Brussels, 24 March 2020

COST 038/20

DECISION

Subject: **Memorandum of Understanding for the implementation of the COST Action "Connecting Education and Research Communities for an Innovative Resource Aware Society" (CERCIRAS) CA19135**

The COST Member Countries and/or the COST Cooperating State will find attached the Memorandum of Understanding for the COST Action Connecting Education and Research Communities for an Innovative Resource Aware Society approved by the Committee of Senior Officials through written procedure on 24 March 2020.

Funded by the Horizon 2020 Framework Programme
of the European Union

MEMORANDUM OF UNDERSTANDING

For the implementation of a COST Action designated as

COST Action CA19135
CONNECTING EDUCATION AND RESEARCH COMMUNITIES FOR AN INNOVATIVE RESOURCE AWARE SOCIETY (CERCIRAS)

The COST Member Countries and/or the COST Cooperating State, accepting the present Memorandum of Understanding (MoU) wish to undertake joint activities of mutual interest and declare their common intention to participate in the COST Action (the Action), referred to above and described in the Technical Annex of this MoU.

The Action will be carried out in accordance with the set of COST Implementation Rules approved by the Committee of Senior Officials (CSO), or any new document amending or replacing them:

a. "Rules for Participation in and Implementation of COST Activities" (COST 132/14 REV2);
b. "COST Action Proposal Submission, Evaluation, Selection and Approval" (COST 133/14 REV);
c. "COST Action Management, Monitoring and Final Assessment" (COST 134/14 REV2);
d. "COST International Cooperation and Specific Organisations Participation" (COST 135/14 REV).

The main aim and objective of the Action is to create a cross-domain, cross-stack, community of researchers and practitioners who will join forces around the common theme of resource-aware computing. The activities of the Action will lead to long-term capacity-building scientific, technical and economical advances in resource-aware computing for multicore/many-core systems. This will be achieved through the specific objectives detailed in the Technical Annex.

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 104 million in 2019.

The MoU will enter into force once at least seven (7) COST Member Countries and/or COST Cooperating State have accepted it, and the corresponding Management Committee Members have been appointed, as described in the CSO Decision COST 134/14 REV2.

The COST Action will start from the date of the first Management Committee meeting and shall be implemented for a period of four (4) years, unless an extension is approved by the CSO following the procedure described in the CSO Decision COST 134/14 REV2.

TECHNICAL ANNEX

OVERVIEW

Summary

Parallel computing platforms have revolutionised the hardware landscape by providing high-performance, low-energy, and specialized (viz. heterogeneous) processing capabilities to a variety of application domains, including mobile, embedded, data-centre and high-performance computing. However, to leverage their potential, system designers must strike a difficult balance in the apportionment of resources to the application components, striving to avoid under- or over-provisions against worst-case utilisation profiles. The entanglement of hardware components in the emerging platforms and the complex behaviour of parallel applications raise conflicting resource requirements, more so in smart, (self-)adaptive and autonomous systems. This scenario presents the hard challenge of understanding and controlling, statically and dynamically, the trade-offs in the usage of system resources, (time, space, energy, and data), also from the perspective of the development and maintenance efforts.

Making resource-usage trade-offs at specification, design, implementation, and run time requires profound awareness of the local and global impact caused by parallel threads of applications on individual resources. Such awareness is crucial for academic researchers and industrial practitioners across all European and COST member countries, and, therefore, a strategic priority. Reaching this goal requires acting at two levels: (1) networking otherwise fragmented research efforts towards more holistic views of the problem and the solution; (2) leveraging appropriate educational and technology assets to improve the understanding and management of resources by the academia and industry of underperforming economies, in order to promote cooperation inside Europe and achieve economical and societal benefits.

Areas of Expertise Relevant for the Action	Keywords
● Electrical engineering, electronic engineering, Information engineering: Embedded systems, cyber-physical systems ● Electrical engineering, electronic engineering, Information engineering: Computer systems, parallel/distributed systems ● Computer and Information Sciences: Theoretical aspects of pervasive and ubiquitous computing ● Economics and business: Management of Technology and Innovation ● Other engineering and technologies: Applied mathematics, statistics, non-computational modeling for other engineering and technology	● Predictable, safe and reliable computing ● Resource-aware computing ● Information and program analysis ● Knowledge transfer ● Technology transfer

Specific Objectives

To achieve the main objective described in this MoU, the following specific objectives shall be accomplished:

Research Coordination

● Resource analysis and associated costs: the goal is to initiate and coordinate research efforts around a new family of tools, approaches, and methods for analysing the resource demands and costs of applications and systems to raise awareness of current and future demands.

● Models for resource use and requirements: To reason about resource requirements and the trade-offs involved, CERCIRAS aims to coordinate efforts focusing on the construction and validation of cost models and optimizations.

● Trade-offs: Key enablers for resource-aware computing will be the identification, definition, and development of general techniques for trading resources against each other to obtain Pareto-optimal systems, not only in terms of required computing resources, but also considering socio-economic aspects, such as development time and market readiness.

● Use cases: The research efforts that will be enabled by this Action via academia-industry collaborations should allow for fast adoption in several application areas and for a wide-variety of resources, including the Internet-of-Things, home automation, cyber-physical systems, robotics, autonomous vehicles, data centres and cloud computing and the resources mentioned above.

Capacity Building

● Balance and raise the awareness of systems resource usage among researchers and practitioners of computing systems, by promoting and facilitating the integration of aspects of resource usage directly into Computer Science curricula, starting from the undergraduate level.

● Balance and increase the number of trained researchers and PhD students with training in holistic resource-usage analysis, modelling, verification, and optimization techniques for computing systems, by organizing workshops, summer school and other networking activities on resource-aware computing.

● Expand knowledge of resource usage techniques for parallel computing systems among all the stakeholders, students, researchers, and practitioners, by organizing crash courses, training materials, and mutual research visits, to kick-start joint interdisciplinary efforts and projects that will tackle pressing issues of resource awareness.

● Balance and raise awareness in the multi-/many-core hardware community on predictability barriers to mitigate the large discrepancies between resource usage and system predictability, through the active involvement of the hardware expert members of the Action.

● Motivate new collaborative national and international research projects to work on holistic resource modeling and analysis for parallel computing systems, by cross-networking and liaising with public and private research funding agencies, insisting on the strategic importance of the topic.

● Provide timely and sustained knowledge transfer between academia and industry, ensuring that academic research on resource modelling and analysis addresses industry needs, in constant consultation with the industrial partners of the Action and a network of external contact points.

TECHNICAL ANNEX

1 S&T EXCELLENCE

1.1 SOUNDNESS OF THE CHALLENGE

1.1.1 DESCRIPTION OF THE STATE-OF-THE-ART

The wealth and growth of the European and COST member state markets are critically dependent on a profound and, ideally, unified understanding of the potential and challenges of digital technologies [1]. Accelerating the acquisition of such understanding requires us to address the thematic fragmentation that afflicts research communities, and to active seek a more holistic view of both, the problems at hand and of their possible solutions. At the same time, overcoming regional differences in economic and technological conditions requires amplification of knowledge and technology transfer beyond the confines of established academic research.

Research: The advent of multi-CPU processors has revolutionised the computing landscape at all levels, from consumer to special-purpose platforms. Indicators [2] suggest there will be increased acceleration toward computing systems comprising hundreds of processor cores. These may well be heterogeneous, with specialized accelerators for machine learning, AI, cybersecurity etc, all operating in parallel. Understanding, predicting and optimising resource usage in such systems is paramount to ensuring their successful deployment in an increasingly large variety of application scenarios, from level-4+ autonomous driving to advanced real-time medical diagnostics, from mobile embedded computing and cyber-physical systems to data-centres and high-performance computing (HPC).

What is common to all these systems is that they aim to provide high performance (as required for e.g. self-guided and self-adaptive execution), at low energy cost (as required for e.g. extended autonomy). Achieving this goal requires a profound understanding of the usage profile of system resources, including time, energy, the underlying hardware, data placement and transmission, to enable sound trade-offs between them, possibly at runtime and in an unattended manner. While there exists some work on predicting time or energy usage for multi-cores [3-8], and on good data placement in parallel computing, much of this research is still initial, domain-specific, and based on simplifying assumptions.

