European Cooperation in the field of Scientific and Technical Research - COST -

Secretariat

Brussels, 11 December 2008

COST 262/08

MEMORANDUM OF UNDERSTANDING

Subject :	Memorandum of Understanding for the implementation of a European Concerted Research Action designated as Cost Action IC0805: Open European network for high performance computing on complex environments

Delegations will find attached the Memorandum of Understanding for COST Action IC0805 as approved by the COST Committee of Senior Officials (CSO) at its 172nd meeting on 24-25 November 2008.

MEMORANDUM OF UNDERSTANDING

For the implementation of a European Concerted Research Action designated as

COST Action IC0805

OPEN EUROPEAN NETWORK FOR HIGH PERFORMANCE COMPUTING ON COMPLEX ENVIRONMENTS

The Parties to this Memorandum of Understanding, declaring their common intention to participate in the concerted Action referred to above and described in the technical Annex to the Memorandum, have reached the following understanding:

1. The Action will be carried out in accordance with the provisions of document COST 270/07 "Rules and Procedures for Implementing COST Actions", or in any new document amending or replacing it, the contents of which the Parties are fully aware of.

2. The main objective of the Action is to develop an integrated approach for tackling the challenges associated with heterogeneous and hierarchical systems for High Performance Computing (HPC).

3. The economic dimension of the activities carried out under the Action has been estimated, on the basis of information available during the planning of the Action, at EUR 76 million in 2008 prices.

4. The Memorandum of Understanding will take effect on being accepted by at least five Parties.

5. The Memorandum of Understanding will remain in force for a period of 4 years, calculated from the date of the first meeting of the Management Committee, unless the duration of the Action is modified according to the provisions of Chapter V of the document referred to in Point 1 above.

A. ABSTRACT AND KEYWORDS

In different fields of science and engineering it is necessary to solve complex and challenging problems with high computational cost. For this purpose, scientists and engineers normally use homogeneous high performance computers. Nowadays, the emergence of heterogeneous computing allows research groups, enterprise and educational institutions to use networks of processors which are already available. On the other hand, high performance computers have become more and more hierarchical and heterogeneous (e.g., a cluster of multiprocessor nodes using multicore processors). These modern hierarchical and heterogeneous computing infrastructures are hard to program and use efficiently, particularly for extreme-scale computing. Consequently, none of the state-of-the-art solutions are able to efficiently use such environments. The Action will establish a European research network focused on high performance heterogeneous computing in order to address the whole range of challenges posed by these new platforms including models, algorithms, programming tools and applications. Indeed, some of the most active research groups of the world on this area are in Europe. The network will contribute to exchange information, identify synergies and pursue common research activities, therefore reinforcing the strength of these groups and the leadership of Europe in this field.

Keywords: High-performance and extreme-scale computing, Heterogeneous computing, Multicore Central Processing Unit (CPU), Graphic Processing Unit (GPU), Hierarchical environments.

B. BACKGROUND

B.1 General background

In recent years, the evolution and growth of the techniques and platforms commonly used for high performance computing (HPC) in the context of different application domains has been truly astonishing. While parallel computing systems have now achieved certain maturity thanks to high-level libraries (such as ScaLAPACK) or runtime libraries (such as Message Passing Interfaces - MPI), recent advances in these technologies pose several challenging research issues. Indeed, current HPC-oriented environments are extremely complex and very difficult to manage, particularly for extreme-scale application problems.

At the very low level, latest generation CPUs are made of multicore processors that can be general-purpose or highly specialized in nature. On the other hand, several processors can be assembled into a so-called symmetrical multi-processor (SMP) which can also have access to powerful specialized processors, namely Graphics Processing Units (GPUs) that are now increasingly being used for programmable computing resulting from their advent in the video-game industry, which significantly reduced their cost and availability. Modern HPC-oriented parallel computers are typically composed of several SMP nodes interconnected by a network. This kind of infrastructure is hierarchical and represents a first class of heterogeneous system in which the communication time between two processing units is different, depending on whether the units are on the same chip, on the same node or not. Moreover, current hardware trends anticipate a further increase in the number of cores (in a hierarchical way) inside the chip, thus increasing the overall heterogeneity, even more towards building extreme-scale systems.

At a higher level, the emergence of heterogeneous computing now allows groups of users to benefit from networks of processors that are already available in their research laboratories. This is a second type of infrastructure where both the network and the processing units are heterogeneous in nature. Specifically, here the goal is to deal with networks that interconnect an (often high) number of heterogeneous computers that can significantly differ from one to another in terms of their hardware and software architecture, including different types of CPUs operating at different clock speeds and under different design paradigms, and also different memory sizes, caching strategies and operating systems.

At the high-level, computers are increasingly interconnected together throughout wide area networks to form large-scale distributed systems with high computing capacity. Furthermore, computers located in different laboratories can collaborate in the solution of a common problem. Therefore, the current trends of HPC are clearly oriented towards, extreme-scale, complex infrastructures with a great deal of intrinsic heterogeneity and many different hierarchical levels.

It is important to note that all the heterogeneity levels mentioned above are tightly linked. First of all, some of the nodes in computational distributed environments may be multicore SMP clusters. Second, multicore chips will soon be fully heterogeneous with special purpose cores (e.g. multimedia, recognition, networking, etc.) and not only GPUs, mixed with general-purpose ones. Third, these different levels share many common problems such as efficient programming, scalability and latency management. Hence, this Action targets the heterogeneity at all presented hardware levels. Moreover, the Action will take a special care of the scalability issues that is a key dimension in the complexity of today environment. The extreme-scale of these environments comes from every level: 1) low level: number of CPUs, number of cores per processor; 2) medium level: number of nodes (e.g. with memory) 3) high level: distributed/large-scale (geography dispersion, latency, etc.) 4) application - extreme-scale problem size (e.g. calculation-intensive or data intensive). Within the timeframe of this Action, it is realistic to expect that large-scale infrastructures composed of dozens of sites, each composed of several heterogeneous computers, some having thousands of more than 16-core processors, will be available for scientists and engineers. However, as of today, the knowledge on how to efficiently use, program or scale applications on such future infrastructures is still vague. This is one of the main challenges that researchers involved in this Action want to take on.

