Roadmap Report

Project acronym: *MAPPER*
Project full title: Multiscale Applications on European e-Infrastructures.
Grant agreement no.: 261507

DOCUMENT INFO

Date and version number	Author	Comments
V0.1, 28/8/13	A. Hoekstra	Introduction written.
V0.2, 10/9/13	A. Hoekstra	Part on external users added
V0.3, 12/9/13	A. Hoekstra	Part on projects added
V0.4, 15/9/13	A. Hoekstra	Part on projects finished, part on communities added
V0.5, 17/9/13	A. Hoekstra, D. Coster, C. Kurowski	Added text on e-infrastructures and fusion community.
V0.6, 30/9/13	W. Dubitzky, D. Groen, A. Hoekstra	Additional text on collaborations
V0.7, 4-10-2013	A. Hoekstra	Final sections
V0.8, 8-10-2013	O. Hoenen, A. Hoekstra	Additional text on collaboration with ITM
V0.9, 9-10-2013	B. Chopard, A. Hoekstra	Additional text on Swiss collaborations
V1.0, 10-10-2013	M. Schiffer, K. Rycerz, A. Hoekstra	Uptake of tools, more on DRIHM, final editing
V1.1, 14-10-2013	A. Hoekstra, W. Dubitzky, D. Coster	Final additional editing
V1.2, 14-10-2013	A. Hoekstra, W. Dubitzky,	Final formatting and checking

Table of Contents

1 Executive Summary

Following recommendations after the 2nd review, the MAPPER project has produced a roadmap identifying steps required to get new communities on board, identifying training, documentation and support, as well as associated planning. The roadmap was implemented in the final eight months of the project. We approached about 25 individuals, of which six actually started to use MAPPER, while another 4 confirmed that they start to use MAPPER by October 2013. We trained them, offered support, and some of them are already producing results that will result in scientific publications. Moreover, we have actively disseminated MAPPER to eleven projects, resulting in varying levels of collaboration (ranging from actually applying MAPPER to applications in these projects to discussing opportunities offered by MAPPER). We have had a very active collaboration with PRACE and EGI, resulting in true uptake of MAPPER technology. Finally, a number of scientific communities are now aware of MAPPER, and again to a varying level, are considering to adopt solutions as offered by MAPPER.

2 Introduction

After the second review of the MAPPER project, an important recommendation from the reviewers was that “A roadmap should be produced by the end of January 2013 identifying the steps required to get new communities on board (and delivering good science using MAPPER before the end of the project). This roadmap will identify training, documentation and support to be provided, as well as the timescale for each step.” (recommendation 7). Other recommendations stressed the importance of uptake by external users (recommendation 1), more widely uptake of tools and methodology within the scientific community (recommendation 2), put strong effort in getting external applications using MAPPER (recommendation 5), and prepare a best practice guide for external MAPPER users (recommendation 6).

Following up to recommendation 7 we wrote a roadmap in January 2013, also taking into account the other recommendations as highlighted above. This roadmap was sent to the Project Officer and reviewers and is also publically available from the MAPPER website.[1] In that document we identified four stakeholders (individuals, projects, communities, e-infrastructures) and four actions (“get on board”, “produce good science”, dissemination, impact). Below in Table 1 we reproduce the main table from the roadmap that summarizes the planned content of the actions for each of the stakeholders.

The goal of this report is to describe the actions that were taken from January 2013 to August 2013 and their results, for each of the four stakeholders, to shortly discuss the uptake of MAPPER tools and technologies, and finally to summarize impact and sustainability of MAPPER for the stakeholders. In section 3 we describe in some detail the results we obtained for all stakeholders, and in section 4 we summarize the uptake on MAPPER technology. Finally, in section 5 we summarize these results in the light of the planned and actual actions and outcomes, as formulated in the roadmap.

Please note that some of the actions and results described in this report are part of the DoW and translated into specific tasks and key performance indicators (KPIs). These KPIs and their status are described in the dissemination reports, deliverable D2.2, D2.4, and D2.5, and will be summarized here. Other results, such as technical collaboration with EGI and PRACE in the context of the *taskforce* have also been reported in regular deliverables, but also contribute to the goals as set out in the roadmap. We finally note that the actions as defined in the roadmap will not end when closing the MAPPER project in September 2013. As part of our sustainability plan (deliverable D3.6) we will take sufficient measures to sustain MAPPER tools and services after the project, such that stakeholders can and will continue to use these for their multiscale modeling and computing projects. Among those measures are tutorials and best practice guides, public domain availability of all key software, adding of key services to the EGI software stack, and installing and making available tools on a number of NGIs.

[1] See `http://www.mapper-project.eu/c/document_library/get_file?uuid=f15191b7-c956-4d62-ba00-ff3d33282608&groupId=10155`.

Table 1: Content and goals of the four Roadmap actions for all stakeholders (reproduced from the Roadmap document).

	Communities	Projects	Individuals	e-Infrastructures
"Get on board"	Yes, from MAPPER Fusion, Biomedical, Fusion, Biology, Hydrology, for external we target Hydro-Meteo and geosciences.	Yes, we aim to target DRIHM, Thrombus, MeDDiCa, Scalalife, and possibly others, to be identified	Yes, we aim to work with at least 5 external individual researchers.	With EGI an MoU is in place, with PRACE we collaborate within the taskforce, in particular with LRZ, PSNC (partners in MAPPER), SARA and CINECA.
"Produce good science"	Via *individuals* and *projects*.	Where possible we will set up joint work in making MAPPER services available to these project, running multiscale simulations on MAPPER, etc.	Yes, we will train[2] and fully support[3] the external researchers, so that they will be able to run multiscale simulations, not possible before, on MAPPER.	We will publish a few papers, together with PRACE and EGI, reporting on distributed multiscale computing and the services to support this.
Dissemination	Via all channels as described in our dissemination plan, including the seasonal school. Also via the social media.	Via all channels as described in our dissemination plan, including the seasonal school. Also by joining project meetings. Also via the social media.	By direct collaboration.	Via all channels as described in our dissemination plan, including the seasonal school. Also by actively participating in the User forum of both EGI and PRACE. Also via the social media.
Impact	Scientific output (what became possible with the help of MAPPER?), educational (training) material, best practices, tutorials, services and tools.	Scientific output (what became possible with the help of MAPPER?), educational (training) material, best practices, tutorials, services and tools.	Scientific output (what became possible with the help of MAPPER?), educational (training) material, best practices, tutorials, services and tools.	Services and tools for distributed multi-scale computing together with general purpose advance reservation and co-allocation capabilities in EGE and PRACE, best practices.