Execution time/wall-clock time, as a resource, can be treated either at a system level, via schedulability analysis, or at an executable level, via worst-case timing analysis. Despite difficulties in obtaining tight bounds for complex architectures, worst-case timing analysis for single-core processors is now fully mature and is used industrially. For multi-/many-cores, however, research results are still immature [9-11]. Acquiring accurate results within a reasonable time and effort is still a challenge: exhaustive exploration of all possible application states is too complex computationally, and is often infeasible with typical measurement-based approaches. A recent, but promising, approach is to use machine learning for time-estimation [12]. If successful, this would allow predictions to be made early in the design cycle, thereby facilitating effective system dimensioning. The situation for schedulability analysis is similar [3]. One of the main underlying challenges, intrinsic to multi-cores, is that complex contention scenarios arise, at various levels of the execution stack. These effects are hard to bound tightly without excessive pessimism. Addressing this is crucial to performant safety-critical systems.

Energy is also being extensively researched as a critical resource. However, methods for energy/power optimisation are also still in their infancy [4-6]. Many system designs are ad-hoc, and thus prevent general solutions. Energy aspects are usually only considered indirectly, since there is no effective means either to express energy properties systematically and formally, or to provide early validation of energy predictions. Research on trading time vs energy at macroscopic scales is mostly limited to dynamic voltage/frequency scaling (DVFS) and gating (disabling unused resources) [13-14]. Moreover, bounding energy consumption is more difficult than bounding execution time: energy consumption depends on data-related switching activity at the hardware level as well as environmental conditions (e.g., operation temperature). Current static analysis techniques and energy models both fail to capture these effects, or to provide sufficiently tight energy consumption bounds [15].

The choice of hardware platform also affects resource-optimisation options and trade-offs, especially for the increasingly complex platforms that modern markets require. Hardware developers constantly seek to produce platforms with improved capabilities, including multiple options to perform the same computational task with different resource trade-offs. However, determining the best resource utilisation for a given heterogeneous platform becomes increasingly difficult as these capabilities evolve. For most systems, resource control solutions are low-level, very close to the hardware, and require detailed, and often proprietary, knowledge of this hardware. Even though these challenges are beginning to be addressed from a programmability perspective [16, 17], dynamic adaptation based on worst-case execution times and energy bounds remains largely unknown.

Data is fundamental to any computation. Complex applications rely on good data management, especially regarding cost of access [18]. Data availability is often an important prerequisite to overall performance, increasingly so for big-data applications [19]. Data quality (for example, freshness) is another essential property that needs active management. To date, the main focus has been on data containers, data communication channels and their properties. Thus, memory access times, including different cache levels [7] and communication channels [20] have been studied extensively. Similarly, collection and maintenance of high-quality, labelled data (e.g., for training) is well understood. However, for live data sources, trade-offs involving data quality vs. resources required (e.g., network bandwidth, energy) are not as well understood. Enabling these trade-offs could enable more powerful analysis frameworks and more informed trade-offs (e.g., integrity/ freshness, vs. time/energy).

Since market pressures cause most systems to be constructed under pressure (human, financial, time-to-delivery etc.), development effort (time, tools, expertise, etc.) is an important resource in its own right. However, the literature presently treats development-time resource optimization and the development process itself independently [21]. This prevents the exploitation of key trade-offs such as energy-efficiency vs. time-to-market or construction costs vs. system reliability. Resource optimization, excluding development resources, is presently the task of skilled engineers with a deep knowledge of the whole systems stack (software, algorithms and hardware). Such expertise is costly and rare. Even when it is available, resource optimisation is time-consuming, which favours inefficient designs.

Technology and knowledge transfer: The traditional per-domain segregation of research communities, combined with regional imbalances in the economy and education, impair the transfer of technology and knowledge from research to practice. The disconnect between research communities prevents the exploitation of potential synergies. An example of such a loss results from the quest for predictable execution techniques (driven by safety concerns), which is progressing independently and largely unaware of the quest for similar solutions for high-performance computing (driven by stability concerns), e.g., [7, 22, 23] vs. [24-26]. Connecting neighbouring ICT communities, such as these, will enable results achieved in one to be exploited elsewhere. For example, by drawing from techniques devised by the timing analysis community, common techniques could be developed that could also analyse energy usage. Conversely, detailed micro-architectural analysis and novel measurement/statistical techniques used in energy analysis could also benefit timing analysis. Currently, a lack of communication inhibits holistic system views and/or the emergence of cross-cutting techniques such as coherent resource use and prediction strategies. Today, most resource analyses and prediction strategies are tailored towards individual resources. A “whole system” view, however, opens the possibility of exploiting a wider range of more efficient prediction methods (e.g., machine learning) in addition to the transfer of analysis methods between communities.

The diversity across Europe is large enough to create disconnection and fragmentation of domain knowledge, which prevents its fluid circulation [27, 28]. Even in regions that are traditionally strong in fundamental science, the lack of relevant locally-based industry often hinders the awareness of technology transfer opportunities and prevents the emergence of a “technology transfer culture” in education. Most technology transfer remains concentrated in areas that host industrial leaders. Even when leading industrial concerns establish themselves in neglected regions or activate remote collaborations, knowledge and technology transfer frequently lags behind due to the sheer absence or insufficiency of education and professional development. A holistic rebalancing approach across the concerned research and technology sectors, domains and regions would remedy this situation. This requires intense exchange among educational, scientific and industrial actors, especially in knowledge transfer, to achieve sustainable development and overcome regional and domain separation.

1.1.2 DESCRIPTION OF THE CHALLENGE (MAIN AIM)

The CERCIRAS Action (Connecting Education and Research Communities for an Innovative Resource Aware Society), aims to make a major contribution to all these issues. On the research front it aims to promote the development of new techniques and technologies that will enable us to predict and optimise the usage of resources and understand their trade-offs. The overall goal is to ensure the effective and efficient execution of software applications on state-of-the-art and future computing systems. A major challenge lies in understanding and managing the trade-offs in application resources, at run time, deployment time and development time, in a systematic, holistic manner. While CERCIRAS encompasses all kinds of resources that may be involved in the development and usage of systems, it will focus primarily on time, energy, platform, development effort, and data. CERCIRAS will promote new approaches where hardware/software designers, implementers, and even computing systems themselves, are aware of the impact their actions have on resources, and will act to ensure that the system respects the complex, possibly conflicting, set of resource requirements that reflect real user needs. This goal involves investigating system resource usage (e.g. by analysis), as well as advising about that usage. It also involves developing methods, tools, and protocols by which the system itself becomes aware of resource trade-offs. Such autonomous control over resource usage is essential for the large class of automated systems that will need to operate unattended and reliably over extended periods of time, including fully autonomous vehicles (trucks, aerospace, railways [29]), edge-cloud infrastructure, wearable medical devices, etc..

On the technology and knowledge transfer fronts, a significant challenge is to spread excellence across research, application and geographical areas. Societies across Europe (and broader) are at different levels of development. They therefore face different problems, sometimes unknown to, or already solved in, other geographical regions. CERCIRAS will bridge such gaps between domains, sectors and regions, to bring best practices to societies lacking resource awareness, by spreading the available knowledge and experience, by developing new knowledge and experience, by expanding the network of experts, and by tapping into trans-national innovation potential. Fundamental advances beyond the current state-of-the-art are not possible without deep integration of research efforts and solutions from different research communities and user groups, overcoming disjointedness and isolation across research and application areas and national borders. CERCIRAS will be a catalyst for overcoming such boundaries and enabling the desired societal and economic development goals.

An additional challenge is to ensure the short-, medium- and long-term exploitation of achieved results. Initially, academic activities such as professional training and crash courses in Undergraduate, Master and PhD programmes, must be oriented to the needs and demands that arise in the commercial sector. A critical mass of disseminators and maintainers must be created through specific training, adapted to the domain, sector or region targeted. In the long term, the aim of CERCIRAS is to foster the production of new technical solutions, ranging from software tools and products to applications and platforms, that will deliver major socio-economic benefits. Resource-awareness will reduce overall energy consumption and development efforts, ensure necessary performance, while minimizing hardware and total system costs. CERCIRAS will produce technical and scientific roadmaps that will highlight the main research challenges that must be overcome to achieve this.