Therefore, a COST Action for high performance and extreme-scale computing in such complex environments is very timely. The main reasons are:

1. There is a huge demand in terms of computational power for scientific and data-intensive applications.
2. The architectural advances offer the potential to meet the application requirements
3. None of the current state-of-the-art solutions in high performance computing allows exploitation to this potential level.
4. Most of the research carried-out in this area is fragmented and scattered across different research teams without any coordination.

COST is indeed an appropriate framework for this Action. Its main goal is to overcome the actual research fragmentation on this very hot topic by gathering the most relevant European research teams involved in all the scientific areas described above (from the CPU core to the scientific applications) and coordinate their research.

Due to the vertical aspects of this project - integration of different research made at different level of the architectural layers - most of the other European research frameworks are not suited as they focus on effectively doing the research or they are targeted to only one layer. The organization of meetings and workshops with participation of researchers working on the different levels of this vertical structure would contribute to forecast the ideas to progress in the common direction of the research groups at the different levels and to the efficient use and integration of the systems composing the different levels. Furthermore, the organization of schools (with participation of some of the most relevant researchers in the field in Europe and all over the world) would contribute to the formation of a group of young scientists specialized in the complex field. Finally, there is also a clear horizontal component for this Action The groups involved are dealing with both the theoretical and applied aspects of heterogeneous computing and therefore their interactions within the Action are expected to result in cross-fertilization among different fields of the theory of heterogeneous computing, and also among different application domains which can significantly benefit from the research and training activities of this Action.

Summarizing, this Action allows expecting some potential benefits such as: high-level scientific results in the very important domain of high-performance and extreme-scale computing in complex environment; strong coordination between different research teams with significant expertise on this subject; a better visibility of the European research in this area and a strong impact for other scientists and high performance applications.

B.2 Current state of knowledge

In the past, the research that has been conducted in high performance computing in complex environments addresses four key topics shaping this field: modelling and algorithm design; performance analysis; programming tools and applications.

First, since mid 90s, when early pioneering works were published, the design of fundamental algorithms for high-performance computing on a set of heterogeneous processors has made a significant progress. Designed algorithms were mainly aimed at numerical linear algebra and scheduling and mapping of applications on heterogeneous computational clusters. New models were required to develop such algorithms. Most studied are the algorithms based on the simplest performance model representing the target complex heterogeneous platform by a set of heterogeneous processors, each characterized by a single positive number representing its speed. More realistic models, such as the functional model and the band model, which take into account the complexity and heterogeneity of the memory hierarchy, have been also proposed. A number of fundamental efficient algorithms based on these models have been also designed and analyzed.

Second, the problem of performance analysis of the new heterogeneous parallel algorithms has been also addressed, and a number of methods, both theoretical and experimental, have been proposed. The heterogeneous algorithms are based on heterogeneous performance models characterized by a large number of parameters. Therefore, their implementation is not a trivial task as it includes the accurate and efficient estimation of these parameters. The problem of accurate and efficient estimation of the performance models of heterogeneous processors has been also studied and some acceptable solutions have been found for many practically important cases. The performance models used in the design of these algorithms only allowed for very basic optimization of the communication and computation cost. The work in this area is mainly based on analytical communication and computation performance models. The main body of research on the optimization of the communication cost deals with MPI and the optimization of MPI collective communication operations. In this research, traditional performance models, such as the Hockney model and LogGP, were used. Moreover, some research has been done so far in modelling and optimization of communication operations in distributed environments (e.g. GridMPI, MPICH-G). However, the full implementation of all MPI collective communications for heterogeneous and hierarchical environments, such as heterogeneous clusters, clusters of multicores, global networks interconnected via Internet and/or dedicated communication channels is lacking.

Third, very few programming tools facilitating implementation of parallel and distributed scientific applications in heterogeneous environments have been designed and implemented so far. The examples are the mpC programming language and the HeteroMPI library designed for parallel computing on heterogeneous computational clusters. In general programming environments address one issue (multithreaded programming with OpenMP, high-performance linear Algebra with ScaLAPACK or message passing with MPI, large-scale with distributed objects) but there are no tools that cover every complexity level of the environment: there is a clear lack in bringing together application designers and programming system designers in order to discuss issues, to teach each other and finally to come up with programming systems that would be accepted and used by high-performance scientific programmers.

Finally, while some high-performance scientific applications for complex heterogeneous environments, mainly for heterogeneous clusters, have been designed and implemented, this area is definitely underdeveloped. For instance, most of the successful applications using large-scale systems, such as seti@home, use a trivial model of parallelism where each task composing these applications can be processed in parallel. Other examples are workflow engines for executing parallel applications on large-scale system, but there is a clear lack of standards. Hence, executing tightly coupled applications is a major challenge that has not been fully addressed so far.

Today, there is some research carried out at each level of the architecture. For instance the ICT-2007.3.4 call on computing system focuses on multicore architecture. Other 7th Framework Programme (FP7) calls cover higher level such as ICT-2007.1.2 that targets service oriented architecture. It is noticeable that none of these calls have a vertical dimension that tries to cover all the layers with the goal to provide integrated solutions that use and master the complexity of each level of today complex infrastructures. It clearly appears that a new approach is necessary: gather researchers from every architectural layer and propose integrated solutions (algorithms, libraries, applications) that efficiently take into account the architectural properties of each level.

B.3 Reasons for the Action

Despite the inherent complexity of current heterogeneous systems such as those described above, an efficient use of such complex environments is necessary to enable the solutions of large computational problems. Unfortunately, whereas applications require more and more computational power, none of the available state-of-the-art solutions for programming distributed systems can fully exploit the available features of modern heterogeneous environments. As a result, there is a clear need to develop new tools for effective programming of highly heterogeneous HPC systems, an activity which has been largely fragmented in Europe due to the lack of effective coordination mechanisms between groups which have been partly addressing this problem from different perspectives.

