

*Open European Network for High-Performance Computing in Complex
Environments*

Proposer:

Name: *Emmanuel Jeannot*

Affiliation: *INRIA*

Address: *615, rue du Jardin Botanique
54602 Villers les Nancy.*

Tel: *+33 3 54 95 85 11*

Fax: *+33 3 83 27 83 19*

E-mail: Emmanuel.Jeannot@loria.fr

DRAFT
MEMORANDUM OF UNDERSTANDING
For the implementation of a European Concerted Research Action designated as

*Open European Network for High-Performance Computing in 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 coordinate European groups working on the use of heterogeneous and hierarchical systems for HPC as well as the development of collaborative activities among the involved research groups.
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 39 million Euro 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

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 goal of the Action is to 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 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 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 environment comes from every levels: 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 the proposed project, 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 proposal want to take on.

Therefore, the opportunity to launch 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 the proposed Action. The main goal of this Action 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: 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 the proposed action. The groups involved in the proposal are both dealing with the theory and applied aspects of heterogeneous computing and therefore their interactions throughout their proposal 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 proposed in this action.

Summarizing, this project within the COST framework allows to expect 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: modeling 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. To develop

such algorithms it has been required to create new models. 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 modeling 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 exists no tools that covers 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, educate one other and finally 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 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: gathering researchers from every architectural layer and propose integrated solution (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 proposal 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 school, workshop, meetings, visit, 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 European countries have already expressed their interest in this network. However, the network is intended to be open to new participants interested in this research area, with special emphasis on the incorporation of partners coming from countries which have recently joined the EU and partners from less-favoured regions.

B.4 Complementarity with other research programmes

It is noticeable that the proposed network is complementary to the other European research effort in this area. First and foremost, the focus of the proposed network is on research coordination while most of the calls within the FP7 are focused on research and infrastructures which are not specific targets in this proposal. 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 Obj 3.4 Computing Systems: 8 projects have started including an NoE and 7 STREPs 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) 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 coordinate European groups working on the use of heterogeneous and hierarchical systems for HPC as well as the development of collaborative activities among the involved research groups.

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 centered on the themes of the different Working Groups. These studies would contain:
 - a. 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.
 - b. 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.
 - c. 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.
 - d. A proposal of the necessary tools for the efficient management of heterogeneous systems: mapping tools, tools for the dynamic analysis of the system, etc.

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.

2. A web site 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 web page, 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 network will be opened to new groups which could collaborate in all the proposed aspects, but especially in the application areas.

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.

C.3 How will the objectives be achieved?

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

Summer schools

Two summer schools will be organized. One could 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 in the project would participate, as members from other groups. The schools would provide young researchers with a good opportunity for them to share information and knowledge, and to present their current research. These schools should contribute to increase the computing community and spread EU knowledge. These schools would be held in the second and third years, as it is important to train researcher early in this Action. The talks and the works presented would be made available on a web page of the project.

International workshops

Heteropar and HCW are the most recognized workshop 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 project usually collaborate to these workshops. The budget of this proposal will be used to sponsor and partially fund these 2 workshops. Key speakers will be invited to participate to these workshops. Such speakers would indicate the main lines in the research in computational aspects of heterogeneous and distributed computing, and would point to future research directions. Between four to eight invited presentations will be organized per year. A summary of the presentations would be published in workshops proceedings.

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.

European working meetings

Four working meetings would be organized, one each year. These meetings would be devoted to studying some specific aspects of heterogeneous computing and will not be specific to a Working Group but of common interest to all of them. The theme to work with would be decided by the members of the group, and opinions would be gathered over the Internet using the project web page, prior to holding the meetings, with the ultimate goal of enhancing participation of external groups and potential new partners to the consortium. After each meeting, the studies would be summarized and placed on the web page.

Management Committee meeting

The first meeting could take place at Strasbourg and the other 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.

Short-terms visit and fellowship

Approximately sixty short-term research visits would be organized. There will be visits by young students to foreign laboratories and departments, and visits of high level researchers to centers with a large number of young researchers. At this point, it is important to emphasize that this Action will be mainly targeted at young researchers, which 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. Twelve long-term visits will also be open to the network participant.