1.2 PROGRESS BEYOND THE STATE-OF-THE-ART

1.2.1 APPROACH TO THE CHALLENGE AND PROGRESS BEYOND THE STATE-OF-THE-ART

A comprehensive approach to the above challenges requires multidimensional coordinated efforts in three fields: education/training (knowledge transfer), research, and industrial application (technology transfer). Knowledge transfer and technology transfer are tightly related, while they both rely on results in the research domain and vice-versa. For research, CERCIRAS will span the whole system stack, from hardware to application software [30], encompassing the entire life cycle (including usage and maintenance), drawing from relevant theory both from within and beyond Computer Science. The objective of CERCIRAS will be to be a major enabler of advances in the field of resource-aware computing in several ways:

- Development of novel resource models and analyses,
- Integration of multiple resources into a single modelling, analytical and prediction framework; Providing new insights into the complex interactions caused by resource trade-offs;
- Producing new mechanisms to explicitly handle resources at the source code and other levels; Realization of system-predictability and system-wide prediction strategies;
- Researching new ways to inform and advise developers/maintainers about resource usage;
- Focusing on non-traditional resources, including hardware platform, development effort, data, and supporting metrics, as well as traditional resources, such as execution time and energy;
- Developing software engineering methodologies that incorporate resources from the outset, and ensuring that they are transferred to industry;
- Developing smart solutions to advise users of resource usage;
- Developing intelligent solutions that can (self-)adapt according to resource trade-offs

CERCIRAS will enable researchers and industry to liaise and join forces to provide specific solutions for individual resources, their usage and trade-offs:

System execution time/wall-clock time as a resource is already well studied, with a large body of knowledge on methods and analyses being available. Progress in this domain will concentrate on the relation between time and other, often competing, resources: e.g. for hardware platforms, how to minimize the set of hardware components that must be active (i.e., powered) in order to meet the given timing constraints; or for development effort, how to derive early reliable timing estimates with limited expertise and effort. In parallel, concepts from energy analysis, which require detailed hardware models, can be used to improve the precision of current timing analysis techniques. The ultimate goal is to help Hardware/Software designers, implementers and users to make judicious choices using accurate estimates, e.g. about the number of cores and threads to use, the operating frequency and voltage, and, last but not least, about the trade-offs with other resources.

For the very first time, hardware platforms will be treated systematically as a resource in their own right. This will entail sophisticated analyses, so that platforms, or to be exact, the associated costs for hardware and development, can be traded against other resources, such as development effort and wall-clock time. This includes, for instance, design-space exploration, guiding the selection, configuration and optimization of the architecture depending on the time and energy requirements, while considering the available hardware budget. Knowledge of existing trade-offs, and how to exploit them, will lead to better platform designs. Similarly, precise resource estimates require platform design that is amenable to precise analyses, in particular when high single-thread performance involves complex hardware for speculation and other instruction-level parallelism extraction techniques. Novel approaches, such as machine learning, will be used in the analysis of complex systems. A promising initial direction is machine learning prediction of different resource parameters (energy, execution time, etc). The essence is to observe the state changes in the system for a set of training data and synthesis models to "learn", draw correlations and then effectively predict the performance of a new software program.

Data is widely recognised as a critical resource for computing systems. However, data properties have rarely been used in resource optimisation frameworks; integrating them will allow the exploitation of trade-offs concerning data. Obvious trade-offs are between data and platforms, e.g., trading available data size against storage size, or trading data integrity and availability against the costs of memory architectures differing in reliability and speed. Other uses include trading data size, integrity and trustworthiness against the execution time and energy needed to pre-process data. Systematic treatment of data and of its properties will create a new research field targeting data analysis in the context of energy/time/platforms and development efforts.

Modern software engineering practices often waste system resources such as performance, energy, memory and processor power for the sake of reducing development effort, in terms of developer-hours, system reliability, and/or reduced time-to-market. Thus, resource-efficiency of the system and development efficiency almost always conflict. To address this challenge, CERCIRAS will guide the emergence of reliable tools to predict the consumption of various resources early in the design and development process. Models of energy and other system resources, and their transparency through all levels of the system stack, will help developers to combine high-level software engineering abstractions with resource-awareness. A key trade-off is in tool automation vs. resource usage (greater automation brings lower development costs, but also less precise resource estimates).

Regarding knowledge and technology transfer, the primary challenge of CERCIRAS is to bootstrap the "institutionalization" of the Action initiatives so that they outlive the Action and become independent. To this end, CERCIRAS will constitute an integrated team of academic and industrial members, tasked to design, deploy and continually improve the Action's training offering. That team will observe weak points in resource usage patterns and methodologies across the target sectors, domains and regions so as to identify potential opportunities and areas of application/intervention.

CERCIRAS will liaise with local contacts in countries/regions that require specific training, to secure the ability to arrange local events optimally for logistics and outreach in these areas. CERCIRAS will use a staggered approach to enhance its portfolio of training and professional development offer, starting from basic courses to attract young researchers and practitioners, adapted to different target audiences and regions, with feedback ensuring continual improvement. CERCIRAS will leverage this baseline portfolio to increase the outreach, number of offerings, and pilot trainings.

CERCIRAS will address two fronts: (1) a research-oriented front, specifically targeted to early researchers, centred on summer or winter schools; (2) an education and technology transfer front, targeted to an active exchange with current and future industrial practitioners and innovators, and using two primary means, face-to-face sessions or webinars. This will improve research and innovation potential for resource-awareness within European Society as a whole.

The fundamental advances enabled by CERCIRAS will collectively lead to improved resource usage for parallel computing systems, including major improvements in software robustness, reliability and resilience, and to a new educational/training focus on resource awareness in the curricula and training. This will increase the competitiveness of the European society, enabling fully resource-aware and parallel applications especially for multi-/many-core platforms. Considering holistically the most critical resources for system design, deployment and executing, bringing together researchers from all relevant communities and creating a new generation of researchers capable of contributing across multiple domains, CERCIRAS will counteract the shortage of experts in key application areas such as the Internet-of-Things (IoT), Cyber-Physical Systems, Autonomous Vehicles, 5G Networks and other areas.

CERCIRAS will foster a new way of thinking about resources as a general and overarching requirement, leading to lower costs, more secure, more reliable, more adaptable, more responsive, more energy-efficient and more capable computer systems. CERCIRAS will ultimately enable the development of new applications and provide a competitive advantage for the COST member state industries. At the same time, CERCIRAS will enable the optimization of all resources that are connected to the development and deployment of IT systems in general, be they hardware, development, maintenance or power, leading to economic, environmental and societal benefits.

1.2.2 OBJECTIVES

1.2.2.1 Research Coordination Objectives

CERCIRAS will create a cross-domain, cross-stack, community of researchers and practitioners who will join forces around the common theme of resource-aware computing to:

- Master the trade-offs between different resources in modern (self-)adaptive and autonomous computing systems, from small IoT devices to large HPC centres, which increasingly use parallel, heterogeneous, architectures;
- Create the foundations for the development of systems that make better use of these trade-offs;
- Improve resource-awareness and utilisation of next-generation compute-intensive applications;
- Know the requirements and challenges for resource usage of key application areas; and
- Actively promote research results from resource-aware computing into industrial practice and into less aware communities and regions.

CERCIRAS will pursue the following research coordination objectives:

RO.1: Resource analysis and associated costs: the goal is to initiate and coordinate research effort around a new family of tools, approaches and methods for analysing the resource demands and costs (e.g., WCET, energy, data use and its quality, storage) of applications and systems to raise awareness of current and future demands. Ideally, this research will lead to generalized and automated techniques for resource and cost assessments to guide adjustments and optimizations.

RO.2: Models for resource use and requirements: To reason about resource requirements and the trade-offs involved, CERCIRAS aims to coordinate efforts focusing on the construction and validation of cost models and optimizations.

RO.3:Trade-offs: Key enablers for resource-aware computing will be the identification, definition and development of general techniques for trading resources against each other to obtain Pareto-optimal systems, not only in terms of required computing resources, but also considering socio-economic aspects, such as development time and market readiness.