According to the DoW of MAPPER, after two years into the project all services and tools for distributed multiscale computing would be ready to facilitate production runs of multiscale applications. Moreover, the application portfolio of MAPPER would be fully running in DMC production mode, so that that third year could be spent on two tasks, performance measurements and production runs. As was acknowledged during the second review of the project, MAPPER indeed achieved this goal after two years. With that, the project was also in the right position to reach out, and liaise with external parties (be it individuals or other projects) to offer them access to MAPPER tools and services, and MAPPER infrastructure, and allowing them to develop their own multiscale applications and execute distributed multiscale computations.

Although we already engaged in contacts with a number of projects and individuals on possibly using MAPPER during year 2 of the project, only in year 3 we systematically started to work with external parties, basically as anticipated in the dissemination plans, and further triggered by the recommendations after the second project review. In section 3 we provide a short account of these activities and their results, followed by an analysis of uptake of

[2] We plan to organise a face-to-face and hands-on training event between external users and MAPPER experts.

[3] By the MAPPER operations team.

MAPPER tools and technologies in section 4. Finally in section 5-7 we summarize, look ahead, and shortly discuss sustainability in the light of the roadmap activities.

3 Stakeholders

3.1 Individual Users

In our dissemination plan we have defined an important Key Performance Indicator (KPI), namely KPI Q10 (number of external users of MAPPER technologies) which we set to 5 by the end of the project. As part of the roadmap activities we decided to stretch this KPI to mean a minimum number, and to target a larger number (10 to 15).

In January and February 2013 we have identified and approached potential external users and offered them the opportunity to 'get on board' and start using MAPPER. We also offered them support, both for training and for actual implementation and execution of multiscale applications. Partners in MAPPER either identified these individual scientists, or they approached us during conferences or workshops. Also as a result of the 2nd MAPPER summer school in June 2013 we got into contact with more individuals interested to use MAPPER.

In the period from January to February 2013 we talked to about 25 individuals. Out of these, six actually started to use MAPPER in the period March to August 2013. Another four confirmed that they start to use MAPPER by October 2013, in line with their own project planning. For these four users we offered our support. Finally two individuals requested to be kept informed on MAPPER, as they may be interested to start using MAPPER for their research. Table 2 shows more details of these users.

The active users were trained on the spot, by MAPPER personnel, relying on our on-line documentation and tutorials[4]. As we also noticed during the second summer school, our external users quickly grasped the basic ideas behind the MAPPER methodology, and in general appreciated and understood the tools and services, and on relative short notice became productive using these tools and services (see e.g. deliverable D2.5.2 reporting on the assessment of MAPPER tools by the participants of the summer school). The factor that delayed the external users was not so much the MAPPER technology, but just the time they could allocate to work on their multiscale applications.

[4] www.mapper-project.eu/web/guest/best-practices

Table 2: Overview of external users (VPH – Virtual Physiological Human; HMR – Hydro-Meteorology).

	Name	affiiliation	community	current status
Active Users	Dr. C. Wang	Wildau University of Applied Sciences, Germany	VPH	His cardiomyocytes multi-scale model, based on the CardioWave framework is currently being deployed on the MAPPER e-infrastructure.
	Dr. L. Rauch	AGH University of Science and Technology, Cracow, Poland	engineering	Currently, the multiscale application for metallurgical industry (multiscale modelling and simulations of car body parts) is ported to GridSpace. We have already built working GridSpace Experiment experiment that executes Abaqus and SSRVE software packages used for numerical simulations of car parts. The work is summarized in a poster "Efficient Execution of Grid-based Multiscale Applications for Metallurgical Industry" presented at Summer School on Grid and Cloud Workflows at Budapest, Hungary, 1-7.07.2013.
	K. Kowalik	Center for AstronomyNicolaus Copernicus University, Poznan, Poland	Astrophysics	Both are active QosCosGrid users in the MAPPER/PL-Grid e-infrastructure. Their code (MHD code called Piernik) has been already integrated with MUSCLE: http://apps.man.poznan.pl/trac/muscle/wiki/Piernik
	D. Woltanski	Center for AstronomyNicolaus Copernicus University, Poznan, Poland	Astrophysics	
	J. Osborne	University of Oxford, UK	VPH	The MUSCLE framework has been deployed and tested on all HPC resources currently used by the team. They are now in the process of porting their scenario to MAPPER.
	M. Bernabeu	University of Oxford, UK		
Confirmed Users	Dr. A. Perodi	CIMA Research Foundation, Italy	HMR	
	Prof. A. Clematis	Istituto di Matematica Applicata e Tecnologie Informatiche Consiglio Nazionale delle Ricerche, Genova, Italy	HMR	
	Prof. L. Garrote de Marcos	Departamento de Ingenieria Civil: Hidraulica y Energetica Escuela de Ingenieros de Caminos Universidad Politecnica de Madrid, Spain	HMR	
	Dr. A. Galizia	Istituto di Matematica Applicata e Tecnologie Informatiche Consiglio Nazionale delle Ricerche, Italy	HMR	
Interested	T. Glatard	CREATIS, INSA Lyon, France	VPH	Virtual Imaging Platform project. This project has a need for PRACE, EGI, and EUDAT resources. CREATIS is also a partner in the THROMBUS project.
	Dr. M. Swain	Aberystwyth University, UK	Biology	Modelling microbial communities, to understand the process of plant degradation in the digestive systems of ruminants, which are important emitters of green-house gases. Our aim is to build predictive, mathematical models to understand these process within the rumen, and here multi-scale approaches offer a number of attractive features: we would like to build individual models of plant metabolism and the metabolism in various microbial species, which would be coupled together using MAPPER to build models of the rumen, and these models we would like to initially calibrate, and then ultimately couple with whole animal models of gas production.