In this regard, the main goal of this Action is to coordinate European groups working on the use of heterogeneous and hierarchical systems for HPC and to develop collaborative activities among the groups involved in this research topic. Thanks to this effort for defragmentation and coordination of the European research in this area, the expected results of this Action focus on efficiently addressing the problems of high-performance and extreme-scale computing in heterogeneous, hierarchical and distributed environment (efficient and easy usage of these environments, grounded software and applications, etc.).

To achieve these goals, the Action will launch several concrete activities such as schools, workshops, meetings, visits, etc. where scientists will meet each other and work on the scientific problems identified as major challenges in this domain.

It should be noted that more than 30 research groups from 16 COST countries have already expressed their interest in this Action. The Action remains open to new participants interested in this research area, with special emphasis on the incorporation of members coming from countries which have recently joined the European Union and partners from less-favoured regions.

B.4 Complementarity with other research programmes

It is noticeable that this Action is complementary to the other European research efforts in this area. First and foremost, the focus of the Action is on research coordination while most of the calls within FP7 are focused on research and infrastructures which are not specific targets of this Action. Second, this Action targets high-level software design and applications while most of current research projects are targeted towards the architectural and hardware level such as (Call 1 ICT Objective 3.4 Computing Systems: 8 projects have started including a Network of Excellence and 7 Specific Targeted Research Projects covering the topic “Novel architectures for multicore computing systems”). Therefore the goal of this Action is to coordinate researchers working at different levels (e.g. multicore, GPU, clusters, multicomputers and large-scale environments) in order to identify common aspects of their works in the development of software optimized for these systems.

C. OBJECTIVES AND BENEFITS

C.1 Main/primary objectives

The main objective of the Action is to develop an integrated approach for tackling the challenges associated with heterogeneous and hierarchical systems for High Performance Computing.

C.2 Secondary objectives

This general objective will be accomplished through a set of specific objectives listed below:

1. Improving the quality of the research activities in this field in the countries participating in the Action.
2. Identifying the trends and needs.
3. Tracking the current activities in the field.
4. Establishing a common view, which may ultimately lead to an FP7 proposal.
5. Creating and stimulating a discussion forum.
6. Fostering standardization procedures in the field.
7. Developing prototypes systems to test the proposed standards.

A summary of the main deliverables is:

1. Studies of the research in this field, investigation of the trends as well as the industrial and scientific interest. These will be centred on the themes of the different Working Groups. These studies would contain (these proposals must take into consideration the fact that the systems are organized in a hierarchical way, and so the proposals must be general or contain proposals of integration between the different levels):
- A proposal for standardization of models for heterogeneous computing, which could be used to develop efficient routines, and for adjusting the models to different types of systems: networks, distributed systems, peer-to-peer systems.
- A proposal for the development of efficient and friendly heterogeneous routines based on the previously addressed models. The routines should have user-friendly interfaces to encourage wide use in solving scientific problems.
- A proposal for the development of libraries for solving problems on heterogeneous and dynamic systems. These libraries should integrate the previously analyzed routines, and could be used to solve basic problems to solve scientific problems, such as basic communications routines, dense and sparse linear algebra problems, optimization problems, etc.
- A proposal of the necessary tools for the efficient management of heterogeneous systems: mapping tools, tools for the dynamic analysis of the system, etc.

2. A website with all the information: the previously mentioned studies, information about the workshops, the meeting, the schools, and the mobility programme.

3. Reports on the training activities and the meetings. All the material generated in the schools will be made available on the website, as will the guide lines decided in the meetings.

4. To form a group of young researchers in this complex and challenging area, which could lead the future research in Europe.

5. Detailed studies on the impact of applying heterogeneous computing practices to real problems of relevance in Europe.

An important outcome of the effort will be the coordination and guidance of the research for the development of standardized routines to solve scientific problems in heterogeneous systems. The previous proposals could guide the research of the different research groups in a common direction. The application of the proposals to problems in which the participants' teams are currently involved will be analyzed. Because application programmers (solving computationally intensive scientific and engineering problems) are the end-users of this project, the Action will be open to new groups which could collaborate in all the mentioned aspects, but especially in the application areas.

C.3 How will the objectives be achieved?

The main goal is to promote collaboration through science meetings, workshops, schools and mobility of researchers. This will allow the exchange of ideas and the mobility of researchers. To achieve the goal of this Action, the following concrete activities are foreseen.

1. Summer Schools: two Summer Schools will be organized. One will focus on management and programming of large-scale environments and another on heterogeneity and hierarchy in high performance computing systems. Ph. D. students and young researchers from the groups involved in the Action will participate, as well as members from other groups. The schools will provide young researchers with the opportunity to share information and knowledge and to present their current research. These schools should contribute to increase the computing community and spread

European knowledge. The schools will be held in the second and third year, as it is important to train researchers at an early of the Action. The talks and the works presented will be made available on the website of the project.

2. International workshops: Heteropar and HCW are the most recognized workshops in the area of heterogeneous computing. They are co-located with Euro-Par and IPDPS, respectively, two major High Performance Computing conferences. Most of the participants in the Action usually attend these workshops. The Action will actively contribute to these 2 workshops, for instance through invited speakers. Such speakers will outline the main research directions in computational aspects of heterogeneous and distributed computing, and contribute to defining future roadmaps. Between four to eight invited presentations will be organized per year. A summary of the presentations will be published in the workshops proceedings.

3. Working Groups meeting: The scientific work plan will be divided among different Working Groups. Each Working Group will have a substantial autonomy in terms of research projects. A leader nominated by the Management Committee will drive each Working Group. Members of a given Working Group will meet from one to twice a year to discuss and exchange on specific scientific issues and problems.