RO.4: Use cases: The research efforts that will be enabled by this Action should allow for fast adoption in several application areas and for a wide-variety of resources, including the Internet-of-Things, home automation, cyber-physical systems, robotics, autonomous vehicles, data centres and cloud computing and the resources mentioned above. CERCIRAS achieves this by encouraging the development of use cases in academic / industrial partnerships.

CERCIRAS will pursue these objectives by coordinating efforts from several existing research projects and communities on the new challenges of resource-aware computing. The Action will bring together researchers and industrial experts who are interested in, inter-alia:

- Resource analysis: worst-case execution time analysis; energy consumption; data as a resource; cache and memory-aware computing;
- Architectures: multi-core/many-core (possibly heterogeneous) parallel hardware design;
- Applications: the Internet-of-Things, including home automation; Cyber-Physical Systems, including robotics, autonomous vehicles; Data Centres and cloud computing;
- Programming language abstractions for capturing and reasoning about resources; and
- Mathematical models of resource usage, and fundamental theories.

Practically all these engineering disciplines have succeeded in identifying approaches that find a balance between the competing constraints within their own vertical domain. However, rapid development of multi-core/many-core computing has resulted in complex software systems where available resources span multiple domain-specific concerns. Abstractions that facilitate research on such complex situations have isolated researchers with rare interaction across the levels of the stack. Trade-offs between resources are often no longer clearly visible.

To achieve these objectives, CERCIRAS will create synergies between communities that are currently confined to their own domain/abstraction level, allowing them to employ their techniques more widely. Such synergies will enable the development of new holistic approaches and provide solutions to complex resource usage challenges for the upcoming architecture generations, covering and integrating hardware, operating systems, middleware, compilers, development tools, and applications.

Specific research coordination activities of the CERCIRAS COST Action will include:

1. Facilitating cross-domain research visits, and by organizing technical workshops, industrial meetings, collocations and ad-hoc collaborations among key academic researchers, industrial experts and technologists to discuss, define, and explore the trade-offs in hardware/software resource use and to ultimately develop new methods that take advantage of explicit and implicit trade-offs between different resources in these systems.
2. Organising cross-cutting workshops, specific scientific missions, developing common research roadmaps, and encouraging novel cross-cutting research publications and research proposals to foster the development of synergistic views of the resource awareness problem.
3. Promoting interactions between funded national and international research projects focusing on resource awareness for multi-core/many-core systems. This will be achieved by inviting representatives of these projects to propose/attend research meetings and to engage with Short-term Scientific Missions, Training Schools and other CERCIRAS activities.
4. Ensuring a common understanding of scientific terminology and correct methodological practice as it relates to resource awareness. The lack of common language and methods is one of the major blocks in resource-aware computing. This will be achieved by producing white papers and other authoritative documents that distil expertise across individual research areas.

1.2.2.2 Capacity-building Objectives

The activities of the Action will lead to long-term capacity-building scientific, technical and economical advances in resource-aware computing for multicore/many-core systems, including:

- the establishment of cross-sectorial new collaborations, designed to outlast the lifespan of the Action and to inspire new research initiatives on public and private funding;
- a portfolio of training and educational resources designed to feed forward knowledge and technology transfer efforts;
- white papers to document the rationale, the design and the implementation of CERCIRAS, to present its outcomes and lessons learned, and to propose a roadmap of initiatives that will build on the proceeds of the Action;
- and, of course, a significant number of joint research publications that will become foundational material for research into resource-aware parallel computing.

The specific capacity-building objectives of CERCIRAS are to:

1. Balance and raise the awareness of systems resource usage among researchers and practitioners of computing systems, by promoting and facilitating the integration of aspects of resource usage directly into Computer Science curricula, starting from the undergraduate level.
2. Balance and increase the number of trained researchers and PhD students with training in holistic resource-usage analysis, modelling, verification, and optimization techniques for computing systems, by organizing workshops, summer schools and other networking activities on resource-aware computing.
3. Expand knowledge of resource usage techniques for parallel computing systems among all the stakeholders, students, researchers, and practitioners, by organizing crash courses, training materials, and mutual research visits, to kick-start joint interdisciplinary efforts and projects that will tackle pressing issues of resource awareness.
4. Balance and raise awareness in the multi-/many-core hardware community on predictability barriers to mitigate the large discrepancies between resource usage and system predictability, through the active involvement of the hardware expert members of the Action.
5. Motivate new collaborative national and international research projects to work on holistic resource modelling and analysis for parallel computing systems, by cross-networking and liaising with public and private research funding agencies, insisting on the strategic importance of the topic.
6. Provide timely and sustained knowledge transfer between academia and industry, ensuring that academic research on resource modelling and analysis addresses industry needs, in constant consultation with the industrial partners of the Action and a network of external contact points.

2 NETWORKING EXCELLENCE

2.1 ADDED VALUE OF NETWORKING IN S&T EXCELLENCE

2.1.1 ADDED VALUE IN RELATION TO EXISTING EFFORTS AT EUROPEAN AND/OR INTERNATIONAL LEVEL

Effective networking is essential if the Action is to successfully address the challenges identified above. Their scope and complexity mean that they cannot be addressed via a couple of localized projects or top-down initiatives. Rather, their solution requires a fusion of existing communities from relevant challenge domains, covering the entire COST area. The strength of CERCIRAS lies precisely in the cross- fertilization that will be engendered among the various emerging research and user communities. It also lies in the combination and unification of state-of-the-art techniques and terminology, across these different communities. CERCIRAS will help to overcome the significant fragmentation between the disparate communities that can benefit from resource-awareness. It will also create a focus on resource-awareness beyond traditional research, educational and technological areas and geographical boundaries, including the COST Inclusiveness Target Countries and International Partner Countries.

CERCIRAS is an exceptional opportunity to promote close collaboration between so far isolated research communities and industrial players. This will be supported by the two-way interaction with the educational participants, enabling a holistic treatment across the full range of resources. There is no similar initiative either at the European level or beyond. Current or past projects that are related to resource-aware computing, e.g. in FP7 or H2020, have been vertically confined to a single resource and to a single research community and its related industry. For instance, worst-case execution time analysis for embedded systems, has been the focus of several collaborative research activities, including ARGO (H2020-688131), PROXIMA (FP7-611085), P-SOCRATES (FP7-611016) Hercules (H2020-688860), and T-CREST (FP7-288008), or the former COST Action TACLe (IC1202); AQUAS (H2020-737475) explores tools to improve development time for safety-critical systems when security and resilience concerns arise; and energy as a resource has recently gained significant momentum through projects such as ENTRA (FP7-FET-318337), EXCESS (FP7-611183), ANTAREX (H2020-671623) and ICT-Energy (FP7-611004), or the COST Action EE-LSDS (IC0804). Other projects focus on computational efficiency and, occasionally, programming, but without explicitly considering either the computing platform or the human programmer as a limited and valuable resource. Moreover, they neither make a connection to time or energy as relevant resources. Examples include POLCA (FP7-610686), ParaPhrase (FP7-288570), COMPASS (FP7-287829), DreamCloud (FP7-611411), Rephrase (H2020-644235), and the COST Actions Game Theory (CA16228), ComplexHPC (IC0805), NESUS (IC1305), cHIPSet (IC1406), and Reversible Computation (IC1405). CERCIRAS is also related to HiPEAC, the European Network on High Performance and Embedded Architecture and Compilation. CERCIRAS is significantly smaller in size than HiPEAC but much more cohesively focused. This will allow more effective research inter-play and articulation. Nevertheless, HiPEAC, with its annual conference and regular Computing Systems Weeks will provide a complementary environment for co-locating CERCIRAS events, including plenary and management meetings, workshops and training sessions. Moreover, CERCIRAS also links with relevant initiatives in new approaches to computing, such as the IEEE Rebooting Computing initiative, which looks at new technologies in a broader sense (e.g., neuromorphic or quantum computing) and with multiple Special Interest Groups (SIGs) of the Association for Computing Machinery (ACM). Finally, CERCIRAS relates to ACM and its SIGs through its educational and training aspects, as well as to appropriate international education and research programmes, such as ERASMUS+, the H2020 Marie Skłodowska-Curie or regional activities (e.g., from DAAD, or FNR's AFR project).