In August 2013 we have asked the active external users to complete a short questionnaire. We received four replies, by Wang, Rauch, Kowalik and Woltanski (jointly), and Osborne and Barnebeu (jointly). Appendix A contains the feedback that we received. It is hard to draw general conclusions from this, as the number of respondents is just to low. However, for these users the main observations seems to be that all of them were able to be productive with and on MAPPER in a relative short period of time, and that they appreciate the tools and services for many different reasons, but in the end they acknowledge the added value for their research. The more in depth study that we did on our own applications, as reported in deliverable D7.3, seems to corroborate this. Our own applications experienced strong benefits, but also for many different reasons.

The external users all expect to obtain some form of scientific results as a result of working with MAPPER, and some of them already published a paper or are writing up their results. Given the relative small time span this is a very encouraging result.

3.2 Projects

Right from the start of MAPPER we have been communicating and collaborating with projects that showed interest in our developments, started to adopt our solutions, or use our tools and services for their applications. In the DoW we already listed projects with which partners in MAPPER collaborate(d), and some of these initial contacts have evolved into very fruitful collaborations, fostering impact, dissemination and sustaining MAPPER.

The type and level of collaboration with projects has been quite diverse. We signed an MoU with a few projects (EGI-Inspire, DRIHM, IGE), worked together in detail on the technical level (PRACE, EGI-Inspire), or on the application level (MeDDiCa, SCALALIFE, DRIHM, MMM@HPC), or only shared knowledge and explored common interest (THROMBUS, CHIC). Notwithstanding the wide spread in engagement and level of collaboration, in all cases these contacts have contributed to creating impact of MAPPER as well as to sustaining MAPPER, as is described in a bit more detail below.

In the sequel of this section we will describe in some detail the results our collaborations with all these projects, except with PRACE, EGI-Inspire, and IGE, which is postponed to section 3.4. Table 3 lists all relevant projects and their URL. We refer to the project websites for more information on all these specific projects.

Table 3: Overview of projects that collaborated with MAPPER.

Name	URL
SCALALIFE	www.scalalife.eu
MMM@HPC	www.multiscale-modelling.eu
MeDDiCa	www.meddica.eu
DRIHM	www.drihm.eu
CHIC	chic-vph.eu
THROMBUS	www.thrombus-vph.eu
ITM	www.efda-itm.eu/
UK 2020 Science	http://www.2020science.net/research/vascular-remodelling
VPH-Share	www.vph-share.eu
CER2EBRAL	ccs.chem.ucl.ac.uk/cerrebral
CRESTA	cresta-project.eu

3.2.1 Direct collaborations via MAPPER's application portfolio

MeDDiCa is a Marie Curie ITN, and UvA is both partner in MeDDiCa and MAPPER. A direct collaboration was established between these two projects via UvA. The in-stent restenosis (ISR) application was developed in MeDDiCA (and earlier projects), was subsequently ported to the MAPPER environment, and used in collaboration with MeDDiCa to establish further scientific results. This not only resulted in a few exciting scientific results on the application, but also to dissemination of MAPPER tools and services to MeDDiCa partners. MeDDiCa is a VPH project, and among others via this collaboration MAPPER has gained visibility within the VPH community (see also section 2.3). Another relevant result was a small task between UvA and Fluent, another partner of MeDDiCa. Fluent is a USA-based company selling Computational Fluid Dynamics software and offering consultancy in the use of this software. They are now developing multi-physics and multi-scale capabilities for their software base, and we explored the potential of the MUSCLE environment, by wrapping the FLUENT code and then coupling it to biological solver codes. This idea unfortunately never got past the drawing board, but the contacts have been established and options are still being discussed.

Collaboration between the Integrated Tokamak Modelling group (ITM), in which both IPP and Chalmers are involved, and MAPPER is centered on physics codes. Fusion codes used in MAPPER have been developed within the ITM around a common data structure, giving a generic framework for coupling different physics. The successful integration of these codes within the MAPPER infrastructure allowed us to run use-cases where the ITM was facing different issues (see section 3.3), and participation to several ITM events improved MAPPER dissemination in the community. Ongoing ITM work concerning the serialization of the data structure is a direct consequence of this collaboration.

The Hydrology application from partner UNIGE leads to a number of further collaborations and projects. MAPPER tools will be used to connect the two fastest Swiss supercomputers (CSCS and CADMOS) to simulate the Rhone river with a 3D free surface flow model. MAPPER tools will also be used in an operational model of the canal de la Bourne, in France, by the team at ESISAR / INP-Grenoble.

3.2.2 Collaborations with other projects funded under the same call

MAPPER is funded under call "INFRA-2010-1.2.2: Simulation software and services". Three other projects received funding under the same call, MMM@HPC, SCALALIFE, and SHIWA (www.shiwa-workflow.eu). At the start of MAPPER we explored options to collaborate with SHIWA. Partner CYFRONET has been in regular contact with SHIWA's project coordinator and visited some of the SHIWA meetings, disseminating the tools as developed in MAPPER to SHIWA partners. However, these contacts with SHIWA have not resulted in technical collaborations.

The contacts with MMM@HPC and SCALALIFE did result in common actions, aiming to prove interoperability between the projects, and creating synergy both in dissemination and sustainability of the results. In Appendix B we provide all details of these actions and results. Here, a summary is provided.

