physics, where a few problems (e.g. dark matter, validation of the standard model) consume large parts of the respective communities, materials science consists of many big problems (e.g. photovoltaics, superconductivity, magnetism, nanoscience in various flavours, microelectronics, etc.) that have direct impact on technology and thus society, but the field may appear defocused to the uninitiated.

In particular, electronic structure simulations have emerged as seminal tools for scientific explorations in materials science. This is due to a powerful combination of theoretical developments (in which Europe has been leading the way), new algorithms, and the stunning increases in computer performance as all scales (from workstations to supercomputers). The impact of this development is measurable with a simple metric: 16 of the 20 most cited papers in the history of the American Physical Society (all Physical Review Journals of the APS starting in 1893 [1]) are concerned with quantum mechanical simulations of condensed matter, atoms, and molecules.

Electronic structure simulations of materials are based on first principles, i.e. they do not depend on empirical parameters and are thus predictive. Modern methods based on density functional theory (DFT) [2,3] are very reliable in describing ground state properties of most classes of materials, such as structure, magnetic phases, and as well as basic electron transport properties within linear response theory. In fact, many of the advances in semiconductor electronics are due to the success of band structure theory and corresponding electronic structure simulations of semiconductor heterostructures.

On the other hand, there are classes of materials, such as strongly correlated transition metal oxides, and problems, such as electron transport far from equilibrium or the proper description of electronic excitations in spectroscopy and during chemical reactions, where traditional DFT based electronic structure methods fail systematically. During the past decade, new methods have emerged that can tackle these problems, however there come at a tremendous computational expense. Fortunately, these new methods map well on to petascale systems, both in terms of scalability and efficiency. In fact, the most efficient codes on petascale systems today all fall in this category, as is demonstrated by a series of Gordon Bell Prize finalists and winners in recent years [4,5,6,7,8].

The European electronic structure community is leading the way in terms of theory and community code development [9,10]. However, the ideas that made quantum simulations so successful on petascale computing systems have not yet made it into the European community codes. The main goal of the service activities of this work package will be to bring this high-end technology into community codes and to prepare these codes for the coming age of multi-core and hybrid multi-core/GPU systems.

3.3.2 Numerical Approaches and Codes

Modern electronics structure codes in materials science build on Kohn-Sham Density Functional Theory (KS-DFT) [2] and some approximation to the exchange and correlation energy, such as the Local Spin Density Approximation (LSDA) or the Generalised Gradient Approximation (GGA) [11]. Recent developments in orbital dependent and time dependent density functional theory (TD-DFT) have also found their way into community codes and are routinely used to study electronic excitation spectra of materials, nanoscale systems, and molecules [12]. The codes are used by theorists, simulation experts, as well as experimentalists, and thus have to be robust, easy to use, and portable.

The numerical methods used in these codes are mostly determined by the desired accuracy and the material type on is investigating. For example, plane wave pseudo potential methods (PWPP) [11,13] are used when moderate accuracy is required, and localised core and semi-core states can be absorbed into a pseudo potential, thereby smoothening the potential and drastically reducing the number of plane waves in the expansion of the wave functions. When the accuracy needs are low, one can typically use numerical basis sets that a localised in space (but are often not systematic in the way plane waves are), which leads to sparse problems that can be solved in order *N* time (most electronic structure methods scale with the cube of the number of electrons). On the other hand, when the desired precision is high, even the use of pseudo potentials can be problematic and one has to resort to all-electron methods, where localised electrons feature in the calculations and expansions in plane waves are no longer feasible. Spherical harmonic expansions around the nuclei that are augmented with plane waves in the intermediate region between atoms are used instead – the linearised augmented plane wave LAPW method [13] has emerged as the gold standard in electronic structure simulations.

All these methods rely on the effective use of (in most cases dense) linear solvers. The optimal algorithms may vary, as the basis sets of the various methods differ, i.e. the fractions of eigenvalues and eigenvectors needed to represent the electronic structure depends on the method.

In plane wave codes, the basis set is prohibitively large and the entire Hamiltonian matrix cannot be stored at any given time. Instead, one uses (distributed) Fourier transforms to switch between real space and Fourier space representations, in order to compute different contributions to the Hamiltonian in the most efficient way. One extracts a sub-space with the electron orbitals of interest, in which a sub-space diagonalisation performed, where typically more than 50% of the eigenvalues/eigenvectors have to be computed. Subsequent iterative solves, the distributed Fourier transforms, and the sub-space diagonalisation, constitute the numerical bottlenecks of the method. With an ingenious reorganisation of the way how orbital and plane wave expansion have been distributed over a parallel supercomputer, Kent [14] was able to run large simulations using the standard electronic structure package VASP with 5000 MPI ranks on productive simulations already in 2007. It turned out the sub-space diagonalisation becomes a major computational/scalability bottleneck when scaling the calculations to thousands of nodes. This calls for more efficient multithreaded eigenvalue solvers.

In the LAPW method, the basis set is much more efficient and the entire Hamiltonian of the system can be computed and stored in a uniform way, i.e. there is not need for distributed transforms with global communications as in the plane wave based methods. The basis, however, is not orthogonal and one is left with solving a generalised eigenvalue problem for a Hermitian matrix. Depending on the required precision, between 50% and 10% of the eigenvalues/eigenvectors have to be computed. In extreme cases, where benchmark results used by other methods are computed, the chosen basis set may be so large than only about 1% of the eigenvalue/eigenvectors may be needed to compute the desired electronic structure.