CERCIRAS participants are already actively involved in several European and international activities of relevance to the Action. This involvement offers a formidable initial boost to know-how, team building, and networking. In addition, the Action will proactively seek to reach out to participants in other relevant projects. These links will bridge existing activities with the new Action research coordination effort. As noted, these activities have orthogonal goals to CERCIRAS, and have related, but different, aims. CERCIRAS will thus provide an opportunity to combine the current, rather isolated, research, training and technology transfer initiatives and to create a new research field focusing on overall resource-awareness without boundaries. The latest research results will be made available via CERCIRAS workshops and roadmaps to a broader audience of fellow researchers and industrial practitioners, who would otherwise be unaware of this ongoing research. Industrial guidance on research directions, validity, usability and applicability will be highly valuable in setting the research send. Finally, innovative knowledge will be transferred among the education, research and industry by an audience-adaptive training program. This will create synergies and opportunities for future follow-up research, educational and technology transfer projects.

2.2 ADDED VALUE OF NETWORKING IN IMPACT

2.2.1 SECURING THE CRITICAL MASS AND EXPERTISE

CERCIRAS already includes relevant stakeholders from academia and industry, spanning all layers of the hardware/software stack. Industrial stakeholders come from both large enterprises and SMEs. Many CERCIRAS members are engaged in standardization activities and well experienced in research coordination, research/industry liaison, and industrial adoption activities, but CERCIRAS includes also many young researchers and rising talents, and intends to involve even more. CERCIRAS will focus on outreach to relevant research communities, industry stakeholders, and the general public. Furthermore, academic participants will promote the introduction of resource concerns into curricula and courses including, e.g., software technology, computer architecture or technical computer science. In this way, CERCIRAS will have a strong impact on students who will then enter industry and research and thus propagate the importance of resource-awareness, especially in regions with low resource awareness. Multidisciplinarity of the network will be achieved through tight connections between the Action participants and experts from other research areas (e.g., mathematics, statistics, and economics) in their everyday work environment. Where relevant, CERCIRAS will invite these and other relevant experts to its training schools and workshops.

Specific activities will be performed during the Action to further increase visibility and impact among researchers and industrial practitioners, in order to create a critical mass that will build a self-sustaining influential community. At the research and innovation level, scientific and professional workshops, panels and tutorials at relevant events will be organized. At the technology transfer and application level the Action will provide technical white papers, joint adaptive training materials, will organise trainings and workshops, and will promote professional training either as dedicated events or within relevant industrial events. Based on this training material, several profiled training activities will be generated (crash courses, training schools, short specialization curriculums, academic courses within existing curricula, separate dedicated Masters programmes, etc.) and will be organized in low- awareness regions. Relevant outputs (e.g., training materials, white papers, case studies, benchmarks) are intended to continue delivering impact, long after the Action lifespan.

CERCIRAS will make a constant effort to identify emerging European CSA and R&D projects, both in directly related domains, and in orthogonal domains, so as to include them in its effort.

2.2.2 INVOLVEMENT OF STAKEHOLDERS

The most relevant stakeholders are recognized at all three levels: (1) at the research and innovation level, academic stakeholders include universities and research institutes with a focus in computer science and engineering, in particular cyberphysical, real-time and embedded systems, as well as centres for HPC and big data with the involvement of mathematics, statistics, management, etc.; (2) at the technology transfer and application level, the most relevant stakeholders are found in industry sectors that (a) will be using high-performance, resource-constrained embedded systems, or (b) will provide high-performance or big data (cloud) computing services, and (c) are not aware of potential benefits from applying resource trade-offs. Potential beneficiaries include large parts of the high-tech sector such as automotive, aerospace, avionics, automation and robotics sectors, and medical technologies, as well as all third party users of these sectors within society as a whole. Stakeholders at the lower levels are vendors of tools for embedded systems development (compilers, analysis tools, integrated development environments (IDEs), etc); (3) at the educational and training level, stakeholders include (a) educational institutions and innovative organisations representing the resource-aware communities, who are tasked to spread awareness, (b) representatives of technology transfer initiatives whose task is to guide the training process to deliver the greatest practical benefit, and (c) educational institutions and innovative organisations representing less resource-aware regions, whose task is logistical and dissemination.

2.2.3 MUTUAL BENEFITS OF THE INVOLVEMENT OF SECONDARY PROPOSERS FROM NEAR NEIGHBOUR OR INTERNATIONAL PARTNER COUNTRIES OR INTERNATIONAL ORGANISATIONS

CERCIRAS will create a focus on resource-awareness beyond traditional research, educational and technological areas and geographical boundaries, including the Near Neighbour Countries and International Partner Countries. It has a specific focus on inclusion for these countries, with the goal of enhancing research and innovation activities and promoting high-level interactions with leading companies, researchers and institutions. This will create economic and societal benefits within these countries, enabling them to engage on equal terms with researchers and industry in the broad COST area and beyond, creating new research and innovation opportunities, and facilitating the creation of the key technological advances that are necessary for the creation of a modern technological society built around the deployment of advanced AI and other technologies in transport, manufacturing, medicine and other sectors. During the Action implementation, through dissemination and knowledge and technology transfer activities, the community will overcome existing boundaries and reach out to near neighbour countries and international partner countries, spreading, at the same time, resource awareness and all impacts and benefits in the focus of this Action.

3 IMPACT

3.1 IMPACT TO SCIENCE, SOCIETY AND COMPETITIVENESS, AND POTENTIAL FOR INNOVATION/BREAK-THROUGHS

3.1.1 SCIENTIFIC, TECHNOLOGICAL, AND/OR SOCIOECONOMIC IMPACTS (INCLUDING POTENTIAL INNOVATIONS AND/OR BREAKTHROUGHS)

CERCIRAS will enable major scientific and technical breakthroughs in the understanding of complex interactions between resources and their usage for parallel computing systems. There are obvious and significant scientific, technological and socio-economic advantages to achieving the Action's aims. Computing systems are fundamental to most aspects of modern life, with the software industry worth over €375 billion per annum [31]. Emerging trends, e.g. in IoT and AI, highlight the dramatic impact of computing technology on everyday life. Parallel multi-/many-core systems are crucial to these advances: achieving the CERCIRAS breakthroughs will be fundamental to ensure the successful adoption of new and emerging processor architectures. CERCIRAS will enable the development of new classes of parallel computing applications, based on, e.g., combining low energy and high performance. At present, understanding and managing the resource landscape for parallel systems is beyond the comprehension of the vast majority of users. There is currently only limited work in this area, it is uncoordinated, it lacks focus, and research needs to reach a level where it can be adopted by industry as a whole. While Europe leads research in this area, with excellent research groups, these must be brought together to tackle the challenges discussed in Section 1.1.2 CERCIRAS brings a pan-European, COST-level, perspective to bear on resource analysis for parallel computing, raising the Technology Readiness Level (TRL) of industrial work in this area, and providing a focus for research and development effort in universities and industry in the COST area.

Initial impact will be perceivable at lower TRLs with medium term impact on actual industrial use. A major current technological drive is the interplay of distributed embedded systems and cloud computing, where distributed and embedded software units work together for a common goal, and collectively consume a variety of resources. This concern is fundamental in such areas as: i) Big Data, ii) the Internet of Things (IoT), and iii) Cyber Physical Systems (CPS). The next step in this evolution is confluence into new smart "places", such as buildings, transportation systems, city areas, etc., with the challenge of tackling the span and diversity of the system requirements. For example, while time is a major concern in Cyber Physical Systems, IoT focuses on power consumption and communication, and Big Data is concerned with throughput and energy, and may possibly benefit the most from new high-bandwidth memory technologies. The aim of CERCIRAS is to aggregate experts from different resource domains and different sectors of the society into shared, common and integrated consideration of holistically effective responses to resource-related challenges. As an immediate outcome, CERCIRAS will produce significant scientific and technological results with short-term impacts, yielding a number of important social and economic cascading impacts in the longer term.