MAPPER, SCALALIFE and MMM@HPC have some commonalities in their project goals. MMM@HPC focuses specifically on applications from material science, and has created a software stack based on Unicore and GridBeans technology to create acyclic (loosely coupled) multiscale simulation workflows. Their approach and vision resonates strongly with MAPPER's vision on multiscale computing. However, the choice for technology has been quite different. Both projects have looked at interoperability of their solutions and demonstrated that this is possible. SCALALIFE has created an easy accessible and useable set of what in MAPPER are called single scale models geared towards the life sciences. SCALALIFE and MAPPER have jointly created a loosely coupled multiscale application based on MAPPER technology, using SCALALIFE single scale models. Finally, the three projects have presented lectures and tutorial on their respective project meetings, workshops and summer schools.

Although the three projects offer complementary and mature technologies, the underlying vision is strongly aligned and fits well together. Having demonstrated interoperability as reported in Appendix B, the three projects have laid a strong foundation for multiphysics multiscale simulations on European e-infrastructures, supporting varying scenarios, from monolithic computations of workflows on PRACE or EGI resources, to tightly coupled scenarios distributed over PRACE and EGI resources. The next steps should be to attract more users and build a strong use base around these technologies. We believe that some form of support action is needed, aiming at:

- Recruiting external users
  - Keep explaining the added value
  - Spend more time with them
- Establish virtual research communities
- Create teaching material for 'Computing Skills' curricula
  - Educational outreach, both on the graduate and post-graduate levels

We are aware of the new instruments that are being developed for Horizon2020, such as the Centers of Excellence for computing applications, and we believe that the outcome of these three projects could form a strong basis.

3.2.3 Collaborations with unrelated projects

During the European Geosciences Union's General Assembly 2013 in Vienna the DRIHM project signed a Memorandum of Understanding with MAPPER. The goal of this MoU is to expand the use of services and tools developed by the MAPPER project within the European hydro-meteorology community. To this end partner LMU has already presented the overall MAPPER vision and technology, and has characterized DRIHM applications in terms of a Scale Separation Map and related Multiscale Modeling Language. The next step is to actually use MAPPER technology to the benefit of selected DRIHM applications. The DRIHM project has scheduled this task to commence after their second review (end October 2013) as alternative to their demonstrator model couplings, and although the MAPPER project has ended by that time, some MAPPER partners (LMU, UvA, PSNC) have committed themselves to offer support if needed.

MAPPER and DRIHM share the common vision of exploiting existing European e-infrastructures for distributed multiscale computing to improve research capabilities for investigating phenomena from the microscopic perspective to a global one and everything in between. Climate, weather and hydro-meteorological modeling are important examples, as they need to consider the global state of the phenomena, but also local conditions. While MAPPER focuses on the fundamental technologies to facilitate multiscale computing, the DRIHM project (and its EU/US-extension DRIHM2US) intends to provide an e-Science environment for hydro-meteorological research which not only eases the access to heterogeneous data and models at various scales, but also the collaboration between meteorologists, hydrologists, and Earth science experts. Although DRIHM is currently built on a proprietary service stack, it became apparent that the DRIHM scientific workflows – basically combining numerical weather prediction models, downscaling/nowcasting models, and discharge models – are characterized by either space and time separation or by time overlap and space separation and that the MAPPER generic tools and services could valuably be leveraged. However, they also need to be supported by standardization efforts, which may be science specific. Consequently, the DRIHM project will adapt tools and services provided by MAPPER to their specific model landscape and integrate them with their technology stack where appropriate during the next project steps. The goal is a sustainable stack of services and tools for hydro-meteorological research.

MAPPER has provided support to the fluid-solid interactions project within the 2020 Science project, which is a British initiative for multiscale modelling in the biomedical domain. This support has resulted in one external VPH application (VPH UCL), which is currently being tested and developed. In addition, the Application Hosting Environment has found uptake in several other projects, such as VPH-Share and the VPH Virtual Imaging Platform.

3.2.4 Emerging collaborations

During the lifetime of MAPPER we have had many contacts with projects that showed more than just interest in MAPPER. However, these initial contacts have not yet resulted in uptake of MAPPER technology by those projects, but we do want to list these projects here, as we believe that we may get these projects ‘on board’, although this will not happen during the life time of MAPPER.

We specifically mention two VPH projects, THROMBUS and CHIC, which have expressed interest in MAPPER. Partners UvA and UNIGE also participate in THROMBUS. This project aims to validate models for thrombus formation in intercranial aneurysms. THROMBUS is now developing multiscale models for some key biological and biophysical processes and is considering applying MAPPER tools and services to build and execute these models. CHIC is a project that aims to create multiscale models for cancer, in a secure environment, ready for in-silico medical trials. Prof. M. Viceconti of CHIC has been discussing with MAPPER the option of applying tools and services of MAPPER, to be integrated with the existing hypermodelling tools developed in the group of Viceconti.

MAPPER partner UU has recently received £11m funding in the area of stratified medicine. It is expected that modelling efforts in this area will face inherently multiscale problems. We are currently working closely with newly appointed scientists in this area to explore how MAPPER results and technology could be used in their research.

The CER2EBRAL project, funded by the Qatar National Research Foundation, has also expressed interest in MAPPER and is considering applying the MAPPER technologies to construct a coupled clinical simulation workflow. However, as the CER2EBRAL project started in the month that the MAPPER project finished, we have not yet observed any concrete uptake.

We also exchanged thoughts with partners in the EU-funded CRESTA project about the potential disruptive aspects of multiscale computing, in the context of the CRESTA disruptive technologies meeting. We do not expect any direct uptake within CRESTA, as CRESTA is primarily an infrastructure project for scaling up single applications to the exascale. However, we did provide them with detailed background information on MAPPER, and we did receive indications that they may identify multiscale computing as one of the potentially disruptive technologies, which they intend to study more closely later in the project.