In most other methods, the basis set is more efficient than in the plane wave and LAPW methods just discussed, although at the expense of precision. That is, for a given problem size, the Hamiltonian matrix will be smaller and a larger fraction of the eigenvalues/eigenvectors (typically 40-60%) has to be computed to describe the electronic structure.

Studies have shown [15] that in most cases direct solvers are preferred for electronic structure methods, although the algorithm of choice (e.g. divide an conquer, MRRR, etc.) may vary depending on the method and basis set size. Most electronic structure codes today simply rely

on LAPACK or ScaLAPACK for these solvers. In the coming age of multi-core and hybrid multi-core/GPU systems it is thus desirable to use new algorithms and library implementations that are adapted to highly multithreaded nodes. Furthermore, it will be equally important that the successful strategy to run standard electronic structure codes like VASP at ORNL on many thousands of MPI ranks be introduced into the main European plane wave pseudo potential codes, such as ABINIT and Quantum Espresso.

While the eigenvalue problems and dense linear solves play a fundamental role in almost all modern electronic structure packages, the introduction of TD-DFT and self-consistent GW methods to describe electronic excitations lead to the need to solve integral equations, which, as it turns out, lead to large distributed matrix multiplication problems. Kozhevnikov et al. [Kozhevnikov] demonstrated how the screened Coulomb interaction, W expressed within the LAPW basis set can be computed at scale with very high efficiency. While the work was originally intended at the computation of frequency dependent Hubbard interactions, the real interest in the implementation is its use for the computationally demanding part within self-consistent GW calculations in European software packages such as EXCITING. The combination of these novel implementations of the GW method with very large supercomputers will open up new avenues in electronic structure simulations.

3.3.3 Community involvement, expected outcomes and their impact

The electronic structure community in Europe has been organised in the PsiK network, which started in the early 1990s. While funding from the European Commission (EC) for the network ended many years ago, it continues to strive with contributions from individual chairs across Europe. A significant fraction of active developers of the European electronic structure community is participating in the EC funded European Theoretical Spectroscopy Facility (ETSF), which will be one of the main communities driving the activities of this part to WP8.

The ETSF is (see etsf.eu) “a knowledge centre for theoretical spectroscopy carrying out state-of-the-art research on theoretical and computational methods for studying electronic and optical properties of materials. The ETSF gathers the experience and know-how of more than 200 researchers, facilitating collaborations and rapid knowledge transfer. The ETSF offers its experience to researchers, industry and students in the form of collaborative projects, free scientific software and training.”

It is be this free software that will be targeted by the materials science service activities of the WP8 of PRACE 2IP, in particular the ABINIT package, which includes a PWPP code and BigDFT (a wavelet based code), and the EXCITING package, which is based on the LAPW method. Professors Xavier Gonze of the Catholic University of Louvain, Angel Rubio of the University of the Basque Country, and Claudia Ambrosch-Draxl of the University of Leoben are supporting the efforts from the ETSF side.

A second significant community has formed around Quantum ESPRESSO, a distributed multinational initiative, open to all those who want to contribute to develop it further, or simply use it. Current and past developers are mainly based at the Scuola Internazionale Superiore di Studi Avanzati (SISSA), the Abdus Salam International Centre for Theoretical Physics (ICTP), and CNR-IOM DEMOCRITOS Simulation Center in Trieste; at the Princeton University; at the University of Oxford; at the Ecole Polytechnique Fédérale de Lausanne (EPFL); at the Università di Udine; at the University of California at Davis; at the University of Minnesota; at the CINECA supercomputing centre in Bologna; at the ICHEC, Irish national supercomputing centre; at the Université Pierre et Marie Curie in Paris.

The QE user community consists of hundreds of users worldwide. Relevant person from representative institutions of the community such as Prof. Nicola Marzari, director of the Laboratory for Theory and Simulations of Materials (THEOS) at EPFL; Prof. Stefano Baroni

head of condense matter department at SISSA; Prof. Paolo Giannozzi from Udine University (Italy) and Prof. Stefano De Gironcoli from DEMOCRITOS expressed significant interest in the goals of this project.

The following potential activities will be further investigated and a work plan will be developed in the coming months:

(1) Scaling up of the PWPP codes in ABINIT and QE by reorganizing the way plane waves and orbitals are spread over distributed memory machines according to the prescription given by Kent [Kent]. The expected outcome is significant improvement in scalability of PWPP codes.

(2) Incorporate novel development in eigenvalue solver technology [Solcà], which has recently been implemented into MAGMA in an ongoing collaboration between the Innovative Computing Laboratory at University of Tennessee in Knoxville and the Computational Materials Group at ETH Zürich, into all above-mentioned electronic structure packages mentioned above. The expected outcome will be significant improvement of time to solution on all multi-core and hybrid multi-core/GPU platforms at all scales (from clusters to large supercomputers).

(3) Explore developing basic libraries for PWPP and LAPW with the goal of reducing future porting works of electronic structure software packages to new architectures. A further expected outcome will be that other LAPW based packages besides EXCITING can adopt the work done in this work package with reasonable programming efforts by the respective communities.