Tangible, quantifiable scientific results with academic and technological short-term effects include:

1. Advancing research on software for parallel and distributed multi-/many-core systems from different resource perspectives by developing a multi-objective optimization process;
2. Guiding novel H/W platform development by exposing resource bottlenecks in existing systems; Creating an international focus and community around resource analysis for parallel computing;
3. Conceiving new ways to model software systems, able to analyse and integrate multiple dimensions of computing resources;
4. Providing new algorithms to optimize resource consumption based on specific metrics;
5. Providing new algorithms to configure resources at design time and optimise them at run-time; Enhancing existing collaborations, and creating opportunities for new ones;
6. Fostering European research to crossbreed approaches & solutions, reducing research silos/risk.

Tangible, quantifiable technological results with medium-term impact on industrial practice, include:

1. Development of new design tools for optimizing resource consumption;
2. New techniques for operating systems, integrating advanced resource management;
3. Resource usage comparison benchmarks supporting cross resource evaluation mechanisms.

Tangible, quantifiable, longer-term economic/societal benefits include:

1. Pushing Europe to the forefront of parallel and distributed multi-core/many-core computing;
2. Pushing Europe to the forefront of CPS, IoT and Big Data integration and smart environments;
3. Increasing the market penetration of state-of-the-art design tools and runtimes, with consequent improvement in the performance and efficiency of (distributed) computing systems;
4. Reducing resource requirements of (distributed) computing systems and energy footprint and ICT increasing the associated industrial productivity;

5. Reinforcing Europe's capacity for creating and attracting highly-skilled employment in the development of high-performance systems and smart environments

3.2 MEASURES TO MAXIMISE IMPACT

3.2.1 KNOWLEDGE CREATION, TRANSFER OF KNOWLEDGE AND CAREER DEVELOPMENT

In the area of education and knowledge transfer, the major mechanism for meeting the Action's long-term objectives will be to bootstrap the Action's educational initiatives onto an institution/entity that will outlive the Action. This will ensure that these educational offerings will be recurrent and independent of the Action. In the shorter term, during the lifespan of the Action itself, a mixed team composed of both academic and industrial partners, is to be formed, with the task of designing, deploying and continually improving the Action's education offering. Additionally, by liaising with local contacts in each country/region that is a special target of its educational efforts, the Action would arrange local events, optimised for logistics and local reach and focus.

3.2.2 PLAN FOR DISSEMINATION AND/OR EXPLOITATION AND DIALOGUE WITH THE GENERAL PUBLIC OR POLICY

Regarding the training and career development objectives, a staggered roadmap is envisioned. In the first couple of years, the Action will offer a baseline number of education/training opportunities. The focus will be on building a flexible portfolio of such offerings, easily adaptable to different target audiences/regions. The feedback and experience from those will permit continual improvement and consolidation. In the next couple of years, the Action will leverage the baseline portfolio to increase the outreach, the number of offerings, and the pilot trainings, with continual improvement.

The Action distinguishes between two fronts, in terms of knowledge and technology transfer. On the research-oriented front, the Action's efforts will be especially targeted to early researchers, centred on Training Schools. On the education and technology transfer front, the Action's efforts are targeted to current and future industrial practitioners. Described activities appear in both, live and electronic form, making it even more adaptive to domainial, sectorial and regional needs.

The foreseen dissemination activities include:

General materials: Logo, brochures, templates, web site, and other actions that create identity, consistency and awareness of the CERCIRAS COST Action;

Event participation: CERCIRAS should actively participate in 6+ relevant European and worldwide Events (Conferences, Fairs, Workshops, etc.) per year, such as the HiPEAC Conference and HiPEAC Computing System Weeks, DATE, and area-specific conferences (e.g. ECRTS, CPSWeek, ESWeek);

Publications: with an Open Science model in mind, at least 7 articles/papers should be produced per year for academic and industrial conferences and technical journals, preferably under open access models, jointly authored by both academic and industrial participants. Targeted journals include the Journal of Embedded Computing, Software: Practice & Experience, IEEE Transactions on Computers, Software and Systems Modelling, IEEE Access, among others;

Industrial engagement: a regular series (2 per year) of dissemination events will be organised, targeting relevant SMEs and large companies;

Training: a training curriculum will be developed for a wide dissemination and targeted at two professional groups: industrialists and researchers. Innovative packaging of project knowledge including examples, exercises, support material and knowledge appraisal will be delivered through Internet-based training and incorporated into existing course provisioning;

Social media: to better inform the general public, CERCIRAS will engage with Twitter, Facebook, LinkedIn, and any other relevant social media, keeping appraised the central COST channels;

Website: a "portal" serving as a public repository for outputs (training materials, benchmarks, case studies). Post-Action sustainability will be addressed by Working Groups (within, e.g., ACM SIGs, IFIP WGs, EIT or EDAA), following other examples, such as EMSIG, cHiPSet and NESUS;

Increasing awareness: the Action will actively increase awareness of the Action activities, targeting ITC countries, under-represented industrial domains, undergraduate/masters courses, industry consortia and standardization bodies (IEEE 2415 and IEEE 2413, among others).

Exploitation Plan. Results will be directly exploited by participants, including in Undergraduate, Masters and PhD education programmes, professional training events, research and commercial projects. technical and scientific roadmaps will be produced to highlight the main research challenges that must be overcome, thus guiding future research activities. Industrial participants are key to identify and promote the development of relevant case studies. The synergies that will be created by joining the different communities will, in the long term, lead to the production of new software tooling, new products and new applications with major socio-economic benefit. Moreover, by increasing the density and spread of the CERCIRAS network, the Action will engage with, promote outreach to, and ensure better transfer between countries and communities that could benefit from resource awareness. Outreach is a key factor in the Action, with WG5 specifically dedicated to these activities.

4 IMPLEMENTATION

4.1 COHERENCE AND EFFECTIVENESS OF THE WORK PLAN

4.1.1 DESCRIPTION OF WORKING GROUPS, TASKS AND ACTIVITIES

The organisation into Working Groups (WGs) reflects a vision of collaboration among previously segregated communities. Instead of the “obvious” and previously applied approach of aligning along resource types, CERCIRAS WGs are cross-cutting and explicitly promote exchange among communities, via specific but interconnected Tasks (T), with mixed participation of experts from different projects, domains and sectors to maximize cross-fertilization. WG1, WG2, and WG3 focus on research goals directed to preparing innovative knowledge, techniques and tools for knowledge and technology transfer, while WG4 and WG5 focus on performing this transfer and on disseminating the knowledge gained. WGs consolidate results and coordinate their effort during the Action’s workshops, and also on a bilateral basis, and add to the annually reported achievements.

WG1. Identifying costs associated with resources

T1.1. Cost assignment for use cases: For the different use cases/application scenarios covered by CERCIRAS, the relevant types of costs associated with each resource type will be identified, ranked by importance, and quantified where applicable (e.g. energy consumption might be negligible for stationary sensors connected to the power grid, but is critical for mobile or battery-powered sensors).

T1.2. Generalizing cost assignment techniques: A logical extension of T1.1 is to aim for general techniques to assign costs to the usage of resources in arbitrary use cases. These techniques must be complete (i.e., capable of assigning costs to all relevant resources) and consistent.

T1.3. Validation of cost models: CERCIRAS will analyse and compare different cost models for resources. To this end, the Action aims to collect a list of models that are validated on real systems and make them available to a wider research community and to industry.

Deliverables cut across the WGs, consolidating their findings and aligned to project milestones. See description below.

WG2. Analysing trade-offs based on individual resource analysis

T2.1. Trade-off identification: It is important to make trade-offs explicit by identifying the resources that compete against each other and by characterizing how this competition impacts system behaviour and properties. CERCIRAS focuses on the implicit trade-offs that lead to suboptimal results if not considered properly, e.g., the trade-off between precision of a timing or energy analysis versus development time, which will often lead to either imprecise results or delayed system release.

T2.2. Formal definition of trade-offs: CERCIRAS will examine commonalities between trade-offs to devise a formal concept of trade-offs that can replace the otherwise mostly intuitive notion. A formal definition will enable correct reasoning/discussion of trade-offs, and thus also better exploitation.

T2.3. General techniques for trade-off identification: The Action will research and develop general techniques to identify trade-offs within the development process and at system run-time. General techniques are required in order to capture novel resource types and trade-offs that have not yet been identified, in a systematic and rigorous way.