Researchers from the Centre for Scientific Computing at the University of Warwick have shown interest in taking up the MAPPER tools as well, and are considering to use them to establish fast and reusable couplings between boundary element method (BEM) solvers, and other numerical codes. The first major discussions on this emerging collaboration are due to take place in late October.

MAPPER tools have been selected for the Swiss HPC project "Biology across scales" in which several research institutions in Switzerland participate. MAPPER tools are also considered in the Institute for Environmental sciences, at UNIGE, for coupling codes and to develop new multiscale applications.

3.3 Communities

The outcome of the MAPPER project is beneficial for large parts of the Earth Sciences community as their primary goal is to couple models and adapt data on several scales. Consequently, the MAPPER project itself and the results produced during the course of the project were presented at the European Geosciences Union General Assembly (EGU) 2013 in Vienna, Austria – both in their session on ICT-Based Hydrometeorology Science and Natural Disaster Societal Impact Assessment and in the special session hosted by the DRIHM project. Participants of the American Geophysical Union (AGU) expressed interest for a similar presentation at their annual venue (Fall Meeting beginning December 2013).

The fusion modeling community in the EU is centered around the European Fusion Development Agreement (EFDA) Task Force on Integrated Tokamak Modeling (ITM). With regard to workflows, the ITM has chosen to concentrate on Kepler, and using a technology partly developed under the FP7 project EUFORIA and partly developed under the ITM to launch code modules on remote HPC resources (using UNICORE). This was functional on the HPC-FF supercomputer in Juelich, but, because of complications with regard to the management of the IFERC-HELIOS supercomputer in Japan, is not yet functional there. The

MAPPER approach developed in this project is, however, functioning on HELIOS, as yet without inter-site functionality. This is allowing the ITM to use resources that have been allocated on HELIOS, and also allows for a possible fallback if it proves impossible to get UNICORE installed on HELIOS - i.e. the MAPPER technology offers the possibility of an alternate approach for coupling to remote jobs, though this will require further development to couple to Kepler workflows. What MAPPER has demonstrated to the ITM is that having alternatives available is very valuable, particular with regard to HPC, since many of the issues faced on these machines are less technical and more political/management issues.

The FP7 ICT initiative Virtual Physiological Human has created a very strong European community on Computational Medicine, by funding a large collection of projects with an overall amount of funding in the order of 150 M€. This community has organised itself, first through the EU funded VPH Network of Excellence, which was led by Prof. Peter Coveney from UCL, who is also partner in MAPPER. Next, the community has created the VPH Institute, which through individual and institutional memberships is self-supporting. MAPPER partners UvA, UNIGE and UCL are also very active in this VPH community, and as such have strongly disseminated MAPPER. Some of the core technologies of MAPPER (such as GridSpace, AHE, and MUSCLE) have been adopted by the VPH toolkit that was established by the VPH NoE. Within this community there is a growing interest in multiscale computing on high end computing infrastructure, and on applying this for pre-clinical or even clinical applications. We are currently talking to all main stakeholders on using MAPPER tools and services to support this strong VPH computing needs. At the time of writing this has not yet resulted in formal agreements, but the informal contacts are strong, and in preparing for H2020 MAPPER has a good position to create true and lasting impact in this community.

In the recent decade, the life sciences have undergone a major paradigm shift, from a reductionist to a holistic science. The new biology is called systems biology. Systems biology is inherently multiscale in nature. It aims to explain, predict and control the properties, behaviour and functions of living entities, phenomena or systems. Biological systems range from molecular subsystems to large-scale ecosystems. Typical biological systems and processes studied in systems biology include mitochondrial production of energy; ribosomal protein synthesis; networks biochemical reactions regulating gene expression; within-cell and cell-to-cell signalling and communication pathways and metabolic pathways (e.g., glycolysis and the cell cycle); organs; genes, gene complexes, genomes and meta-genomes; cells, unicellular organisms, multicellular organs, multicellular-multiorgan organisms; populations and demes (e.g. a local population of a species); microbial communities; ecosystems such as lakes, rainforests, deserts or predator-prey ecosystems; and the biosphere (the totality of all ecosystems While many approaches to systems biology involve single-scale models, there is a growing body of work that aims at modeling of life phenomena across several scales on planet Earth). Biological systems could be viewed along a hierarchy of levels of biological organization, ranging from the environment to the quantum level. Ulster is actively working in the field of systems biology. Of particular interest are systems-based approaches to degenerative diseases (e.g. in stratified medicine). The knowledge and technology developed in MAPPER will help Ulster to engage in new collaborations in this area. Recently, Ulster has been coordinating a special issue in Computation on multiscale modelling and simulation in computational biology.

3.4 E-Infrastructures

The MAPPER project aimed to deploy a computational science environment for distributed multiscale computing on and across European e-infrastructures. In order to further the project's aim, MAPPER initiated a collaboration with PRACE and EGI in May 2011. The coordination of activities relevant to this collaboration was performed via an expressly formed Task Force comprising specialists from each of the three projects. The short-term goal of the taskforce was a rapid prototyping (fast track development) of two distributed multiscale computing applications, to be delivered by the first MAPPER review (November 2011). During the review two distributed multiscale scenarios were demonstrated showing the full interoperability between EGI and PRACE. MAPPER has delivered a set of core middleware services QosCosGrid (QCG) extending existing capabilities (e.g. advance reservation) and application tools (MUSCLE and GridSpace) to allow users to run and control distributed multiscale applications on selected sites from EGI and PRACE.

Additionally, MAPPER has performed various loosely coupled and tightly coupled distributed application tests and benchmarks (see respective deliverables). The successful results achieved during the first period of the Task Force encouraged participants to enter the second phase of MAPPER contribution to operational tools developed by EGI and PRACE as well as participation in the Unified Middleware Distribution (UMD). Therefore, the following objectives in the second phase of the Task Force have been defined:

- to continue multiscale applications tests and benchmarks on production sites from PRACE and EGI,
- to integrate QCG Nagios probes into the SAM release and to assess which of these probes will affect site Availability/Reliability in case of failure,
- to assess the status of accounting, and the developments needed to integrate what exists with the EGI accounting infrastructure,
- to integrate QCG support activities into the EGI mainstream support tasks (1st level, 2nd level and 3rd level).