4. European working meetings (Joint MC&WG meetings): four working meetings would be organized, one each year. These meetings will be devoted to studying some specific aspects of heterogeneous computing of common interest to all WGs. The main theme will be decided by the members of the group, and opinions will be gathered via Internet using the project website, beforehand, with the ultimate goal of enhancing participation of external groups and potential new members of the Action. After each meeting, the studies will be summarized and placed on the web page.

5. Management Committee meetings: The first meeting will take place in Brussels and the others will take place at the same time as the European meetings. They will be devoted to the organization of the network and will ensure the scientific quality of the network.

6. Short-Terms Scientific Missions: There will be visits by young students to foreign laboratories and departments, and visits of high level researchers to centres with a large number of young researchers. This Action will target mainly young researchers, who would receive a cross-disciplinary training and take advantage of existing resources (e.g., training programmes locally offered to MSc and/or PhD students and post-doctoral fellows at the local institutions). This will increase competitiveness and career development of those scientists in this rapidly developing field through cutting-edge collaborative research on the topic.

C.4 Benefits of the Action

The overall benefit of the Action is to foster the collaboration of the most European active research groups in the field of heterogeneous and hierarchal computing. More precisely, on the scientific side, it is expected that, the Action will:

1. Overcome research fragmentation.
2. Facilitate the use of complex high performance computing environments.
3. Use more efficiently these environments.
4. Bridge the gap between theoretic research and real world applications thanks to cross-fertilization.
5. Train a critical mass of young researchers that will represent the new generation of specialists in the use of this emerging technology.

C.5 Target groups/end users

The Action will also provide significant benefits to external research groups. Specifically, collaboration with different scientific groups in the efficient solution of highly computationally demanding problems on parallel systems will be fostered. The infrastructure, methodology, software, knowledge and expertise, developed in the context of the network would provide significant benefits to those groups and their scientific fields. Examples of the scientific fields which are expected to obtain significant improvements from the techniques and methodologies developed within the Action are: satellite imaging (i.e. parallel processing of scientific data from the surface of the Earth using remote sensing), with applications on climate change analysis; environmental monitoring; analysis of natural and man-induced disasters such as oil spills and other types of chemical and atmospherical contamination; meteorology and other highly relevant fields such as signal processing, quantum chemistry, econometrics, nuclear physics, multimedia processing and medicine, etc.

D. SCIENTIFIC PROGRAMME

D.1 Scientific focus

The expected scientific impact of the Action: is to create a community of researchers dedicated to HPC: to focus the research on hot topics and applications of interest for Europe; to propagate the collaboration of research groups with industry; to stimulate the formation of new groups in new European countries; to facilitate the solution of highly computationally demanding scientific problems as mentioned above. For this reason, the groups involved in this Action will collaborate with several scientific and industrial groups that could benefit from the advances propitiated by the Action itself, and the incorporation of new groups to the network would be prompted.

In details, the Action will target work on the following scientific tasks:

- Numerical analysis for hierarchical, heterogeneous and multicore systems (such as linear algebra, differential equation solvers, self-adaptive numerical libraries).
- Efficient use of complex systems with an emphasis of computational library and communication library.
- Algorithms and tools for mapping and executing applications onto distributed and heterogeneous systems.
- Applications.

To achieve these research tasks, different leading European research teams will participate to the concrete activities detailed in section C. Summer Schools will help early-stage researcher to acquire the basic knowledge to pursue these research activities and to gather together groups that not necessarily know each other. The international workshops will serve as a forum for spreading new ideas and exchange results on these tasks. Working Group meetings will be a place for the exchange of new ideas and results on specific topics. Each European meeting will be devoted to one aspect of these scientific tasks and therefore will stimulate exchange on these topics between the research groups. Finally Short-Term Scientific Missions will enable bi-lateral research to be conducted on highly focused topics.

D.2 Scientific work plan – methods and means

Four Working Groups will be organised to coordinate the scientific research. Naturally these Working Groups follow the scientific tasks described in section D.1:

1. Numerical analysis for hierarchical and heterogeneous and multicore systems.
2. Libraries for efficient use of complex systems with an emphasis of computational library and communication library.
3. Algorithms and tools for mapping and executing applications onto distributed and heterogeneous systems.
4. Applications of hierarchical-heterogeneous systems.

The organization of these Working Groups will be detailed in section E.2. However, it is important to note that these Working Groups target vertical aspects of the architectural structure outlined in section B.1. For instance, the Action will carry out work on numerical analysis at the multicore level but also at the heterogeneous systems level as well as at the large-scale level. The last Working Group (Application) will benefit from research of the other three ones.

Here is for each Working Group a detailed description of the objectives as well as examples of research issues that will be tackled and of how the research tasks will be carried out.

WG1. Numerical analysis for hierarchical and heterogeneous and multicore systems.

The goal of this Working Group is to tackle the scientific problems that arise in the fields of linear algebra, differential equation solvers, pre-conditioners, etc. Such problems are of key importance when high-performance is the target. Moreover, in this Working Group, participants will address these issues at every level of the architecture (from the core to the large-scale). The long-term objective is to have a set of libraries and tools that will efficiently use nowadays complex systems. To achieve this long term objective the Action will gather European researchers who are specialized in this area and have them work together on subtopics. Here are two examples of synergies that have already been identified for this Working Group and the appropriate work plan.

a. Linear algebra for multicore systems

The main emphasis of linear algebra (LA) libraries for future multicore processors will be on shifting the focus in expressing numerical algorithms towards approaches based on (large grain) data flow analysis and out-of-order execution methods. The approach will rely on data dependency analysis. This work seeks a simple, yet flexible and powerful mechanism for expressing data dependencies (initially in the limited context of numerical linear algebra). The main goal will be to maximize performance by making maximum use of the potential for parallel execution of non-dependent operations. It is anticipated that the mechanism will facilitate adaptivity in the level of data granularity and granularity of the parallelization in order to accommodate different levels of coupling of the processing element within the system. The central challenge in developing a methodology for programming multicore systems is to create models that allow high flexibility for experimenting with techniques (e.g., addressing vectorization and asynchrony, dynamic/adaptive out of order execution, hiding memory latency, blocked storage, heterogeneity, etc, as described in the following subsections) and consequently, provide a framework for discovering optimal solutions. In addition, these models will raise the level of abstraction in developing efficient implementations for multicore systems and thus reduce the time and effort to develop libraries of highly optimized performance.