WG3 – Exploiting resource trade-offs

T3.1. Incorporating design-time advice: CERCIRAS will identify and understand current and novel approaches for the design-time exploration of trade-offs, considering overall system requirements and its potential dynamic behaviour. This involves identifying contexts for static resource access models and static resource access criteria, to achieve predictable and efficient resource access, as well as exploring hardware-architecture support for resource trade-offs.

T3.2. Runtime management: Analogously, CERCIRAS will identify and understand current and new approaches for run-time management of resources, considering both guaranteed levels of resources as well as best-effort approaches, including software and hardware mechanisms. CERCIRAS will also study mixed-criticality access models to provide asymmetric resource guarantees.

T3.3. Multi-objective resource allocation evaluation metrics: CERCIRAS will identify resource allocation approaches are suitable for optimisation with respect to multi-objective utility metrics. Examples are, e.g., scalability towards a high number of users and efficiency for a low number of users. CERCIRAS will develop hybrid resource allocation strategies consisting of different allocation strategies enabling adaptation to multiple resource allocation evaluation metrics.

WG4 – Knowledge and technology transfer

T4.1. Connecting within ICT disciplines: Connecting resource-aware analysis with related ICT disciplines: Parallel programming can be complicated and the analysis of such programs even more so. The parallel programming community has devised informal and ad-hoc analysis techniques but these are neither general nor sound. CERCIRAS will focus on identifying such informal analysis techniques, and investigate how they might help construct more effective resource-aware analytical techniques. Other areas of ICT of interests include business process management, software quality and project management, while areas that might benefit from the results of CERCIRAS are IoT, CPS, Autonomous Systems, Robotics, (Self-)Adaptive Systems and Cyber Security.

T4.2. Connecting to non-ICT disciplines: As in T4.1, CERCIRAS will reach out to non-ICT disciplines: An impressive variety of methodologies, from the areas of mathematics, probability and statistical theory, could potentially be used for multi-resource-aware analysis. Researchers in system design are only familiar with few of these and the mathematical community is largely oblivious to applications of its theory to resource aware design. So far, most analysis techniques applied to computing systems have their roots in engineering and various branches of mathematics and computation theory. It is therefore worth exploring techniques other domains (e.g., economic theory). CERCIRAS will cross-educate these different communities via seminars, events and learning materials, to foster cross-pollination of ideas and results.

T4.3. Connecting and knowledge transfer to industry partners: Research results acquire value when put to practice for innovation. CERCIRAS will bring together academic and industrial actors inside and outside the consortium with the intention of problem-solving among these two sectors.

WG5 – Outreach and External Relations

T5.1. Industrial liaison: The Action will hold: a) technical days to disseminate results and to foster specific collaborations; and ii) informal workshops, at the end of the Action's lifetime, to capture how the industry interprets and ranks the Action's goals. CERCIRAS targets key industrial events e.g. Embedded World, HiPEAC, Artemis-IA, EDAA and BDVA events (to which the Action members already have links).

T5.2. Relations with other networks and communities: The Action will be made visible and attractive, in terms of thematic focus, membership, initiatives, achievements, etc. Targets include the networks and communities already created on all three fronts (research, education and industry) including relevant actors from all European regions. Furthermore, relevant H2020 CSAs, Erasmus+ partnerships and COST Activities will run in parallel with the Action to make connection between communities stronger according the general interests of European Society.

T5.3. Case studies and benchmarks: An important objective of the Working Group will be to identify a set of representative industrial case studies and benchmarks that can be openly used by researchers and practitioners. The involvement of the industrial experts in the Action will help achieve this goal. The Action will make these publicly available through a specific web site, and procedures will be put into place for the web site to outlive the Action (as described in Section 2).

T5.4. Impact on training and education: The expertise gathered in the Action will be used to reach out to students at all levels from undergraduate to PhD, as well as practitioners, to enable exchange of knowledge . This will afford them a rich, fresh, and authoritative exposure to education in resource awareness. CERCIRAS achieves this by creating a unique training material adaptive to trainees from different domains, sectors or regions, suitable for: a) organizing Training Schools; ii) organizing professional development programs; and iii) engaging other Higher-Education institutions/bodies to revise their curricula to systematically cover resource awareness and related disciplines.

4.1.2 DESCRIPTION OF DELIVERABLES AND TIMEFRAME

Milestone M1 (month 12) will identify resource costs from use cases (WG1), and resource trade-offs (WG2), for use by all WGs. Activities in WG3 and WG4 will survey existing approaches and alternatives. Deliverables: First Annual Report of the Action, including WG status (D1.1); Proceedings of the first CERCIRAS Cost Action workshop (D1.2); Initial version of training materials from WG5 (D1.3) training material will be subsequently improved with new results and lessons learned from internal and external training sessions; State-of-the-art and roadmap with inputs from WG1-4 (D1.4) .

M2 (month 24) will identify initial approaches to design-time and run-time management of multiple resources (WG3), to be analysed by all WGs. Gap analysis will be applied to all WGs/activities to identify important missing topics and alter the work plan. A pilot Training School will be organized. Deliverables: Second Annual Report of the Action, incl. WG reports (D2.1); Proceedings of the second workshop (D2.2); Improved training material (D2.3); Case studies and benchmarks (D2.4).

M3 will shift the focus of the activities to validation of cost models (WG1), generalisation of techniques (WG2) and multi-objective evaluation metrics (WG3). Training materials will be improved based on experiences from the Training School, as well as on input from sector and domain experts (WG4 and WG5). CERCIRAS will cross-check all WGs, to identify and prioritize the activities for the final year of the Action. Deliverables: Third Annual Report of the Action, including WG reports (D3.1); Proceedings of the third workshop (D3.2); Improved training material (D.3.3).

M4 marks the end of the Action activities. This final period is oriented to knowledge/technology transfer, intensive exchange among domains, sectors and regions, dissemination and ensured future exploitation of Action results. Deliverables: Final report of the Action, including WG1-5 reports (D4.1). Published proceedings of the fourth workshop (D4.2). Publication of adjusted training material (D4.3).

4.1.3 RISK ANALYSIS AND CONTINGENCY PLANS

Risk	Description	Prob.	Severity	Mitigation
R1	Insufficient industry involvement	M	M	The CERCIRAS consortium already includes relevant industry experts, from different domains, with track records of collaborations with academia. Should industry not respond to the planned activities during industrial events, then efforts will intensify, leveraging links of participants to industry and industrial members. Furthermore, the Action will promote the use of industrial studies and benchmarks in research activities to increase the industrial relevance and visibility of the research.
R2	Lack of engagement by different communities	L	H	The CERCIRAS work plan foresees joint events co-located with different communities to promote the Action and community engagement. The Action's activities will be proactively taken to existing events. Active participation by the different communities during the proposal phase provides further confidence in successful engagement.
R3	Difficulties in reaching consensus about resource definition models	L	H	Different communities already displayed good prospects for developing cross- community work before the start of the Action. This will be reinforced during CERCIRAS through *cross-cutting* working groups, constructed orthogonally to specific resource types and expert areas, with the goal to maximize collaborative efforts.
R4	Difficulties in including new resources in the holistic model	L	H	Different communities have already analysed before the start of the Action how different resources can be jointly considered. The structuring of the workplan in *cross- cutting* working groups will reinforce this integration, allowing to quickly identify difficulties in integration, and proposing solutions for addressing the holistic model.
R5	Lack of engagement of early researchers from ITC	L	H	From the outset, the preparation of CERCIRAS has involved young researchers in leadership activities, including from ITC. As presented in the COST Mission and Policies, activities will include a focus on further promoting and early on involving other young researchers from ITC.
R6	Difficulties in attaining international visibility	L	M	CERCIRAS participants include renowned researchers, with high-level links to relevant international forums and communities. This gives international visibility of the Action work. The Action also includes participants from non-COST countries, which allows linkage with world-wide communities.
R7	Failure to characterize resource costs or trade-offs	L	M	WGs 1 and 2 may face difficulties in characterizing certain resources or resource combination trade-offs. This can be mitigated by connecting experts knowledgeable in different resources and application areas. Any shortfalls are expected to be specific to certain resources or their combinations, which limits impact.
R8	Failure to adapt resources	M	L	Some resources may not be manageable at run-time / design- time. This risk will be compensated for, by reverting to the relevant alternative management scheme.