C.4 Benefits of the Action

The overall benefit of the proposed network is to foster collaborations of most European active research group in the field of heterogeneous and hierarchal computing. More precisely, on the scientific side, it is expected that, thanks to this network the Action will:

1. Overcome research fragmentation.
2. Facilitate of 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 outcome of this network 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 proposed network 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 project is to raise the specific community, to focus the research on hot topics and applications of interest for the EU, to propagate the collaboration of research groups with the industry, to stimulate the formation of new groups in new EU countries, and to facilitate the solution of highly computationally demanding scientific problems as mentioned above. For this, the groups involved in this proposal, collaborate with several scientific and industrial groups that could benefit from the advances propitiated by this Action, 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 meeting will be a place for 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. Short-term visits and fellowship will enable bi-lateral research to be conducted on highly focused topics. Moreover, fellowships are viewed as the proper way to ensure strong links between institutions and hence the sustainability of the network.

D.2 Scientific work plan – methods and means

Four Working Groups are proposed to coordinate the scientific researches. 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 detailed, for each Working Group, its objective as well as examples of research issues that will be tackled and how, the research task will be carried out.

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

The goal of this Working Group is to tackle scientific problem that arise in the field 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 researcher that are specialized in this area and have them work together on subtopics. Here follows example of synergies that have already been identified for this Working Group and the appropriate work plan.

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. The proposed 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 anticipate 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 would 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 techniques proposed, 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 proposed research).

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 modeling 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 modeling 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.

2. 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. In this Working Group, partners will work on the design, implementation and test of libraries for an efficient use of modern complex systems. In particular, an emphasis on GPUs (Graphical Processing Unit), multicore processors and heterogeneous network is given.

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 work 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.

Multicore Toolkit

The main features 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.

3. 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 carries-out in the Working Group are given next.

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.

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 project will also target the simultaneous execution of several application workflows that each operates on continuous streams of data sets.

The story gets more complicated when energy minimization is 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.
- 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.

4. 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 of memory. In this Action, the challenge of efficiently executing applications onto complex infrastructures will be taken. A set of applications that members of this Working Group already agree to develop is given here. Due to the nature of the target infrastructures (hierarchical, heterogeneous and large-scale), addressing this challenge requires solving a lot issues. Therefore, this Working Group will use the results of the three above Working Group to design and implement the proposed applications when necessary. It is expected that the proposed application described above, will raise new problems that will be addressed in the other Working Group.

Molecular Dynamics Simulations on large Scale Heterogeneous platforms

Molecular Dynamics simulations require huge amount of computations and standard parallel versions (such as NAMD), based on domain-decomposition techniques, fail to scale well due to the huge number of required synchronization. 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.

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 members of the consortium are already collaborating on the application of heterogeneous computing to this problem and would share their expertise within this program. The main features of the work related with remote sensing data processing that will be carried out in the proposal 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 proposes a coordination programme within the research program of each existing groups and will foster synergies. This way, it will be easier to obtain standards of interest for the global scientific community. Given the nature of the proposed collaboration, COST Research Networking Programme is the most appropriate Action.

An outline of management activities is provided below:

- A Management Committee will be composed of one member of each country (see appendix). The Management Committee will meet 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 this Committee. Another chair will be in charge of the project web pages. Moreover, the MC will select the members in charge of each Working Group.
- The chair of the Management Committee will steer the project day by day, interface the activity requests with the Management Committee as well as document the intermediate and final reports.
- Each year, the MC meeting will establish the activity for the coming year (workshops, schools...). Chairs of each particular activity will be appointed. Each chair will be responsible to promote the activity and to coordinate with other groups.

Concerning research coordination, The Action will work as follows.

- Each group and institution of this Action will register to at least one Working Group. The Working Group membership will be decided on the basis of 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 is focused on one aspect of the research tasks this Action is addressing, a good synergy between the participants will be created.
- A leader will be appointed 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 research interests that fits 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 meeting will be focused on a specific aspect of the Working Group while other meeting will cover the entire scope of the group. In summary, Working Groups will be the corner stones 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 visit activities. Such short-term visit will last several weeks and will be an ideal framework to conduct deep research on a given subject.
- Once a year there will be a working 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 interests. 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

would 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, members of the Action will find:

- Aims and objectives of the project.
- List of researchers involved and their research objectives.
- Public document such as reports, talks and international publications.
- An internal wiki and blog that the network can use to document, comment and interact via.
- Internal documents such as white paper, research report, meeting agendas, meeting reports, drafts, consortium agreement, etc.
- Deliverables as described in section C.1
- Possibility to apply for a short-term visit (with CV and research program)