Key objectives:
- To first develop prototypes: loose on frameworks/rules, intended as race car implementations to study the limits of the used techniques, and gain insight into the problems faced.
- To base implementations on explicit threads (vs. higher level multi-threaded API) within a single chip multicore and on explicit communication.
- To design a draft for the multicore framework that would allow transfer of the technology to a large collection of numerical libraries (and thus enormously increase the impact of the research).

b. Solution of differential equations for heterogeneous systems

Many simulations from scientific computing are based on the numerical solution of ordinary differential equations (ODEs) or partial differential equations (PDEs). For homogeneous environments, efficient parallel solution methods have been developed which rely on a fixed assignment of computations to processors. For heterogeneous systems, such a fixed assignment often leads to non-satisfactory execution times. In this context, the main focus of the work will be (1) development of adaptive techniques for the scheduling of computations of typical ODE or PDE solvers; (2) development of a precise cost model for execution modelling of computations of ODE and PDE solvers; (3) development and integration of self-adaptivity techniques for ODE and PDE solvers with an automated selection of an efficient implementation variant for a given execution situation; (4) integration into a software library for multicore cluster systems.

Key objectives:
- Provision of an adaptive scheduling toolset for ODE and PDE solvers; the toolset should be generic in the sense that it can be parameterized and used for arbitrary ODE and PDE solvers;
- Provision of a toolset for the dynamic acquisition and modelling of cost information of computational parts of ODE and PDE solvers;
- Development of a support library for the integration of self-adaptivity into ODE and PDE solvers;
- Development of a software library for the parallel solution of typical ODE and PDE solvers on heterogeneous multicore cluster systems providing different variants of the different solution methods with different execution characteristics.

WG2. Efficient use of complex systems with an emphasis of computational library and communication library.

Efficiently programming complex systems is an extremely difficult challenge. The development of applications that use all the power of the available architecture and that are also portable requires the development of low-level libraries. With these libraries, programmers will easily design new applications that will work on today and future architectures. This Working Group will work on the design, implementation and test of libraries for an efficient use of modern complex systems. In particular, emphasis is given to GPUs (Graphical Processing Unit), multicore processors and heterogeneous network.

a. Heterogeneous systems supporting stream-based computing

Stream programming naturally expresses the parallelism inherent in a program by decoupling computation and memory accesses. Data is firstly gathered into a stream and processed by one or more processing kernels, each kernel exploiting its inherent data parallelism, in order to scatter back the live data into an output data stream. Nowadays, an increasing number of powerful streaming-oriented processors, such as the Graphics Processing Units (GPUs), are already available in commodity computers. Current GPUs perform more than one order of magnitude faster than dual-core CPUs and have been also applied for general purpose processing (GPGPU), namely for physics simulation, signal processing, computational geometry, database management, computational biology, computational finance, and in computer vision. The main goals of this Working Group will be 1) to establish a general model for stream-based computing; 2) to develop tools for programming and remotely using streaming-oriented processors, namely GPUs; 3) to integrate the developed models and tools with established programming tools and interfaces, such as CUDA and MPI; and 4) to evaluate the developed tools by programming applications with different characteristics.

Key objectives:
- Study on how to model and program these powerful heterogeneous computational systems.
- The developed models and tools for stream-oriented distributed computing will be integrated with well established programming tools and interfaces, such as CUDA, Caravela, and MPI.
- To program applications with quite different characteristics, in what concerns the amount and type of exhibited parallelism and processing regularity.
- The relative performance of many-core GPUs will be evaluated for the programmed applications.

b. Multicore Toolkit

The main goal of the work will be to 1) establish the viability of multicore systems as a mainstream parallel/distributed computing engine, 2) develop a useable open source toolkit for applications development, 3) bind the toolkit into high-level programming interfaces such as MPI and OpenMP, 4) demonstrate the use of the toolkit in developing significant parallel applications in numerical and non-numerical problem domains, 5) make the toolkit available to academic community (and through appropriate licensing commercial vendors/users). The Multicore Toolkit (or MCT) will contain the essential tools and libraries needed for developing programs and applications executing on commodity clusters of multicore processors. To scope the tool set, the restriction will be to data parallel and message passing programming models, which underpin a significant proportion of current high performance applications.

Key Objectives:
- Based on the current and emerging multicore system architectures produce a generic low-level API, with accompany drivers, to cope with issues such as process synchronisation, caching strategies, thread affinity, and overlapping communication and computation.
- Bind the new multicore API into to work with OpenMP and MPI,
- Create a basic numerical style library that is optimised for a set of commodity multicore systems.
- Create a support library for irregular based computational problems such as data mining and combinatorial optimisation problems.

- Create support for visualization such as video coding, movement tracking, image manipulation and reconstruction.

WG3. Algorithms and tools for mapping and executing applications onto distributed and heterogeneous systems.

The resources that compose the complex systems targeted in this Action are of three types: computing resources, network resources and storage resources. The performance of an application greatly depends on the way these resources are chosen to execute it. The goal of this Working Group is to work on algorithms for mapping, scheduling and executing applications on large-scale and complex systems. Members of this Working Group will work on creating new algorithms for this area taking into account the complexity of today infrastructure as well as the different criteria that need to be optimized in this context. Examples of researches that will be carried out in the Working Group are given next.

a. Mapping of parallel applications

In order to achieve high performance, critical aspects of the heterogeneous computing systems are to efficiently assign application tasks to the most suitable machines and to determine a scheduling execution order of the tasks on each selected machine. Existing models are not sufficient as they mostly focus only the job level for homogeneous (HPC) or heterogeneous (large-scale systems) environments. However, future infrastructures need an integrated approach with a more fine grain resource management that includes more insight knowledge on the parallelism of the applications. The complex hierarchical structure of communication bandwidth and latency in hierarchical system will need suitable scheduling strategies. Moreover due to their size these environments are subject to uncertainties. Within this Action, new approaches for mapping and scheduling parallel applications on large-scale HPC systems with thousands of nodes consisting of thousands of cores will be explored. The approach includes multi-criteria optimization, adaptive re-allocation of applications, fault tolerance and reliability.