(4) Explore using the distributed matrix techniques used by Kozhevnikov and co workers to compute the screened Coulomb interaction in GW and TD-DFT methods that are implemented in all of the above-mentioned packages. The techniques are directly transferable to the EXCITING package. Applicability to ABINIT and QE will have to be studied. The expected outcome will be a drastically improved scalability and efficiency of excited state electronic structure simulations.

The work performed in this part of WP8 is not necessarily limited to the above mentioned software packages, and may be extended to other packages, depending on effort and available manpower. Furthermore, depending on work performed for the climate community, we will explore leveraging structured grid type solver libraries for PDEs in the OCTOPUS software package, which is a multi-grid based TD-DFT package within the ETSF.

3.3.4 Relevant Bibliography

[1] http://mml.materials.ox.ac.uk/Main/Top20

[2] Kohn, W., and Sham, L. J., Self-consistent equations including exchange and correlation effects, Phys. Rev. 140, A1133 (1965).

[3] Kohn, W., Nobel Lecture: electronic structure and matter – wave functions and density functionals, Rev. Mod. Phys. 71, 1253 (1998).

[4] Alvarez, G., Summers, S., Maxwell, D. E., Eisenbach, M., Meredith, J. S., Larkin, J. M., Levesque, J. M., Maier, T. A., Kent, P. R., D'Azevedo, E. and Schulthess, T. C., New algorithm to enable 400+ TFlop/s sustained performance in simulations of disorder effects in high-Tc superconductors, Proceedings of the ACM SC08 Conference (2008).

[5] Wang, L.-W., Lee, B., Shan, H., Zhao, Z. Meza, J., Strohmaier, E. Balley, D., Linear scaling divide-and-conquer electronic structure calculations for thousand atom nanostructures, Proceedings of the ACM SC08 Conference (2008).

[6] Eisenbach, M. Zhou, C., Nicholson, D. M., Brown, G., Larkin, J., and Schulthess, T. C., A scalable method for abinitio computation of free energies in nanoscale systems, Proceedings of the ACM SC09 Conference (2009).
[7] Apra, E., Harrison, R. J., de Jong, W. A., Rendell, A., Tipparaju, V., and Xanthess, S., Proceedings of the ACM SC09 Conference (2009).
[8] Kozhevnikov, A., Eguiliuz, A. G., and Schulthess, T. C., Toward first principles electronic structure simulations of excited states and strong correlations in nano- and materials science, Proceedings of the ACM SC10 Conference (2010).
[9] http://www.abinit.org
[10] http://www.quantum-espresso.org
[11] Martin, R. M., Electronic structure: basic theory and practical methods, Cambridge (2004).
[12] http://www.etsf.eu
[13] Singh, D. J. and Nordstrom, L., Planewaves, pseudopentials, and the LAPW method, Springer (2006).
[14] Kent, P. R. C., Computational Challenges of Large-Scale Long-Time First-Principles Molecular Dynamics, J. Phys.: Conf. Series 125, 012058 (2008).
[15] Auckenthaler, T., Blum, V., Bungartz, H. J., Huckle, T., Jahanni, R., Krämer, L., Lang, B., Lederer, H., and Willems, P. R., Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations, preprint (2010).

3.4 Nuclear and Particle Physics-Lattice QCD

Strong interaction phenomena are manifest in areas that at first sight look unrelated such as the early universe, the formation of all the elements and stars and supernovae explosions, the asymmetry of matter and antimatter in our universe, quark matter, the formation and stability of atomic nuclei that build the atoms and molecules forming the familiar matter on which life originates. Sophisticated experiments are devoted to the understanding of these complex phenomena at large experimental facilities such as the LHC at CERN, RHIC and Jefferson Lab in the US, GSI in Germany and SuperB in Italy. A number of experimental searches of new physics require good understanding of the strong interactions that in turn demands considerable computational power for carrying out calculations to address key questions that may lead to such an understanding. The design of accelerators and detectors for carrying out these experiments and analysis of the results also require large computer resources [1] (note that all the references in section 3.4 and subsections, refers to section 3.4.4).

A particularly demanding area of nuclear and particle physics requiring large scale simulations is the non-perturbative solution of Quantum Chromodynamics (QCD), the fundamental theory of the strong interactions. Lattice QCD is the only known *ab initio* method for performing QCD calculations in the low-energy regime, and for acquiring a quantitative description of the physics of hadrons and nuclear forces. It involves a direct numerical evaluation of the path-integral of QCD using large-scale computers.

A number of fundamental issues are being addressed within lattice QCD using advanced algorithms and current supercomputers, such as the hadron spectrum and structure, the calculation of meson and baryon interactions, the calculation of hadron matrix elements for weak processes, the phase structure of strongly interacting matter and the study of theories beyond the standard model.

3.4.1 Scientific Challenges

Nuclear and Particles physics is an extremely broad scientific area. In this work package we will focus on lattice QCD.

Lattice QCD

After the successful calculation of the hadronic mass spectrum and the zero-density phase structure of QCD, lattice QCD aspires to address more complex problems such as hadron structure and transitions, resonances, nuclear forces and systems of multiple hadrons in the coming years. These problems require large 4-dimensional lattices, of the order of a few hundred sites in each direction. The statistical signal in such calculations is inherently noisier, therefore, a multiple of the statistics currently used will be required.