E.2 Working Groups

The research program presented in section D is decomposed in 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 of 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 the same way. A Working Group leader will be appointed by the Management Committee. Participant of the Action will decide to join a Working Group based on the scientific focus of this Working Group and the research interest of the participant. Working Group meeting will be organized at least once a year. The “*Applications*” Working Group 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 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 of the COST teams 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 exchange of relevant data and methods, since both EGI_DS and the COST actions 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 from the most efficient European HPC service centers providing researchers with access to capability computers and forming the top level of the European HPC ecosystem. The COST action will provide use cases of interest for PACE.
- In the particular field of multicore systems, the COST 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 COST 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 COST action members early in the testing phase.

- In the application field, several project cooperations are 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. No discrimination based on sex, race, religion or citizenship will be considered. 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 Cost Action is planned for four years. A MC meeting is planned once a year. Each Working Group will meet between once and twice a year. The planning of short-term visit 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 and Kick-off meeting.	Elect WG leaders, decide on the upcoming activities, nominate chairs 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	Start of the first short-term visit campaign.	15 visits will be allocated.
M+6	Working 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 it will be taken back for the second year (with the inclusion of the first school, and the start of long-term visits) otherwise another organization will be proposed. 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 chair of the summer school, and decide its location.
M+12	Start of the first long-term visit (fellowship) campaign.	4 fellowships will be granted.
M+14 to M+15	Third Working Group meetings.	
M+6	Start of the second short-term visit campaign.	15 visits will be allocated.
M+18	Working 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: Belgium, Bulgaria, Denmark, France, Germany, Greece, Hungary, Ireland, Italy, Poland, Portugal, Romania, Russia, Slovenia, Spain, United Kingdom.

On the basis of national estimates, the economic dimension of the activities to be carried out under the Action has been estimated at 39 Million € for the total duration of the Action (4 years, an average of 11.2 participants per country).

Estimations are done as follows:

	No.	Total PMs	Monthly €	Total M€
Members of Management committee	32	800	10000	8
Members of the Working Groups	48	1800	7500	13.5
Early-stage researchers	100	4800	2500	12
Travels	150			0.5
Scientific event organization	8			0.4
Equipments	16			3
Overheads	16			1.6
Total				39

Total person years/country: 11.2

(2 for management committee, 3 for working group, 6.2 for early-stage researchers)

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.

H. DISSEMINATION PLAN

H.1 Who?

This Action aims 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. Therefore, companies such as NAG (Numerical Algorithms Group) in the UK, the ESI group or the Scilab consortium in France will benefit from the solutions proposed in this Working Group. The audience is also composed of smaller scale companies. For instance AURORA Software and Testing, in Valencia could use the algorithms and methods developed in the Action to their problems of microwave application.

The development of computational and communication library for efficiently use these complex system is also of key interests for the academic world. For instance it is expected that the communication library developed within this Action will be implemented in 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 broad. More generally, companies and research teams that develop high-performance software (*e.g.*, game, CAD, simulation) have extreme difficulty to use the full power of these systems due to their complexity and hence will benefit from the Action results.

As stated in section D, efficient scheduling and resource management algorithms are of key importance for enabling performance in large scale-system. Already existing distributed infrastructure 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. This is the case of IBM-France, which wants to optimize energy consumption in their green-data center while maintaining a high level of quality of service.

Applications, developed within this Action, will first and obviously benefit to the users it has been developed for (such users are not 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 or that the computer simulation application could be used by the Turboinstitut of Slovenia.

Lastly, 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, being able 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 where the community is already actively participating.
- Annual Action report
- Working-group report
- Press official 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 document will be published (finalized technical report, other report, standardization proposal, etc.).
- Important event announcement (conference, meeting, 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 person from industry, non-EU country, IEEE or ACM, or connected sciences).

The Action will send representatives of the Action to the major event of the domain of high performance and distributed computing for participation to forums, booths, panel discussions, etc.

The Action will participate in the organization of the Heteropar and HCW workshops and will invite keynote and offer grants to early-stage researchers for participating at this event.

Several experts involved in this network 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 action activities on a large scale.

Part II – Additional Information

A. LIST OF EXPERTS

Here follows a list of experts who have been contacted during the drafting of the proposal and have expressed their interest in participating. The list is sorted by country, for each country the first expert will be member of the Management Committee.