In the meantime, MAPPER users have successfully applied for PRACE grants, and used production HPC resources provided by LRZ, SARA, and EPCC. During the second MAPPER review (November 2012) all MAPPER distributed multiscale applications together with tools were demonstrated in fully operational mode on EGI and PRACE. The remaining tasks have been completed in 2013. Importantly, UMD 3.2.0 as the second update of the third major release of UMD was recently released and it now includes MAPPER core middleware software components:

- QCG ntf,
- QCG comp,
- QCG accounting,
- QCG broker client,
- QCG broker.

MAPPER software has been made available for the EGI production infrastructure and will be supported by PSNC as Technology Provider in the future. Moreover, as the core QCG services have been successfully integrated with multiscale application tools, in particular GridSpace (GS) and MUSCLE, external communities have now the ability to use these frameworks to perform their own multiscale computational experiments on MAPPER e-Infrastructure. At the same time, this represented an important and final stage in Task Force activities as MAPPER has moved forward towards becoming part of the sustainable EGI infrastructure used by the scientific communities at large in Europe.

4 Uptake of MAPPER technologies

From the uptake and sustainability point of view, MAPPER tools and technologies fall into the following three groups:

- Local domain (low level) MAPPER services like QCG-Computing (an advance reservation and execution service) or QCG-Notification (a notification service).
- Grid domain (high level) MAPPER services and user interfaces like QCG-Broker (a load balancing and co-allocation service) or Application Hosting Environment (AHE) which facilitates application specific services to utilize Grid resources.
- Multiscale coupling libraries like MUSCLE (a tool for tightly coupled MMS applications) or GridSpace for facilitating loosely coupled MMS applications. This group also includes a graphical tool for tightly and loosely coupled multiscale application composition i.e. Multiscale Application Designer (MAD) and registry of models descriptions i.e. MAPPER Memory (MaMe)

A list of all tools and technologies together with their capabilities are presented in Table 4.

The first two groups include the QosCosGrid middleware with various new interfaces, extensions and application tools developed during the MAPPER project. QosCosGrid middleware has joined recently the EGI Universal Middleware Distribution (UMD). Technically speaking, UMD is the set of software components developed for EGI by technology providers to provide the innovation needed by EGI to satisfy its users that cannot be found elsewhere and is therefore endorsed by EGI for use within the production infrastructure. In other words, UMD is a software repository for distributed computing and middleware technologies and can be seen as Apple App Store or Android Market. The addition of QosCosGrid to UMD means that the MAPPER software has successfully passed restrictive procedures that proved its high maturity and usefulness.

Table 4: Overview of MAPPER tools and technologies.

Software	**Capabilities**						
Middleware Services	**Basic Jobs**	**MPI Jobs**	**Loosely Coupled Jobs**	**Tightly Coupled Jobs**	**Advance Reservation Support**	**Automated Notifications About Jobs**	**Automated Co-allocation Support**
QCG-Computing	x	x	x	x	x		
QCG-Notification						x	
QCG-Broker	x	x	x	x	x	x	x
AHE	x	x	x		x		
Application Tools							
MAD	x	x	x	x	x (via QCG)		
MaMe	x	x	x	x			
GridSpace	x	x	x	x (via MUSCLE)	x (via QCG)	x	x (via QCG)
MUSCLE		x		x	x		x

Consequently, the MAPPER software will enhance the interoperability and integration between grids and other computing infrastructures over the next few years, strengthen the reliability and manageability of the services and establish a sustainable model to support, harmonise and evolve the QosCosGrid middleware, ensuring it responds effectively to the requirements of the scientific communities relying on it, especially for DMC. Starting from September 2013 supercomputing centers accross Europe can easily and safely deploy the QCG components on their HPC and HTC sites fully compatible with European Grid Infrastructure.

The AHE has been distributed through a number of different channels, in order to reach as wide an audience as possible. The release of AHE 2.0 took advantage of a trend towards server virtualization, with the AHE server provided as a preconfigured virtual machine image, which could be easily downloaded and deployed by a reasonably experienced user. There have been downloads of the AHE client and server software from the RealityGrid site, and further installations made through the OMII stack. In addition, AHE has been made available from other sources, such as the NeSCForge portal, and the VPH ToolKit, and has been used extensively in the EC-funded VPH-Share project, as well as the VPH Virtual Imaging Platform project. The AHE 2.0 has been downloaded approximately 140 times in total. In addition a number of applications have been made available through AHE, for example CPMD, NAMD, LAMMPS, DL_POLY, LB3D and Gromacs.

MUSCLE 2 has found wide uptake within MAPPER, as it is used to facilitate the coupling in ISR3C, Canals, MultiGrain applications, as well as one of the Fusion applications. In addition, MUSCLE has been applied or is being applied in several external applications, including an Astrophysics application by Piernik et al. and the MCardio application, which is part of the VPH initiative. In addition, researchers at UNIGE have been able to run a simulation coupled with MUSCLE on a cloud infrastructure, greatly increasing its range of applicability. MUSCLE 2 is also expected to play an important role in facilitating the coupling in the workflows that are built in the Qatar-funded CER2EBRAL project.

GridSpace was extended within MAPPER to support multiscale applications and it is used to facilitate coupling of nanomaterials and fusion applications. It also is integrated with MUSCLE 2 to support ISR3D, Canals and MultiGrain applications. In addition GridSpace is being applied for Metallurgy application that simulates the manufacturing process of a car crushbox. Multiscale Application Designer and Mapper Memory are used for all MAPPER applications.