Furthermore, discretised formulations of QCD, which preserve chiral symmetry exactly on the lattice are also expected to become more widely used in the coming years as computer power increases. Exact chiral symmetry can either be achieved through the Domain Wall formulation, where an additional dimension is required, or through the so-called Overlap formulation, where the sparcity of the fermion operator is lost. In either case these two equivalent formulations require an order of magnitude more compute power than formulations, which break chiral symmetry explicitly.

The Physic program that can be pursued using large-scale simulation of QCD is very board. A few examples are:

- The evolution of the universe, from its primordial hot and dense state (10^{-32} sec after the Big Bang) to the formation of stars and galaxies to supernovae and star collapse, requires knowledge of the strong interactions. Lattice QCD contributes in the theoretical

understanding of the quark-gluon plasma, the state of the universe at time 10^{-32}s, the description of hadrons that formed at time 10^{-6}s and the abundance of matter to antimatter. The determination of the phase diagram of QCD that covers the transition from the quark gluon plasma to ordinary hadron matter requires large scale computing.

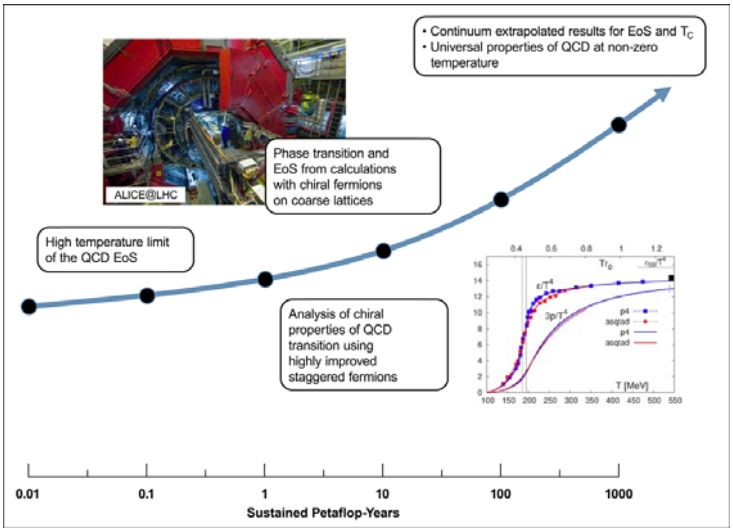

Figure 7. Computational resources required to study the phase structure of QCD.

- The formation of heavier nuclei elements, solar fusion, type-II supernovae, neutron formation stars require knowledge of the nuclear force and the equation of state, information that can be obtained from lattice QCD.

Figure 8. Stellar evolution.

The determination of the nucleon-nucleon force from lattice QCD simultion is now feasible and pioneering studies have aready been performed [2]. Two- and three- nucleon systems have also been studied. The determination of the nuclear force requires going beyond two-boay intractions to three- and four-body forces, which requires large computational capabilities. Exa-scale computing opens the possibility of a unified description of nuclear physics from the underlying theory.

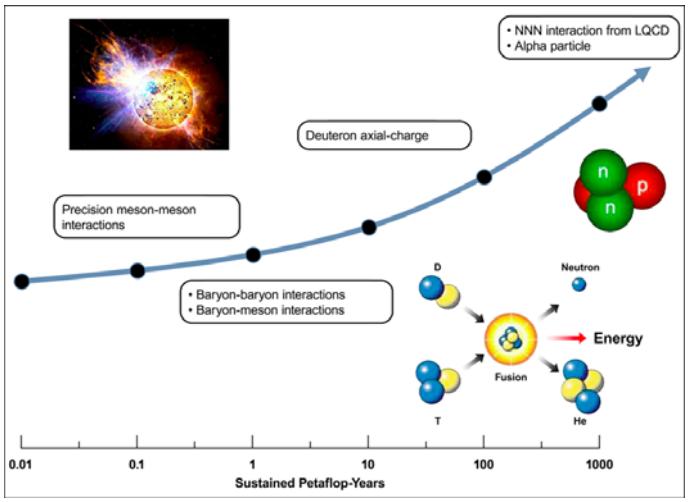

Figure 9. Resources required to calculate nuclear forces.

- The theoretical understanding of the structure of the nucleon and other hadrons that make up the bulk of visible matter in the universe requires numerical solution of the fundamental theory of the strong interactions. Precision results on the nucleon structure functions and form factors, on hybrids and exotics, and on the matter-antimatter arising from simulations of QCD will benefit the experimental program of major labs.

- With the operation of the Large Hadron Collider (LHC) at CERN, there is a new era in Particle Physics and lattice QCD can play an important role. The search for new physics at the intensity frontier requires performing ever more stringent tests of Standard Model predictions ito determine whether the experimental results are consistent with the Standard Model. This requires precise calculations in QCD, and for this numerical simulations of lattice QCD are needed. The ultimate aim is to reduce the error in the lattice QCD calculations to a level significantly below that in the corresponding experimental measurement. Searching for new physics in this indirect manner allows one to reach energy scales that are comparable to, or in some cases exceed, those that are directly accessible at the LHC. A related, though different, role that lattice QCD will play occurs once evidence for new physics. This evidence will likely be consistent with many possible models, and simulation technique similar to those used in QCD will play a crucial role in its interpretation. This is because new physics leads to contributions to rare Standard Model processes that must be consistent with experimental results. As for the Standard Model contributions themselves, one needs non-perturbative calculations in QCD to determine the contribution of a given model of new physics. Once again, lattice QCD is the only tool for many of these calculations. Similar input will be needed in some searches for dark matter at the cosmic frontier. The coupling of candidate dark matter particles to ordinary matter depends on non-perturbative QCD physics. Lattice QCD is thus needed as an integral component of direct dark-matter searches [3].