Belgium

MC member ;Prof. Pierre Manneback,
Dept of Information Technology
Faculty of Engineering
9 rue de Houdain
7000 Mons, Belgium

Email: Pierre.Manneback@fpms.ac.be

Tel: +32 65 374050

Bulgaria

MC member : Prof. Svetozar Dimitrov Margenov
Scientific Secretary of IPP-BAS
Institute for Parallel Processing (IPP-BAS)
Bulgarian Academy of Sciences
"Acad. G. Bontchev" Str., Bl. 25 A
1113 Sofia, Bulgaria

Email : margenov@parallel.bas.bg

Tel: +359-2 979 6612

Denmark

MC member: Prof. Jerzy Wasniewski
Senior Research Professor
Department of Informatics and Mathematical Modeling
Technical University of Denmark
Richard Petersens Plads - Building 321
DK-2800 Kongens Lyngby - Denmark

Email: jw@imm.dtu.dk

Tel: +45 45 25 30 76

France

MC member: Dr. Emmanuel Jeannot
INRIA
615, rue du Jardin Botanique
54602 Villers les Nancy, France

Email: Emmanuel.Jeannot@loria.fr

Tel: +33 3 54 95 85 11

Other experts:

Prof. Olivier Beaumont, University of Bordeaux, Olivier.Beaumont@labri.fr

Prof. Serge Petiton, University of Lille, Serge.Petiton@lifl.fr

Prof. Yves Robert, Ecole Normale Supérieure of Lyon, Yves.Robert@ens-lyon.fr

Germany

MC member: Prof. Gudula Rünger
Chemnitz University of Technology
Department of Computer Science
09107 Chemnitz, Germany

Email: ruenger@informatik.tu-chemnitz.de

Tel: 0371 531 31794

Other experts:

Prof. Thomas Rauber, University Bayreuth, Rauber@uni-bayreuth.de
Dr. Ramin Yahyapour, University Dortmund, ramin.yahyapour@udo.edu

Greece

MC member: Dr. Helen Karatza
Associate Professor
Department of Informatics
Aristotle University of Thessaloniki
54124 Thessaloniki, Greece

Email: karatza@csd.auth.gr

Tel: +30 2310 997974

Hungary

MC member : Prof. Peter Kacsuk
MTA SZTAKI
Laboratory of Parallel and Distributed Systems
1111 Budapest, Kende u. 13-17.

Email: kacsuk@sztaki.hu

Tel: +36 1 329 7864

Ireland

MC member: Dr. Alexey Lastovetsky
School of Computer Science and Informatics
University College Dublin
Belfield, Dublin 4, Ireland

Email: alexey.lastovetsky@ucd.ie

Tel: +353 1 716 2916

Italy

MC member: Dr. Andrea Clematis
Istituto di Matematica Applicata e Tecnologie Informatiche (IMATI) – CNR
Via De Marini 6
16149 Genova , Italy

Email: clematis@ge.imati.cnr.it

Tel: +39-010-6475668

Other expert:

Dr. Ranieri Baraglia, Information Science and Technologies Institute, Pisa, ranieri.baraglia@isti.cnr.it

Poland

MC member: Dr. Maria Ganzha
Systems Research Institute Polish Academy of Sciences
ul Newelska 6
Warsaw, 01-447

Email: maria.ganzha@ibspan.waw.pl

Tel: +48 501754269

Other experts:

Dr. Marcin Paprzycki, Systems Research Institute, Polish Academy of Science,
paprzyck@ibspan.waw.pl

Prof. Przemysław Stpiczynski, Maria Curie Sklodowska University, przem@hektor.umcs.lublin.pl

Prof. Roman Wyrzykowski, Czestochowa University of Technology, roman@icis.pcz.pl

Portugal

MC member: Prof. Leonel Augusto Pires Seabra de Sousa
INESC-ID,
R. Alves Redol, 9, 1000-029,
Lisboa, PORTUGAL

Email: las@inesc-id.pt

Tel: +351213100300

Other expert:

Prof. Jorge Barbosa, FEUP/INEB, University of Porto, Portugal jbarbosa@fe.up.pt

Romania

MC member: Prof. Dana Petcu,
Computer Science Department,
Faculty of Mathematics & Informatics,
Western University of Timisoara,
and Institute e-Austria in Timisoara
B-dul V.Parvan 4, 300223 Timisoara, Romania