5 Summary of Actions and Outcomes

Sections 3 and 4 gave an account of the results obtained in the context of activities carried out as a result of the planned actions of the roadmap. The goal of this section is to provide a more precise account of the planned and actual actions and outcomes, thus putting sections 3 and 4 in line with the roadmap.

The roadmap contained, in its section 3, two tables that summarized the actions and related timeframes to reach the goals of the roadmap. Below we reproduce these tables, but now fill in the column ‘Current Status’, which was as good as empty in the roadmap report. Table 5 applies to reaching out to individual researchers and projects, and table 6 applies to actions related to dissemination and impact.

Table 5: Actions, timing and current status for individual external users and projects.

Time Frame	Action	Comments	Current Status
Jan 2013	Identify	Identify potential individual researchers or projects to collaborate with us.	Completed, 25 users identified, contacts with 11 application projects and three eInfrastructure projects.
Feb 2013	Commitments	Get true commitments, in writing.	Getting commitments in writing was difficult, at best we got email confirmations, in the end 6 users actually started to work with MAPPER, 4 start in October, and 2 are interested. We signed MoUs with three projects, and started scientific and technical collaborations with 14 projects.
Mar 2013	Training	On site training by MAPPER partners. This is not related to the school that will be held in June.	Done, in combination with tutorials this worked very well. Same applies to those projects where we actually have jointly implemented models (e.g. with ScalaLife).
Apr – Jun 2013	Act	Get multiscale simulations on MAPPER and run in production, supported by MAPPER staff.	This worked for all users, although this needed more time, and this is still on-going.
Jul – Aug 2013	Report	We don't expect scientific publications yet on such short notice for the external researchers, but will work on some form of scientific output (report, conference abstract)	This is on-going. We expect scientific output from all collaborations, but the timing is not yet clear. MAPPER keeps supporting the external users and projects.

Table 6: Actions, timing and current status related to dissemination and impact.

Time frame	Action	Comments	Current Status
Feb – Mar 2013	Write	Write up best practices, white papers, etc.	Done, best practices of MAPPER are collected on a website, that works as a FAQ, with pointers to more detailed material and tutorials. Through the experience we gained with the external users and during the second summer school, this material seems to be on target for its intended goal.
Mar – Aug 2013	Act	Publish scientific papers, mapper booths at conferences (e.g. EGI user forum), lectures at conferences (e.g. ICCS, ISC), Multiscale Computing workshop in Leiden, MAPPER school in June, Social Media presence, news items and press releases, etc.	This we did at many occasions, and for details we refer to the dissemination reports.

To wrap up, the roadmap activities have resulted in a stronger focus in the last 8 months of the project on creating impact for the project. This has resulted in more, and more advanced external users and collaborations than originally planned.

6 Future Plans

After MAPPER all project partners will continue to develop their multiscale applications and execute them on e-infrastructures, using best practices, tools and services from MAPPER. They will also continue to support current and new 'external' users, and keep investing in spreading the MAPPER framework and methodologies for multiscale modelling and computing. Moreover, selected partners from MAPPER have committed themselves for continued maintenance and support for key services (the MAPPER Toolset, GridSpace, MUSCLE, QCG-Computing, AHE), in accordance with our sustainability plan.

More important is the further spreading of MAPPER technology. We see strong opportunities in the currently emerging H2020 work programmes in e-infrastructures and FET on the theme of High Performance Computing. The best strategy for MAPPER to create lasting impact, beyond what was already achieved in relation to EGI, is by building on new projects, both driven by application communities and e-infrastructures. We believe that the ideas around HPC centres of excellence would be a perfect instrument to further push MAPPER technology. Some of us are now exploring ideas in this direction.

Most of the MAPPER partners are one way or the other active within large research communities, and also along those lines we are now working to create lasting impact. The H2020 programmes in Health, Materials, Energy, and Climate all provide opportunities for multiscale modelling, and MAPPER partners and collaborators will offer MAPPER as a technology in those programmes.

7 To Conclude

The vision of Distributed Multiscale Computing emerged from a vision on multiscale modelling and computing coming from the FP6 COAST project, and a set of high quality distributed computing services, tools, and middleware from the FP6 projects QosCosGrid and Virolab and the UK funded RealityGrid program. In only three years the MAPPER project not only managed to let the vision materialize in services, tools, middleware and to demonstrate its added value via its portfolio of 7 applications, but also managed to attract a small but growing external use base, as well as to engage in collaborations with 14 projects, and disseminate MAPPER to a number of well established scientific communities. The roadmap activities have considerably contributed to this dissemination and impact of MAPPER.

Appendices

Appendix A: Feedback by external users

Below we reproduce the literal feedback that we received from the external users on our questionnaire.