Key objectives:

- Design of new scheduling models that take into account heterogeneity, the hierarchy and the scale of the environments
- Design of scalable scheduling algorithms
- Manage uncertainties (failures, random submission, etc) at the scheduling level.

b. Multi-criteria scheduling of workflows on large-scale systems

A lot of parallel applications (common in the domains of image processing, computer vision, query processing, etc) that follow simple dependence patterns, such as pipeline, fork-joins or series-parallel dependence graphs can benefit from execution on large-scale heterogeneous platforms. Typically, such applications can be expressed in terms of algorithmic skeletons that the programmer can invoke from a library. Parallel applications expressed as skeletons typically operate on a very large number of data sets. The Action will also target the simultaneous execution of several application workflows which operates each on continuous streams of data sets. Energy minimization complicates the situation when added to the list of the objectives. Instead of selecting the fastest available processor, the goal is to achieve the best trade-off between performance-oriented metrics (be it throughput, latency or stretch) and energy consumption. The goal is to investigate the trade-offs between performance and energy-consumption, and to design robust scheduling strategies that achieve optimal solutions by focusing on multi-criteria approaches (i.e., minimizing the energy under performance-oriented constraints, or the converse). When dealing with several application workflows simultaneously, aiming at mixing objectives and achieving good levels of throughput, stretch and energy consumption for each application is a very challenging algorithmic problem indeed.

Key objectives:

- Overview of platform and application models, choice of performance-oriented objectives
- Survey of current approaches for energy-minimization. o Algorithms and heuristics for multi-criteria scheduling of single pipeline workflows
- Algorithms and heuristics for multi-criteria scheduling of multiple pipeline workflows
- Extensive simulations and experimental comparison of scheduling heuristics.

WG4. Applications of hierarchical-heterogeneous systems.

Nowadays applications process larger and larger data sets. Therefore, they require a larger number of processing units as well as more and more memory. This Action challenges the efficient execution of applications onto complex infrastructures. A set of applications that potential members of this Working Group already agreed to develop is given here. Due to the nature of the target infrastructures (hierarchical, heterogeneous and large-scale), addressing this challenge requires solving many issues. Therefore, this Working Group will use the results of the three Working Groups described above to design and implement the considered applications when necessary. It is expected that the application described above will raise new problems that will be addressed in the other Working Groups.

a. Molecular Dynamics Simulations on large Scale Heterogeneous platforms

Molecular Dynamics simulations require a huge amount of computations and standard parallel versions (such as NAMD), based on domain-decomposition techniques, fail to scale well due to the corresponding synchronization requirements. Recently, new algorithms for molecular dynamics simulations, such as ABF (Adaptive Biased Force) have been proposed. These algorithms involve more computations, but they are better candidates for scalability purposes. The goal of this research is to demonstrate that it is possible, by changing the algorithms, to execute applications, on large-scale platforms, that are considered as tightly coupled using traditional algorithms.

Key objectives:

- Find a set of local parameters that define a realistic model of the platform, and propose a prototype to instantiate these parameters.
- Use these parameters for independent tasks distribution and collective communications. In particular, build overlay networks suited to these applications.
- To efficiently use distributed heterogeneous resources, design high-level services to group resources to perform a given task.
- Design efficient overlays to perform complex requests such as finding intermediate data spread across the platform.
- Build a scalable molecular dynamics simulation software and integrate it to NAMD software suite.

b. High-dimensional data processing in remote sensing

Remote sensing-based satellite imaging is a key application in which heterogeneous computing is starting to play a very relevant role. The development of distributed architectures for efficiently storing and processing the vast amount of satellite data that is now being collected across Europe is expected to introduce major advances in the systems and tools currently being used by agencies and institutions such as the European Space Agency (ESA). Several researchers involved in the Action are already collaborating on the application of heterogeneous computing to this problem and would share their expertise within the Action. The main features of the work related with remote sensing data processing that will be carried out by the Action will be to: 1) establish the viability of using heterogeneous computing platforms (including heterogeneous networks and large-scale platforms) to improve existing methodologies; 2) explore the viability of using heterogeneous specialized hardware platforms (such as networks of GPUs) for onboard data processing and compression; and 3) develop standardized software libraries for enhanced distributed processing of remotely sensed data set.

Key Objectives:
- Explore the integration of large-scale heterogeneous platforms and distributed systems in a new application domain that can greatly benefit from the use of heterogeneous computing practice.
- Based on the current and emerging multicore system architectures and GPUs, develop new heterogeneous algorithms and techniques for advanced remote sensing data exploitation.
- Develop a standardized library for heterogeneous parallel processing of remotely sensed data.
- Validate the newly developed library in the context of different remote sensing-based application domains with high social relevance, such as disaster monitoring and prevention.

E. ORGANISATION

E.1 Coordination and organisation

This Action will focus on overcoming fragmentation and unifying research efforts in the area of high performance in complex systems, guiding them in a common direction. This Action will act as a coordination effort within the research programmes of each existing groups and will foster synergies. In this way it will be easier to obtain standards of interest for the global scientific community. Given the nature of this collaboration, COST is the most appropriate framework.