Email: petcu@info.uvt.ro

Tel : +40-256-592370

Russian Federation

MC member: Dr. Alexey Kalinov
119017 Russia
Moscow, B. Ordynka Str., 44, building 4

Email: akalinov@cadence.com

Tel: +7 495 980 8260

Slovenia

MC member: Dr. Roman Trobec,
Institut Jozef Stefan
Department of communication systems
Jamova 39 1000
Slovenia

Email: roman.trobec@ijs.si

Tel: +386 1 477 3497

Spain

MC member: Prof. Domingo Gimenez
Departamento de Informatica y Sistemas,
Universidad de Murcia
Campus de Espinardo, 30070 Murcia, Spain

Email: domingo@dif.um.es

Tel: +34 968 364611

Other experts:

Prof. Antonio J. Plaza, University of Extremadura, aplaza@unex.es

Prof. Enrique S. Quintana-Orti, Universidad Jaime I, quintana@icc.uji.es

Prof. Antonio M. Vidal, Polytechnic University of Valencia, avidal@dsic.upv.es

Prof. Francisco Almeida Rodriguez, La Laguna University, falmeida@ull.es

The United Kingdom

MC Member: Prof. Mark Baker
School of Systems Engineering,
The University of Reading,
Whiteknights,
Reading,
Berkshire,
RG6 6AY, UK

Email: mark.baker@computer.org

Tel: +44 118 378 8615

Other experts:

Prof. Rizos Sakellariou University of Manchester, rizos@cs.man.ac.uk

Prof. Jack Dongarra, University of Manchester, dongarra@eecs.utk.edu

Prof. Vladimir Getov, University of Westminster, V.S.Getov@westminster.ac.uk

Here follows a list of experts who may well be interested but who have not been contacted, or who

have not yet replied, during the pre-proposal planning. This list is not intended to be exhaustive but, with the above list, shows the large spectrum of potential interested researchers and institutions.

- Prof. Erik Elmroth, Umea University, Sweden, elmroth@cs.umu.se
- Prof. Peter Sloot, University of Amsterdam, The Netherlands, P.M.A.Sloot@uva.nl
- Prof. Henk Sips, University of Delft, The Netherlands, H.J.Sips@ewi.tudelft.nl
- Prof. Raymond Namyst, University of Bordeaux, France, Raymond.Namyst@labri.fr
- Prof. Marios Dikaiakos, the University of Cyprus, Cyprus, mdd@ucy.ac.cy
- Prof. David Walker, Cardiff University, United Kingdom, David.W.Walker@cs.cardiff.ac.uk
- Prof. Rosa M. Badia, Universitat Politecnica Catalunya, Barcelona, Spain, rosab@ac.upc.edu
- Prof. Tomas Margaleff Universidad Autonoma de Barcelona, Spain tomas.margalef@uab.es

B. HISTORY OF THE PROPOSAL

The core research groups involved in the proposal began their work on high performance parallel computing on non-traditional heterogeneous architectures in mid and late 90s. The first collaboration was established in 2000 between the groups led by Alexey Kalinov (Russia), Alexey Lastovetsky (Ireland) and Yves Robert (France). The collaboration involved visit exchanges and discussions of the research results.

In 2001, the first international workshop on algorithms, models and tools for parallel computing on heterogeneous platforms (HeteroPar'01) was organized by Kalinov, Lastovetsky and Robert in Las Vegas, USA. In parallel, collaborative research between a number of French research groups and the Innovative Computing Laboratory led by Jack Dongarra was conducted. The leaders of the core research groups were also involved in the organization of the IPDPS Heterogeneous Computing Workshop since late 90s and early 00s. In 2003, the second HeteroPar workshop (HeteroPar'03) was organized by Alexey Kalinov, Alexey Lastovetsky and Denis Trystram in conjunction with PPAM 2003. By that time, more research groups have gotten interested in heterogeneous parallel computing and this conference was quite fruitful in building up the research community in this field. In particular, contacts were established with the groups led by Domingo Gimenez (Spain), Roman Wyrzykowski (Poland), Marcin Paprzycki (USA/Poland), and some others. This process continued over the following years with the HeteroPar workshop series being particularly instrumental in making new contacts and starting new collaborations. The established collaborations were later supported by national grants such as the collaborative research work between the UCD group (Ireland) and groups from the UT Innovative Computing Lab (USA, since 2004), the University of Murcia (Spain, since 2004), the Technical University of Valencia (Spain, since 2006), the University of Extremadura (Spain, since 2005). Other clusters of collaboration were built in parallel.