3.4.2 Numerical Approaches and Codes

Lattice QCD allows for a quantitative study of the strong interaction at energy scales where perturbative methods fail. Observables are evaluated numerically by evaluating the Feynman path integral with Monte Carlo methods. In the path integral, the action appears in the exponent of the probability weight, much like the Hamiltonian appears in the exponent of the Boltzman weight in statistical systems. In order for this probability to be evaluated numerically, the action is discretised on a four-dimensional Euclidean space-time grid, or lattice. One proceeds by generating representative sets of QCD gauge configuration

backgrounds via Monte Carlo methods, and subsequently by evaluating the statistical average and error of observables of interest using these representative configurations. Therefore lattice QCD allows for a first-principles, quantitative study of strong-interaction phenomena, since the QCD action of the theory is directly simulated. Since numerical simulation requires discretisation of the continuum theory a number of different choices exist associated with different discretisation effects, or lattice artefacts, depending on the discretisation scheme.

A typical calculation in lattice QCD consists of firstly generating the gauge configuration backgrounds, which involves carrying out the Monte Carlo simulation, and subsequently evaluating observables of interest on the stored gauge configurations. In both these steps, a large portion of the computational effort (around 80% of the wall-time) is spent on inverting the so-called Fermion-Dirac operator. Although the discretised Fermion-Dirac operator can be represented as a large sparse matrix, its inverse is dense and computation of all the elements of the inverse is unfeasible, both in terms of computer time and memory required. Usually one can get away with having available just the action of the inverse of this operator on a vector, which can be computed via iterative solver methods, such as e.g. the Conjugate Gradient method. These iterative methods consist of repeatedly applying the operator on a trial vector, and updating the trial vector accordingly. Optimisation of a QCD application code therefore begins by implementing an optimised version of the discretised Fermion-Dirac operator. Such an optimised implementation is typically tuned to achieve optimum data reuse, use specialised vector registers for floating point operations specific to the architecture, and is capable of hiding communication latencies with computation.

Apart from hardware oriented optimisations, the lattice QCD community has developed a range of algorithmic oriented methods which speed up convergence of iterative solvers. These consist mostly of preconditioners, such as even-odd preconditioning, multi-grid methods and low-mode eigenvalue deflation which reduce the condition number of the matrix under inversion. On the other hand, variants of iterative solvers have been developed driven by lattice QCD demands, such as the stabilised Bi-Conjugate Gradient method, the multiple-shift Conjugate Gradient method and multiple right-hand-side methods.

A number of numerical methods have also been developed for speeding up the production of gauge configurations. These are typically oriented towards improving the acceptance ratio of the Monte Carlo update as well as obtaining a better sampling of the configuration space. An early development was the so called hybrid Monte Carlo (HMC) method [4], which is a modification which greatly improves the acceptance ratio over standard Monte Carlo. Within this method, extra degrees of freedom are added to the QCD action, which serve as momentum variables. The expanded action is treated as a Hamiltonian of a molecular dynamics system, with the QCD action part serving as the potential term. An HMC update consists of evolving the system deterministically, with an accept/reject operation at the end of the evolution. This introduces parameters to the simulation such as the time-step of the molecular dynamics update, which can be adjusted to improve further the accept/reject ratio. With the HMC algorithm having been established as a standard in simulating lattice QCD, a large number of modifications have been further developed for improving it by the community. Most notable is the multiple time-scale algorithm, which introduces a hierarchy of time-scales in the molecular dynamics evolution, as well as the so-called Hasenbuch trick [5] which allows reformulating the action such that a larger time-step can be used, without affecting the acceptance ratio.

The lattice QCD community has developed and made public a range of open source software codes. The most comprehensive package is Chroma [6], which is being developed under the US DOE funded SciDAC project. This package implements all major formulations of lattice QCD and various simulation algorithms. Another open source code that uses GPUs is QUDA developed by computational physicists at Boston University.

The lead scientists have experience with Chroma and QUDA [7] as well as the tmLQCD software, an open source code developed by the European Twisted Mass (ETM) collaboration. A copy of the source code can be obtained at http://www.itkp.uni-bonn.de/~urbach/tmlqcd. Implementation details have been published in Ref. [8].

The tmLQCD code implements various algorithms for inversion of the QCD fermion operator, such as Conjugate Gradient, Stabilised Bi-Conjugate Gradient, General Minimised Residual, and multi-shift variants of a subset of these. The code is also the main software used by the ETM collaboration for generation of their gauge field configurations, and implements a range of sophisticated simulation algorithms to accelerate the Polynomial Hybrid Monte Carlo method used, such as even/odd and mass preconditioning and multiple time scale integration.