#	Question	Respondent
1	**Please briefly describe your application. What is its main objective; what is simulated; what are its computing and storage requirements?**	
	Piernik (project homepage: http://piernik.astri.umk.pl/) is an magnetohydrodynamics code using a simple, conservative numerical scheme, which is known as Relaxing, Total Variation Diminishing (RTVD) scheme, and relies on a dimensionally split, second order algorithm in space and time. The RTVD scheme is easily extendible to account for additional fluid components: multiple fluids, dust, cosmic rays, and additional physical processes, such as fluid interactions, Ohmic resistivity effects. Additionally code incorporates selfgravity solver using multigrid method and patch based adaptive mesh refinement. The simplicity and a small number of floating point operations of the basic algorithm is reflected in a performance of 1e5 zone-cycles/s (on single-core 2 GHz processors). The multiscale variant of the code was developed by integrating existing code with the particle module implemented using Monte Carlo method. The code is used in the research of multiple astrophysical phenomena e.g. large scale dynamo processes in galaxies, planetary formation. Due to the complexity and several time and length scales our problems are very computationally demanding, i.e. we can easily utilize thousands of cores and several hundreds of GB per simulation	Kacper Kowalik, Dominik Woltanski
	Our multiscale application is dedicated to modelling of metallic material behaviour in macro-micro scales under loading conditions or/and during heat treatment. This approach is dependent mainly on the size of the problem. Typical examples which compute metallurgical process of stamping (macro scale) and statistically similar representative volume elements (micro scale) take days to finish by using one computing nodes built of two Xeons and 24GB of RAM. The storage is not so important in this case, because the software needs only couple of gigabytes to store the results.	Lukasz Rauch
	I'm interested in modelling and simulation of vascular remodelling processes. This involves coupling models of blood flow with models of tissue mechanics and growth. Blood flow is simulated with a highly parallel lattice-Boltzmann CFD solver (HemeLB). Simulations are run in supercomputers (UK's HECToR mainly). Tissue mechanics and growth are simulated with a sequential code. We currently have a Python script that orchestrates both executions.	Miguel Bernabeu and James Osborne
	MCardio Model is designed to investigate functional consequence of genetic mutation in cardiac ion channels.	MCardio
2	**Did you find the MAPPER infrastructure easy to use?**	
	Very easy to use: 1. Comment: I got great support from Derek Groen in University College London.	Piernik
	Easy to use: 1	Miguel Bernabeu, James Osborne; MCardio
	Not so easy to use: 1 Comment: At first step infrastructure is not so easy and interaction with the software require some explanation, however graphical user interface of supporting tools offer rich functionality.	Metallurgy

4	**From your point of view, what are the main benefits of using the MAPPER infrastructure?**	
	Single, uniform and easy access to heterogeneous computational resources. Switching fully to QCG significantly improved our workflow and saved us a great deal of time.	Piernik
	Possibilities of using HPC infrastructure by scientists, which were previously not involved in such tepy of computing. Facilitation of computing planning, scheduling and monitoring.	Metallurgy
	I can build on the experience of the consortium in running multi-code distributed simulations.	MB/JO
	1. Coupling of heterogeneous codes with MUSCLE. 2. The easiness of using end user tools. 3. The computing resources provided by the infrastructure.	Mcardio
5	**Have you already achieved any scientific or technological results with the support of the MAPPER infrastructure?**	
	Yes: 1	MB/JO
	Not sure: 2. The publication on using MAPPER tools is in preparation; The computing engine is still under development and the obtained results are still not satisfactory for our industrial partners.	Mcardio
6	**If you achieved any scientific and/or technological results, please list them.**	
	Utilizing the MUSCLE framework allowed us to combine two computational methods in a way which was previously impossible. We expect it to result with unique scientific results that will put us in the very forefront of planetary formation field.	Piernik
	I'm currently writing up a paper describing the coupled vascular remodelling model which Derek Groen will coauthor and will talk about the usage of the MAPPER toolkit in our simulations.	MB/JO
7	**If you achieved any scientific and/or technological results, did you publish them online?**	
	No: 3	MB/JO; metallurgy
	Comments: Not yet.	MB/JO;
8	**Do you have any publications for which you used the MAPPER infrastructure?**	
	No: 2	MB/JO; Metallurgy
	Not sure: 1	Mcardio
	Comments:	
	In preparation.	MCardio

Appendix B: More details on interactions between MMM@HCP, SCALALIFE, and MAPPER

On April 22, 2013 delegations from the three project met to discuss their progress and explore options for common actions. This resulted in some technical and scientific collaboration that are reported below, and some ideas on joined dissemination actions.

MAPPER has advanced the state-of-the-art in high performance computing by demonstrating distributed execution of loosely and tightly multiscale models simultaneously on Tier-0, Tier-1 and Tier-2 systems. Through a set of key technologies: a new QosCosGrid middleware, application coupling tools MUSCLE and Grid Space, added to existing system configurations we have shown an elegant approach to the interoperability between PRACE and EGI. Thanks to a set of carefully selected and adopted open communication and data exchange standards in MAPPER we were able to demonstrate the interoperability between MAPPER and external technologies, in particular solutions provided by the ScalaLife and MMM@HPC projects.

Researchers from ScalaLife, in close collaboration with MAPPER, have defined a new multiscale scenario of hybrid QM/MD simulations of molecular properties in complex systems. This advanced use case required a non-trivial dynamic coupling process of two commonly used quantum chemistry and molecular dynamics codes, namely Dalton and Gromacs. Thanks to MAPPER middleware and multiscale application tools in a few weeks we have managed together to implement fully automatic coupling mechanisms between Gromacs and Dalton. More importantly, we moved the codes to production EGI sites with new capabilities added by MAPPER, in particular advance reservation and co-allocation. Both Gromacs and Dalton are powerful molecular dynamics and molecular electronic structure applications commonly used by many users on production PRACE and EGI sites. Therefore, MAPPER jointly with ScalaLife has offered a new added value service in the future for those users who are interested in large scale multiscale experiments based on those software packages.

Additionally, MAPPER has initiated technical meetings with MMM@HPC. As MMM@HPC has focused more on the application side, i.e. developing and running multiscale models, and adopted the UNICORE technology there is still a great opportunity for interoperability in two areas. First, MAPPER could definitely benefit from language-independent and service-oriented MMM@HPC framework, which provides a transparent connection between distributed heterogeneous applications and data storage resources. Consequently, end users could retrieve data from alternative computing resources and databases, often available on campuses, and send them easily for processing to more powerful and co-allocated by MAPPER computing resources in EGI and PRACE. Second, using the OGSA-BES standard provided by QosCosGrid we could demonstrate various loosely coupled multi scale models managed efficiently by the QosCosGrid metascheduler and partially executed on Unicore sites. In this way, we would be able to link together not only PRACE and EGI sites, but also link databases and clusters.

The meeting also resulted in an agreement to have a substantial presence and visibility of MMM@HPC and Scalalife in the second MAPPER summer school. Moreover, the three projects have been inviting representatives to some of their final events (e.g. a CECAM workshop in Jülich and final project meetings).