The very first attempt to build a more structured network of research groups in this area and get support from the EU to implement this idea was made in 2004. The attempt coordinated by Alexey Lastovetsky involved 10 research groups from 6 countries (Ireland, France, United States, Spain, Russian Federation, and Portugal) and the application was submitted to the FP6 program (in response to the call FP6-2004-IST-FETPI as an Integrated Project in the Global Computing activity). While the project got quite decent marks from the reviewers, it was not supported. Submission of a pre-proposal coordinated by Domingo Gimenez was made in response to the COST Open Call 2006. That time, the proposal involved 15 research groups from 10 countries (Spain, Ireland, Romania, United States, Portugal, Poland, France, Italy, India, and Russian Federation). While the pre-proposal received quite high marks and passed the pre-set eligibility threshold for being considered for a full proposal, due to the unusually high number of proposals that year, it was not funded (only 6 pre-proposals out of a total of 100 assessed ones were selected for further consideration). The third attempt coordinated by Emmanuel Jeannot, Domingo Gimenez and Alexey Lastovetsky was made in 2006. A proposal of a European Network for High Performance Computing on Heterogeneous and Hierarchical Environments was submitted to the European Science Foundation (ESF) and involved almost 30 research groups from 14 countries (Belgium, Denmark, France, Germany, Greece, Hungary, Ireland, Italy, Poland, Portugal, Romania, Slovenia, Spain, and the United Kingdom). The

proposal received very high marks from the reviewers. However, it was not funded due to very severe competition in terms of quality and a wide range of research domains that should be evenly covered by successful proposals.

C. PRELIMINARY WORK PROGRAMME

Based on the above MoU, and the fact that most expected MC members already know each other, the work program will be as follow.

1. Some synergies have already been identified and expressed in section D.2
2. Each participant will register to at least a Working Group by giving its main research focus.
3. More synergies will be identified during the first WG meeting and the subsequent one, based on the presentations during these meeting and the expressed research interests during WG registration (step 2.)
4. Once synergies will be identified, short-term visits will enable research to be carried-on on that topics
5. After the first year, fellowship will enable deeper collaborations on those topics.
6. Results will be presented to the Working Group at the next meeting and to the whole network at the European working meeting.

Steps 3 to 6 will be repeated during the whole duration of the Action.

D. RECENT PUBLICATIONS

A selection of recent publications of the proposer (in the topic this Action):

Journals:

- Emmanuel Jeannot, Keith Seymour, Asym Yarkhan, and Jack J. Dongarra, **Improved Runtime and Transfer Time Prediction Mechanisms in a Network Enabled Servers Middleware**, *Parallel Processing Letters*, 17(1): 47–59, March 2007.
- Raphaël Bolze, Franck Cappello, Eddy Caron, Michel Daydé, Frédéric Desprez, Emmanuel Jeannot, Yvon Jégou, Stéphane Lanteri, Julien Leduc, Noredine Melab, Guillaume Mornet, Raymond Namyst, Pascale Primet, Benjamin Quetier, Olivier Richard, El-Ghazali Talbi, and Iréa Touche, **Grid'5000: A Large Scale And Highly Reconfigurable Experimental Grid Testbed**, *International Journal of High Performance Computing Applications* 20(4) 481-494, november 2006.
- Emmanuel Jeannot and Frédéric Wagner, **Scheduling Messages for Data Redistribution: an Experimental Study**, *International Journal of High Performance Computing Applications* 20(4) 443-454, november 2006.
- Johanne Cohen, Emmanuel Jeannot, Nicolas Padoy, and Frédéric Wagner, **Message Scheduling for Parallel Data Redistribution between Clusters**, *IEEE Transactions on Parallel and Distributed Systems*, 17(10): 1163-1175, 2006.
- Vandy Bertin, Joel Goossens and Emmanuel Jeannot, **On the Distribution of Sequential Jobs in Random Brokering for Heterogeneous Computational Grids**, *IEEE Transactions on Parallel and Distributed Systems*, 17(2):113–124, 2006.
- Yves Caniou and Emmanuel Jeannot, **Multi-Criteria Scheduling Heuristics for GridRPC Systems**, *International Journal of High Performance Computing Applications*, 20(1):61–76, spring 2006.
- Eddy Caron, Bruno DelFabbro, Frédéric Desprez, Emmanuel Jeannot and Jean-Marc Nicod, **Managing Data Persistence in Network Enabled Servers**, *Scientific Programming Journal* 13(4):333–354, 2005.
- Michel Cosnard, Emmanuel Jeannot and Tao Yang, **Compact DAG Representation and its Symbolic Scheduling**, *Journal of Parallel and Distributed Computing* Vol. 64, No. 8, pages 921-935. Aug. 2004