An outline of the management activities is provided below:
- A Management Committee will be composed of up to two members from each participating COST country. The Management Committee (MC) will meet at least once a year. It will ensure that the scientific objectives of this proposal are achieved. At the first meeting, one of the participants will be elected as the Chair of the MC. Another member will be nominated for the development and maintenance of the Action website. Moreover, the MC will nominate a coordinator for each Working Group.
- The Chair of the MC will steer the project, interface the activity requests with the MC as well as prepare the intermediate and final reports.
- Each year, the MC will plan the activity for the coming year (workshops, schools...). The MC will appoint Coordinators for each particular activity and those will be responsible to promote it and to coordinate their activities with the other groups.

Concerning research coordination, the Action will work as follows.

- Each group and institution of this Action will take part at least into one Working Group. The Working Group membership will be decided on the basis of the research focus and interest of the group. As Working Groups cover vertical research activities it is expected that different communities (e.g. multicore specialists, large-scale system researchers, etc.) will be part of the same Working Group. Moreover, as each Working Group focuses on one aspect of the research tasks this Action is addressing, a good synergy between the participants will be created.
- The MC will appoint a coordinator for each Working Group. At least one Working Group meeting will occur every year during the timeframe of this Action. The first meeting will be devoted to the presentation, by each member, of the research interests that fit in this Action. Therefore, after the first meeting, synergies between each participant will be defined. The goal of further meetings will be to provide a forum for exchange and develop new ideas on the subject. Young and senior researchers will be invited together to present results on the topics of the Working Group. Some meetings will be focused on a specific aspect of the Working Group while other meetings will cover the entire scope of the group. In summary, Working Groups will be the cornerstones for implementing national research coordination.
- Thanks to these Working Group meetings, bilateral common research topics will emerge. Collaborations between two research groups will be enabled by the Short-Term Scientific Missions. Such Short-Term Missions will normally last several weeks and will be an ideal framework to conduct deep research on a given subject.
- Once a year there will be a joint MC and Working Groups' meeting for all the participants of this Action. This meeting will cover a particular aspect of the research themes of the Action and will also cover all the Working Groups' areas of interest. As an example, the theme of a meeting could be "Libraries for heterogeneous systems", and the points to discuss: format of the routines, user interface, libraries architecture, system description, mapping strategies, etc.
- Since the research would be guided by applications, the most challenging applications will also be analyzed in a specific meeting, where researchers and representatives of the industry will share experiences about how heterogeneous computing can be applied efficiently to real problems.

Finally, a website that gathers important information on the management and organization will be set-up. In this website, the members of the Action will find:

- Aims and objectives of the Action.
- List of researchers involved and their research objectives.
- Public document such as reports, talks and international publications.
- An internal wiki and blog.
- Internal documents such as white papers, research reports, meeting agendas, meeting reports, drafts, etc.
- Deliverables as described in section C.1
- Possibility to apply for a short-term scientific mission (with CV and research program)

E.2 Working Groups

The research program presented in section D is divided into 4 Working Groups:

1. Numerical analysis for hierarchical, heterogeneous and multicore systems (such as linear algebra, differential equation solvers, self-adaptive numerical libraries).
2. Efficient use of complex systems with an emphasis on computational library and communication library.
3. Algorithms and tools for mapping and executing applications onto distributed and heterogeneous systems.
4. Applications.

Each Working Group will be organized in the same way. A Working Group Coordinator will be appointed by the Management Committee. The participants of the Action will decide to join a Working Group on the basis of the scientific focus of this Working Group and the research interest of the participants. Working Group meetings will be organized at least once a year. The Working Group on Applications will also have a special role as it will serve as an integration forum for the research carried-out by the other Working Groups (for instance a library developed in the second Working Group could be used in some of the targeted applications).

E.3 Liaison and interaction with other research programmes

The following key fora are the most likely to benefit and/or provide interactions for mutual interest:
- DEISA2 - Distributed European Infrastructure for Supercomputing Applications (FP7 -Capacities) currently encompasses most high-end systems in Europe available for public research. It aims at delivering a turnkey operational solution for a future European HPC ecosystem. The cooperation is envisioned by the exchange of project information and data, especially useful for the foreseen COST studies.
- EGEE-III - Enabling Grids for e-Science (FP7 - Capacities) gathers the largest multi-disciplinary Grid infrastructure in the world. Several applications studied by the Action will run on EGEE-III infrastructure and the results will be disseminated in the frame of EGEE events.
- EGI_DS - European Grid Infrastructure Design Study (FP7 Capacities). The main foundations of EGI are the national Grid initiatives. The cooperation is envisioned by the exchange of relevant data and methods, since both EGI_DS and this Action are collecting data about the current relevant applications in the field.
- PACE - Partnership for Advanced Computing in Europe (FP7 - Capacities) prepares the creation of a persistent pan-European HPC service, consisting of the most efficient European HPC service centres providing researchers with access to capability computers and forming the top level of the European HPC ecosystem. The Action will provide use cases of interest for PRACE.
- In the particular field of multicore systems, the Action members will cooperate with the projects running in the frame of the FP7-ICT-3.4 Computer Systems programme, by a continuous exchange of information about applications and offering new results of the Action for possible implementation in the frame of the ICT projects. The list of projects of interest includes six projects covering the topic "Novel architecture for multicore computing systems". The links with the Network of Excellence HIPEAC, which covers key aspects of architectures, compilers, operating systems, and programming environments, are essential for the Action. Through the expected good collaboration, the tools for multicore systems provided by the VELOX and Apple-CORE (focusing on parallel programming), respectively MERASA, JEOPARD and ICT-eMuCo (focusing on real-time systems and mobile computing systems) are expected to be adopted by the Action members early in the testing phase.

- In the application field, cooperation with several projects is expected. For example, in the particular field of satellite image processing, experts from the recently funded Marie Curie Research Training Network "Hyperspectral Imaging Network" (FP6), coordinated by the University of Extremadura and with a strong component on high-performance computing, will be invited to explore synergies among the two programs.

E.4 Gender balance and involvement of early-stage researchers

This COST Action will respect an appropriate gender balance in all its activities and the Management Committee will place this as a standard item on all its MC agendas. The Action will also be committed to considerably involve early-stage researchers. This item will also be placed as a standard item on all MC agendas.