3.4.3 Community involvement, expected outcomes and their impact

The community

There are major lattice QCD groups in Europe, the US, Japan and, very recently, China. The community ranges from large collaborations to small groups. In Europe well-known collaborations are UKQCD, QCDSF, ETM and BMW. They differ in the type of discretisation scheme each is using. These collaborations span across Europe engaging scientists from major Universities and research centres, such as Edinburgh, Liverpool, Swansea, Glasgow, Trinity College Dublin, Madrid, Bern, Groningen, Bonn, Berlin, DESY-Zeuthen, Regensburg, Wuppertal, Mainz, JSC, University of Cyprus, Cyprus Institute, Saclay, Marseille, Budapest, Grenoble, ECT*, Rome, Pisa, Parma, SISSA/Trieste, Barcelona, Valencia, etc. A number of these institutions have dedicated machines for lattice QCD as for example the APE-machines at Rome and DESY-Zeuthen. ECT* Trento is managing the AuroraScience project (web.infn.it/aurorascience), which has contributed to the development of the Aurora computing system and to the development and optimisation of parallel and hierarchical algorithms for a variety of scientific fields, including in particular Lattice QCD.

The group that will lead the code development within WP8 consists of scientists at the Simulation Labs at CaSToRC and JSC (including scientists from DESY-Zeuthen). This effort will be supported by members of the community at ECT*, Saclay, and at the National Technical University of Athens through GRNET. The team at CaSToRC has common activities with PetaQCD, a consortium of 7 laboratories (IN2P3, CNRS, INRIA, CEA) and two SMEs, within WP7 ("Enabling Petascale Applications: Efficient Use of Tier-0 Systems) in PRACE-1IP. The lead scientists will be C. Alexandrou and G. Koutsou (CaSToRC), St. Krieg (JSC) H. Simma, DESY-Zeuthen and Tsapalis (National Technical Univ. and Hellenic Naval Academy, GRNET), L. Scorzato (ECT*), J. Carbonell (Saclay).

From PRACE CaSToRC, JSC (including scientists from DESY-Zeuthen) shall lead the efforts. CaSToRC has committed a 24 PM PRACE-2IP contribution. JSC and DESY-Zeuthen will contribute using PM funded internally. GRNET has expressed interest in this task with the intention of allocating 12 PM, and shall contribute through coordination with CaSToRC. ECT* is in contact with CINECA for this activity. GENCI has been supporting lattice QCD projects by allocating a sizeable amount of national computing resources, a policy that will be continued. A list of representative publications by these groups in given in references [10-14].

Outcomes and impact

Lattice QCD simulations are CPU intensive. Due to its local nature, lattice QCD is an ideal application to parallelise and due to its simplicity it has been used as a prototype application on novel computer architectures. Progress in this field depends crucially on having optimal

codes. Therefore preparation of codes for the next generation of computers is critical for maintaining progress in the quest for understanding the physics of strongly interacting elementary particles. Although scientists in this field are well versed in adopting codes for new machines, the community has been benefiting from public-domain software such as CHROMA developed within the US SciDAC project. Therefore developing open source codes for the next generation of computers will have a big impact on the ability of the scientific community to effectively use these computers to produce competitive science.
As a prototype application, lattice QCD has had a profound impact on IT and computer science. Scientists in this field pioneered dedicated machines that led to innovative computers, the best example being the Blue Gene computers of IBM. Therefore this line of research is bound to continue to be useful for the development and benchmarking of innovative computer architectures.

It is crucial for the Lattice QCD community to develop the necessary algorithms and parallelisation strategies in order to effectively make use of future computer architectures, as has been the case in the past years. The linear algebra-like operations, which take up most of the computational time of lattice QCD applications, are inherently bound by access to memory or node-to-node communication. So far, most software codes available employ only a single level of parallelisation, e.g. using one MPI process per core, and are typically unaware of thread-level parallelism (apart from codes which are explicitly written for GPUs).

In this work package our goals will be two-fold. On the one hand we shall identify and implement algorithms which are currently available and which are designed to tolerate larger bandwidth and latency than straightforward implementations of lattice QCD. Such algorithms are, for example, domain decomposition and its application as a preconditioner [15], which has been shown to reduce communication requirements in the inversion of the Dirac operator. On the other hand we shall refactor the most performance critical routines found in lattice QCD codes, such as the calculation of the force in the HMC molecular dynamics update, and the application of the Dirac operator. This refactoring shall be done in order to exploit a two-level parallelism, with node-to-node parallelism on the top level and thread-level parallelism at the lower level. The end result will be a library, which, with appropriate documentation, will be made publicly available to the community. We additionally intend to demonstrate the usage of this library by linking to a few select cases, such as the tmLQCD code.

The library we will develop shall be released as open source software and will be available to the community as well as the general public. It is customary in this community that such technical developments be presented at the annual conference of the field (Lattice symposium), in which the entire community is represented. Depending on the outcomes the implementation details may be posted on the preprint arXiv as a paper and submitted to a peer reviewed journal.
Lattice QCD kernels have traditionally been used as benchmark applications on prototype architectures, since scalability of the underlying algorithms is usually ideal. Furthermore, due to the regular data access patterns and deterministic execution schedule of the linear algebra operations, the exact number of floating point operations is countable and known, meaning the floating point efficiency of these kernels is directly measurable. The library we intend to develop can therefore easily be included in a benchmark suite and profiling information can be built in and switchable through a compilation option, for instance.