Conferences:

- Emmanuel Jeannot, Erik Saule, and Denis Trystram. **Bi-Objective Approximation Scheme**

- for Makespan and Reliability Optimization on Uniform Parallel Machines.** In The 14th International Euro-Par Conference on Parallel and Distributed Computing (Euro-Par 2008), Las Palmas de Gran Canaria, Spain, August 2008. To appear.
- Louis-Claude Canon and Emmanuel Jeannot. **Scheduling Strategies for the Bicriteria Optimization of the Robustness and Makespan.** In 11th International Workshop on Nature Inspired Distributed Computing (NIDISC 2008), Miami, Florida, USA, April 2008.
 - Emmanuel Jeannot and Luiz-Angelo Steffanel, **Fast and Efficient Total Exchange on Two Clusters,** In the 13th International Euro-Par Conference European Conference on Parallel and Distributed Computing (Euro-Par 2007), Rennes, France, aug. 2007.
 - Jack J. Dongarra, Emmanuel Jeannot, Erik Saule, and Zhiao Shi, **Bi-objective Scheduling Algorithms for Optimizing Makespan and Reliability on Heterogeneous Systems,** In 19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'07), San Diego, CA, USA, june 2007.
 - Z. Shi, E. Jeannot, and J. J. Dongarra, **Robust Task Scheduling in Non-Deterministic Heterogeneous Systems,** IEEE International Conference on Cluster Computing, Barcelona, Spain, sept. 2006.
 - Emmanuel Jeannot and Frédéric Wagner, **Messages Scheduling for data Redistribution between Heterogeneous Clusters,** IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2005). Phoenix, AZ, USA. Best paper award in algorithm.
 - Emmanuel Jeannot, **Improving Middleware Performance with AdOC: an Adaptive Online Compression Library for Data Transfer,** IEEE IPDPS'05, Denver, CO, USA.
 - Yves Caniou and Emmanuel Jeannot, **Experimental Study of Multi-Criteria Scheduling Heuristics for GridRPC Systems,** in ACM-IFIP Euro-Par 2004, Pisa, Italy, sept 2004.
 - Yves Caniou and Emmanuel Jeannot, **Efficient Scheduling Heuristics for GridRPC Systems,** in QOS and Dynamic System workshop of IEEE ICPADS (International Conference on Parallel and Distributed Systems) conference, New-Port Beach, California, july 2004. Pages 621-630.
 - Emmanuel Jeannot and Frédéric Wagner, **Two Fast and Efficient Message Scheduling Algorithms for Data Redistribution through a Backbone,** in IEEE International Conference on Parallel and Distributed Processing Symposium (IPDPS) 2004. Santa-Fe, New-Mexico.
 - Bertrand Cirou and Emmanuel Jeannot, **Triplet : a Clustering Scheduling Algorithm for Heterogeneous Systems,** in IEEE ICPP International Workshop on Metacomputing Systems and Applications (MSA'2001), sept. 2001, Valencia, Spain.
 - Michel Cosnard, Emmanuel Jeannot and Laurence Rougeot, **Low Memory Cost Dynamic Scheduling of Large Coarse Grain Task Graphs,** 1998 IEEE International Parallel Processing Symposium (IPPS'98), Orlando, Florida.

Other major publications related to topics of this Action

- Dongarra J. and Lastovetsky A., **An Overview of Heterogeneous High Performance and Grid Computing",** In Engineering the Grid: Status and Perspective, Eds B.DiMartino, J.Dongarra, A.Hoisie, L.Yang, and H.Zima, pp.1-25, American Scientific Publishers, February 2006
- Kumar, R., Tullsen, D.M., Jouppi, N.P., Ranganathan, P. **Heterogeneous chip multiprocessors,** Computer Nov. 2005 Volume: 38, Issue: 11 On page(s): 32- 38
- H. Sutter, **The Free Lunch Is Over --A Fundamental Turn Toward Concurrency in Software,** Dr.Dobb's Journal, 30(3), 2005.
- A. Adl-Tabatabai, C. Kozyrakis, and B. Saha: **Unlocking concurrency,** ACM Queue 4(10):24-33, Dec 2006.
- D. Kuck: Platform 2015 **Software-Enabling Innovation in Parallelism for the next Decade,** Intel White Paper, Technology@Intel Magazine, 2005.
- John D. Owens, David Luebke, Naga Govindaraju, Mark Harris, Jens Krüger, Aaron E. Lefohn,