In all cases, encouragement of applications from all sectors of the community (including minorities) will be clearly stated. There will be no discrimination based on sex, race, religion or citizenship. The systems developed also provide means to improve equal opportunities for disabled persons. This will be further strengthened by encouraging applications from disabled young researchers and, e.g., improvements in e-learning concepts for this purpose.

F. TIMETABLE

This Action lasts for four years. MC meetings will take place at least once a year. Each Working Group will meet between once and twice a year. The planning of Short-term Scientific Missions (STSMs) will start after the first Working Group meetings, once participants will have identified common research interest. Two Summer Schools will take place (one in Year-2 and one in Year-3). Year 4 will also have a final meeting.

In this Action the main milestones for the first year will be the following (M being the month when this Action will start):

Date	Event	Milestones and notes
M+0	First Management Committee meeting	Nominate WG Coordinators, decide on the upcoming activities, nominate coordinators for these activities.
M+2 to M+3	First Working Group meetings.	Identify scientific synergies, present work plan, decide topic of next meeting.
M+4	Launch of the first Short-term Scientific Missions' call.	Allocation of STSMs
M+6	Working Group meeting and Management Committee meeting.	Presentation of early results, identified synergies, website, presentations and discussions on the theme of the meeting.
M+9 to M+10	Second WG meetings.	

At M+12 the Management Committee will meet again and decide the organization and milestones for the second year. If the organization of the first year is found successful this will be applied to the second year (with the inclusion of the first school, and the start of long-term visits) otherwise another organization will be sought. Hence, for the second year, a possible timetable is:

Date	Event	Milestones and notes
M+12	Third Management Committee meeting.	Delivery of first annual report, nominate coordinator for the summer school and decide its location.
M+14 to M+15	Third Working Group meetings.	
M+6	Launch of the second Short-term Scientific Missions' call.	Allocation of STSMs
M+18	Working Group meeting and Management Committee meeting.	
M+19	First summer school.	
M+21 to M+22	Fourth WG meetings.	

G. ECONOMIC DIMENSION

The following COST countries have actively participated in the preparation of the Action or otherwise indicated their interest: BE, BG, DK, DE, ES, FR,GR, HU, IE, IT, PL, PT, RO, SI, UK. On the basis of national estimates, the economic dimension of the activities to be carried out under the Action has been estimated at 76 Million € for the total duration of the Action. This estimate is valid under the assumption that all the countries mentioned above but no other countries will participate in the Action. Any departure from this will change the total cost accordingly.

An institution from Russia has also shown interest in taking part in this Action.

H. DISSEMINATION PLAN

H.1 Who?

This Action aims at high performance computing on nowadays complex infrastructure. Therefore the main target audience is twofold. Results in terms of development of library, algorithms or applications as well as deliverables and standard will be usable both in the academic world and in the industry.

Concerning the Working Group on numerical analysis, many academic groups that have high computational needs could benefit from the results of the Action (e.g., astronomy, physics, biology, etc.). Moreover, software companies that develop scientific tools are currently struggling in using the complex architectures targeted in this Action and especially multicore processors

The development of computational and communication library to efficiently use these complex system is also of key interest for the academic world. For instance it is expected that the communication library developed within this Action will be used in the popular implementation of MPI (such as OpenMPI). As MPI is the leading standard for message-passing programming and is used worldwide, the audience of this result will be very large. More generally, companies and research teams that develop high-performance software (e.g., game, CAD, simulation) have extreme difficulties to use the full power of these systems due to their complexity and hence will benefit from the Action's results.

As stated in section D, efficient scheduling and resource management algorithms are of key importance for enabling performance in large scale-systems. Already existing distributed infrastructures such as the one developed in Europe (EGEE), Germany (D-Grid), Bulgaria (BgGrid) to cite a few, would have better performance with better resource management algorithms implemented in their middleware. Companies are also interested in improving their resources management.

Applications, developed within this Action will first and obviously benefit to the users it has been developed for (such users are neither necessarily member of the Action nor computer scientists researchers). However, since the applications will be developed for a new environment and architecture, the target audience, by side effect, is expected to be widened. Among many possibilities the satellite image processing on heterogeneous systems application could benefit to the ESA-PECS project and to the FP7 SEE-Grid-SCI.
Finally, this Action will have an indirect impact on the society as today performance of desktop computers is mainly increasing by the multiplication of the number of cores. Therefore, the ability to efficiently program such processors will improve everyday usage of everyone machines.

H.2 What?

The results of this Action will be disseminated by the following Activities:
- Research reports and peer-reviewed articles written by members of the Action
- Participation to forum and booth at major event (conferences, symposium, etc.)
- Heteropar and HCW workshops
- Annual Action report
- Working Group reports
- Press statements

H.3 How?

The main way to achieve this dissemination will be made through the Action website. This website will have a private part for internal organization as described in section E.1. The public part of this website will be composed of:
- Description of the Action (goals, means and results)
- All relevant documents will be published (finalized technical report, other report, standardization proposal, etc.).
- Important event announcements (conferences, meetings, etc.)
- News on the life of the Action

- Description on how to join the Action.
- A wiki where members of the Action will be able to collaborate on the edition of a knowledge base.

Moreover two mailing lists will be set: one for internal communication within the Action and one for non-members interested in the results and the activity of the Action.

The annual meeting of the Action will be the perfect occasion for strengthening the dissemination by inviting key actors of the domain (for instance persons from Industry, Non-COST Countries, IEEE or ACM, or connected sciences).

The Action will send representatives of the Action to the major events of the domain of high performance and distributed computing in order to participate in forums, booths, panel discussions, etc.

The Action will contribute to the organization of the Heteropar and HCW workshops, inviting keynote speakers and offering grants to early-stage researchers for participating at this event. Several experts involved in this Action are involved in the editorial boards of well-known journals in the field as well as organizers of scientific events. This fact is a guarantee of the dissemination of the Action activities on a large scale.