3.4.4 Relevant Bibliography

[1] Scientific Grand Challenges, Forefront questions in Nuclear Physics and the role of computing at the extreme scale, U.S. Department of Energy, Jan. 26-28, 2009, Washington, D.C., USA.

[2] S. Aoki, (HAL QCD) ,"Hadron Interactions in lattice QCD", arXiv:1107.1284.
[3] Scientific Grand Challenges, Challenges for understanding the qunatum universe and the role of computing at the extreme scale, U.S. Department of Energy, Dec. 9-11, 2008, Menlo Park, CA, USA. U.S. Department of Energy, Jan. 26-28, 2009, Washington, D.C., USA
[4] S. Duane, A. D. Kennedy, B. J. Pendleton and D. Roweth, "Hybrid Monte Carlo", Phys. Lett. B195, 216 (1987).
[5] M. Hasenbusch, "Speeding up the hybrid Monte Carlo algorithm for dynamical fermions", Phys. Lett. B519, 177 (2001) [arXiv:hep-lat/0107019].
[6] Robert G. Edwards, and Balint Joo, "The Chroma Software System for Lattice QCD", Nucl. Phys. Proc. Suppl. 140:832 (2005), arXiv:hep-lat/0409003.
[7] M. A. Clark, R. Babich, K. Barros, R. Brower, and C. Rebbi, "Solving Lattice QCD systems of equations using mixed precision solvers on GPUs", Comput. Phys. Commun. 181, 1517 (2010), arXiv:0911.3191 [hep-lat].
[8] K. Jansen and C. Urbach, "tmLQCD: A Program suite to simulate Wilson Twisted mass Lattice QCD", Comp. Phys. Commun. 180 (2009) 2717-2738 arXiv:0905.3331 [hep-lat].
[9] "Jülich Blue Gene/P Extreme Scaling Workshop 2009", Bernd Mohr and Wolfgang Frings (Eds.) February 2009 [http://www2.fz-juelich.de/jsc/files/docs/ib/ib-10/ib-2010-02.pdf].
[10] G. Egri et al. "Lattice QCD as a video game", Comput. Phys. Commun. 177 (2007) 631
[11] S. Durr et al. "Ab-Initio Determination of Light Hadron Masses", Science 322 (2008) 1224
[12] H. Baier et al. "QPACE – a QCD parallel computer based on Cell processors", PoS LAT2009 (2009) 001
[13] C. Alexandrou et al. "Low lying baryon spectrum with two dynamical twisted mass fermions", Phys. Rev. D80 (2009) 114503
[14] X. Feng et al. "Two-flavor QCD correction to lepton magnetic moments at leading-order in the electromagnetic coupling" (2011) arXiv:1103.4818
[15] M. Luescher, "Lattice QCD and the Schwarz alternating procedure", JHEP 05 (2003) 052, hep-lat/0304007

3.5 Other communities

The scope of WP8 may be extended to further communities that should propose for collaboration with PRACE, matching the selection criteria described in Section 2.1.
Particularly promising areas, in which some investigation and contacts are already ongoing, are:

- Engineering (e.g. fluid dynamics for design of mechanical systems);
- Geophysics (e.g. seismic wave propagation);
- Life Science (e.g. molecular dynamics for bimolecular systems);
- Plasma Physics (e.g. studies related to nuclear fusion and energy)

The inclusion of specific communities from these fields has to be assessed also in light of the finite amount of resources (PMs) available for WP8.

4. WP8 organisation

WP8 has the complex duty of keeping coordinated PRACE's partners with the scientific community, their research targets and the codes they develop. For this reason a specific organisation has been defined.

WP8 is coordinated by the WP leader (CSCS-ETH) and it is organised into inter-disciplinary teams around the communities. Participating centres contribute programming environment and architecture experts. The community code projects commit software developers with domain expertise.

Each team may be further split in Units, each working on a specific numerical/computational problem and it is led by a Technical Leader (and a deputy co-leader), selected between PRACE's partners, that coordinates the team activities, keep tracks of the progress of the work, reports to the WP leader and collaborate with the Community Leader for an effective interaction between HPC experts and codes developers.

The Community Leader is a scientist representing the scientific community. His/her role is to interface the scientists and their needs to the PRACE framework and to coordinate with the Technical Leader on the software development activities. Finally, the Community Leader have a close interactions with the WP leader, in order to keep the community continuously involved and updated on the project progresses and achievements.