- and Tim Purcell. **A Survey of General-Purpose Computation on Graphics Hardware**. Computer Graphics Forum, volume 26, number 1, 2007, pp. 80-113.
- Akhter, Shameem (2006). **Multicore Programming**. Richard Bowles (Intel Press), pp. 11-13.
 - P.J.C. Aerts, K. Koski, F. Lozano, **A sustainable High Performance Computing Ecosystem for Europe**, HET Final report, January 2007, <http://www.hpcineuropetaskforce.eu/deliverables>
 - L. Johnsson, **European Large-Scale HPC Initiatives: Background and Proposed Actions SNIC study**, May 2007, http://www.snic.vr.se/docs/European_HPC_final.pdf
 - PACE Consortium, **Memorandum of Understanding concerning the establishment of a European Tier-0**, High-Performance Computing Service, April 2007, <http://www.prace-project.eu/documents>
 - NSF - **Cyberinfrastructure Council, Cyberinfrastructure: Vision for the 21st Century Discovery**, NSF 07-28, March 2007 <http://www.nsf.gov/pubs/2007/nsf0728/>
 - Zhe Fan, Feng Qiu, Arie Kaufman, Suzanne Yoakum-Stover, **GPU Cluster for High Performance Computing**, Proceedings of the ACM / IEEE Supercomputing Conference 2004, November 06-12, 2004, Pittsburgh, PA.
 - Wahid Nasri, Hajer Hamad, Hadhemi Fejjari. **A Framework for Adaptive Communication Modeling on Heterogeneous Hierarchical Clusters**. Proceedings of the 2006 IEEE International Conference on Cluster Computing, 25-28 Sept. 2006, Barcelona, Spain.

E. FURTHER REMARKS

While the Action is focused on attracting young researchers and female researchers in particular, it is important to notice that one of the important issues that is detrimental to attracting female researchers to computational sciences is the fact that most "positions of power" are occupied by white Anglo-Saxon males. In this way there are no role-models that would assure female researchers that there is a possibility to reach higher levels of career advancement.

This proposal naturally counteracts the existing stereotypes by putting four women in the Management Committee of our network. These are: Maria Ganzha, Helen Karatza, Dana Petcu and Gudula Runger. These senior researchers will be able to act as role-models for young female computational scientists.

This Action will also take a particular attention to early-stage researcher by allocating in priority funds for their participation of major events such as HCW and Heteropar. Moreover, Working Group leaders and other responsibilities (*e.g.* meeting chair, summer school organization) will be devoted in priority to young researchers taking a special care of the gender-balance.

This proposal aims at gathering the leading European countries in the area of heterogeneous and hierarchical high-performance computing. As already 16 European countries are involved, it is clear that this proposal reflects a wide European dimension. Moreover, as this network is open to new membership, and none all of the possible institutions have been involved at this drafting stage, it is expected that the number of participants will increase after the launch of this Action.

The monitoring of the activities and assessment of this Cost Action will be the responsibility of the MC. Several documents will be written down and evaluated at each MC meeting:

- The Annual Report will gather and summarize all the activities of this Action.
- Researchers, for every new result found within this Action, will write a research report. It is expected that each participant will write at least one research report with another participant of the Action.
- Proceedings, gathering talks of each summer school, will be edited by the Action
- A report for each short-term visit or fellowship, presenting the activities carried-out during the visit will be written. It is expected that a research report will be written for each visit and

that a concrete outcome will be an international publication in high-level conference (for short term-visit) or journal (for fellowship).

- Each Working Group will be responsible to report its activities annually. Such activities will be monitored and assessed at MC meetings.
- Proceedings of the European working meeting will be edited by the Action.

The website will reflect the main activities of this Action and will be used to assess and monitor the Action. Moreover, the website will be used as a repository for all the above documents. To this regard, it will be the main medium for diffusing this Action results.