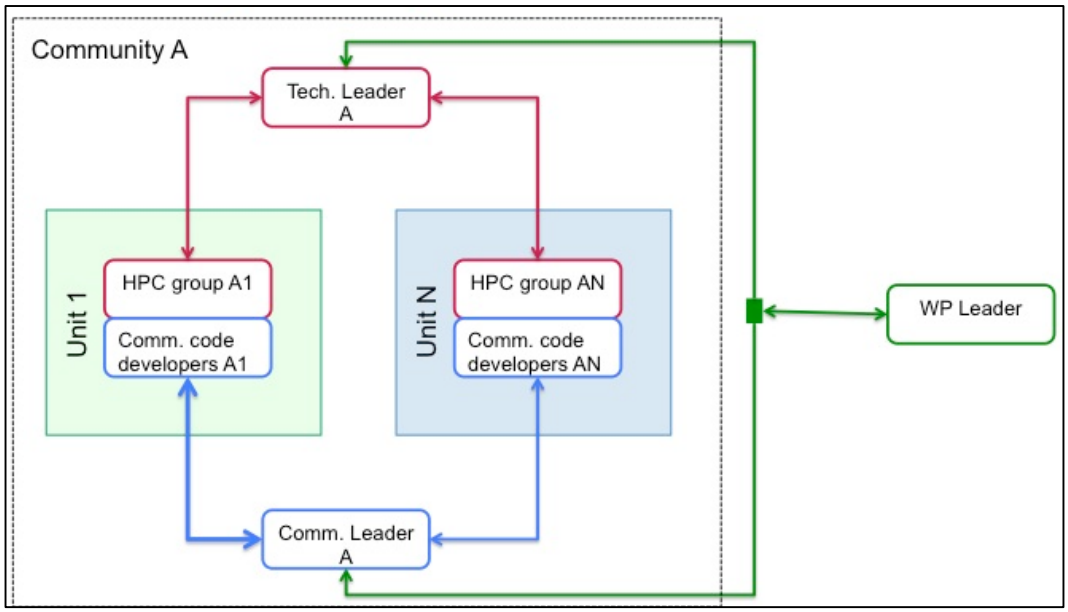

Figure 10: WP8 organisation for each Community Team

The Technical and the Community leaders have been appointed at the PRACE-2IP kick-off meeting (Barcelona, October 15-16):

Table 3: Team leader institutions

	Astrophysics	**Climate**	**Material Science**	**Particle Physics**
Tech. Leader	UC-LCA	NUI Galway	ETHZ	CASTORC
Tech. Co-leader	EPCC	ETHZ	CINECA	GRNET
Comm. Leader	AstroSim	ENES	ETSF	DESY & Uni. of Cyprus

A detailed distribution of work, with the effort devoted by the partners to each community, will be defined at the first WP8 Face-to-Face meeting, that will be held mid-October in Lugano, Switzerland. In the same meeting the Teams will be defined.

5. Conclusions and Next Steps

In its first month, WP8 has focused on the selection of the scientific communities that will be supported and guided in the adaptation and refactoring of their most relevant codes toward the effective exploitation of the next generation of HPC architectures. The same communities have committed to strongly contribute to the success of the workpackage by taking care of the code development, exploiting and leveraging the support of the PRACE's HPC centres.

A set of criteria have been defined in order to identify and select the communities, the most relevant being their impact on science and society, the relevance of supercomputing to the scientific cases and the willingness of the involved communities to contribute to the project and invest in software refactoring and algorithm re-engineering. A further important aspect was the potential impact and effectiveness of WP8 work on each community. Taking into account the finite resources (PMs) available in the workpackage, communities with well defined structure and organisation and with an already ongoing effort compatible with the WP8 objectives were favoured. It is interesting to notice that our analysis revealed that this level of organisation can be very different in different scientific areas. It therefore would be worthwhile for PRACE to consider whether it may actively work in community formation and coordination in future PRACE projects, in order to extend to additional communities.

According to the adopted criteria, Astrophysics, Climate Science, Material Science and Particle Physics have been selected as those scientific domains where communities could be identified in order to start a joint and synergic action along with the PRACE-2IP project. Depending on the WP8 resources, further communities may be added later, extending the scope and the impact of the project.

The next two steps of WP8 (that will lead to deliverables D8.1.2, due at PM3, and D8.1.3, due at PM4) will complete the initial assessment and selection phase. The main objectives are the identification of those numerical algorithms that are key components of community codes and that require specific care to be efficiently and effectively exploited on the emerging HPC architectures. Such algorithms may require a re-design and re-factoring that can be accomplished only through tight interaction between HPC experts and code developers. They will isolate and analyse specific algorithms in order to identify possible solutions to be implemented in the framework of WP8.

For this purpose, we will start defining:

- The detailed contribution of the partners to the different communities, in terms of effort and competencies. The Community Teams will be defined as well.
- The performance model, that will allow to test and verify quantitatively and accurately the characteristics, the performances and the bottlenecks associated to the codes the scientific communities are interested in (listed in Section 3).

These topics will be the subject of a first WP8 meeting scheduled mid-October in Lugano (Switzerland). Community codes experts will participate to the meeting, in order to introduce the proposed scientific software and to start the discussion with the HPC experts, focusing on the features of the different algorithms and on the most suitable performance-tracking model. At the end of the meeting the different partners will be in charge of analysing and profiling one or more algorithms, according to the selected performance model.

After the meeting, the code performance analysis will be accomplished, with the identification of promising and relevant areas for improvements, highlighting those areas that are the most critical, addressing numerical solutions to major scientific problems requiring high-end computational resources, and/or that are common to a number of different HPC codes and

that, therefore, can have a dramatic impact on a large number of applications. Solutions for different architectures will be investigated and the most effective will be identified and proposed for the refactoring plan expected at PM6.

At the same time, the Technical Leaders and Co-leaders will work with the corresponding community representatives, in order to consolidate the collaboration between PRACE and the scientific communities and to start defining the working methodology that will lead, starting from PM7, PRACE HPC experts and community developers teams to interoperate on the design, development and refactoring of the codes. The resulting work plan will be presented in deliverable D8.1.4 at PM6.