4.1.4 GANTT DIAGRAM

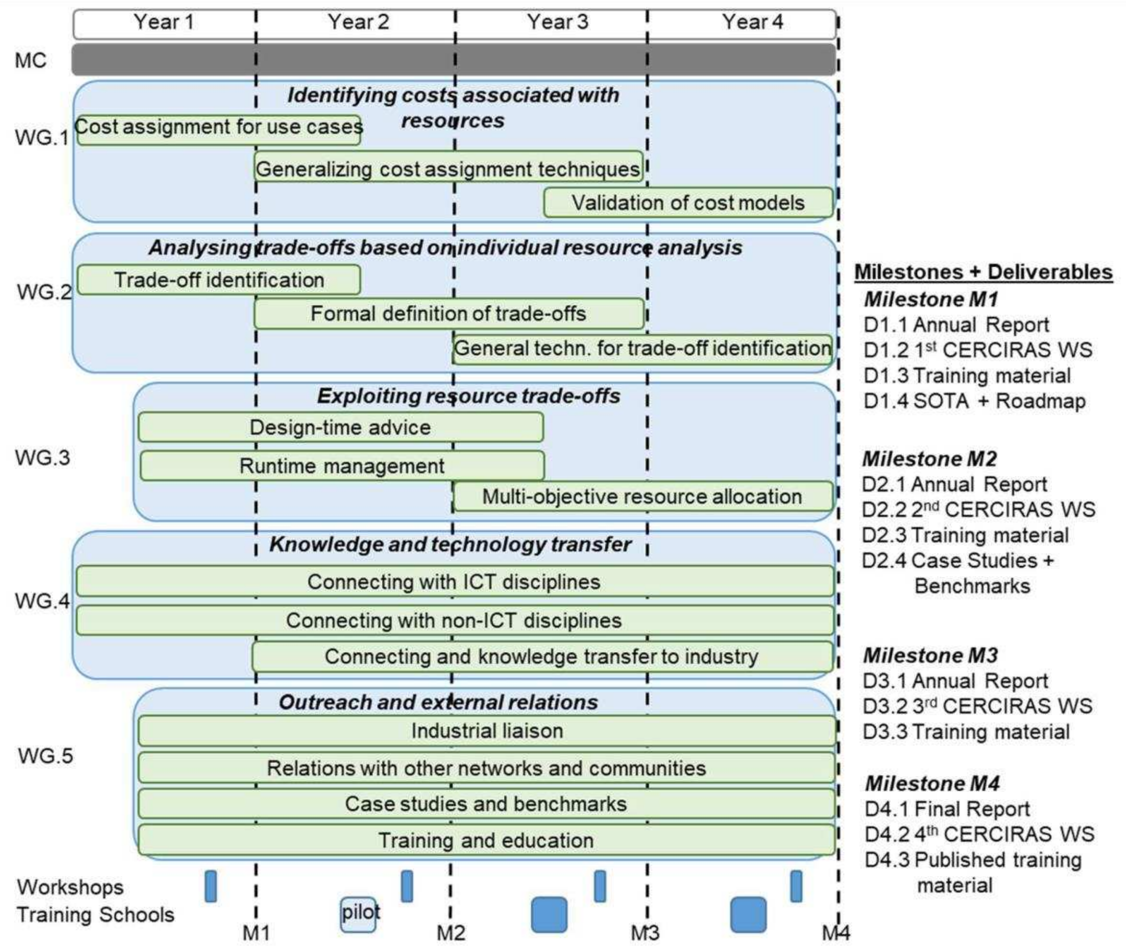

REFERENCES

[1] Duranton, M., De Bosschere, K., Coppens, B., Gamrat, C., Gray, M., Munk, H., Ozer, E., Vardanega, T. and Zendra, O., 2019. The HiPEAC Vision 2019. online https://www.hipeac.net/vision/#/latest/

[2] The International Technology Roadmap for Semiconductors 2.0: 2015, . Semiconductor Industry Association. Online: https://www.semiconductors.org/resources/2015-international-technology-roadmap-for-semiconductors-itrs/

[3] Davis, R.I. and Burns, A., 2011. A survey of hard real-time scheduling for multiprocessor systems. ACM computing surveys (CSUR), 43(4), pp.1-44.

[4] Liu, Y. and Zhu, H., 2010. A survey of the research on power management techniques for high‐performance systems. Software: Practice and Experience, 40(11), pp.943-964.

[5] Orgerie, A.C., Assuncao, M.D.D. and Lefevre, L., 2014. A survey on techniques for improving the energy efficiency of large-scale distributed systems. ACM Computing Surveys (CSUR), 46(4), pp.1-31.

[7] Lv, M., Guan, N., Reineke, J., Wilhelm, R. and Yi, W., 2016. A survey on static cache analysis for real-time systems. Leibniz Transactions on Embedded Systems, 3(1), pp.05-1.

[8] Ababei, C. and Moghaddam, M.G., 2018. A survey of prediction and classification techniques in multicore processor systems. IEEE Transactions on Parallel and Distributed Systems, 30(5), pp.1184-1200.

[9] Wilhelm, R., Engblom, J., Ermedahl, A., Holsti, N., Thesing, S., Whalley, D., Bernat, G., Ferdinand, C., Heckmann, R., Mitra, T. and Mueller, F., 2008. The worst-case execution-time problem—overview of methods and survey of tools. ACM Transactions on Embedded Computing Systems (TECS), 7(3), pp.1-53.

[10] Sha, L., Abdelzaher, T., Cervin, A., Baker, T., Burns, A., Buttazzo, G., Caccamo, M., Lehoczky, J. and Mok, A.K., 2004. Real time scheduling theory: A historical perspective. Real-time systems, 28(2-3), pp.101-155.

[11] Cazorla, F.J., Quiñones, E., Vardanega, T., Cucu, L., Triquet, B., Bernat, G., Berger, E., Abella, J., Wartel, F., Houston, M. and Santinelli, L., 2013. Proartis: Probabilistically analyzable real-time systems. ACM Transactions on Embedded Computing Systems (TECS), 12(2s), pp.1-26.

[12] Bonenfant, A., Claraz, D., De Michiel, M. and Sotin, P., 2017. Early WCET prediction using machine learning. In 17th International Workshop on Worst-Case Execution Time Analysis (WCET 2017). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[13] Le Sueur, E. and Heiser, G., 2010, October. Dynamic voltage and frequency scaling: The laws of diminishing returns. In Proceedings of the 2010 international conference on Power aware computing and systems (pp. 1-8).

[14] Bambagini, M., Marinoni, M., Aydin, H. and Buttazzo, G., 2016. Energy-aware scheduling for real-time systems: A survey. ACM Transactions on Embedded Computing Systems (TECS), 15(1), pp.1-34.

[15] Georgiou, K., Kerrison, S., Chamski, Z. and Eder, K., 2017. Energy transparency for deeply embedded programs. ACM Transactions on Architecture and Code Optimization (TACO), 14(1), pp.1-26.

[16] Sha, L., Caccamo, M., Mancuso, R., Kim, J.E., Yoon, M.K., Pellizzoni, R., Yun, H., Kegley, R., Perlman, D., Arundale, G. and Bradford, R., 2014. Single core equivalent virtual machines for hard real—time computing on multicore processors.

[17] Castrillon, J., Lieber, M., Klueppelholz, S., Völp, M., Asmussen, N., Assmann, U., Baader, F., Baier, C., Fettweis, G., Froehlich, J. and Goens, A., 2017. A hardware/software stack for heterogeneous systems. IEEE Transactions on Multi-Scale Computing Systems, 4(3), pp.243-259.

[19] Batini, C., Cappiello, C., Francalanci, C. and Maurino, A., 2009. Methodologies for data quality assessment and improvement. ACM computing surveys (CSUR), 41(3), pp.1-52.

[20] Genius, D., 2013, July. Measuring memory access latency for software objects in a NUMA system-on-chip architecture. In 2013 8th International Workshop on Reconfigurable and Communication-Centric Systems-on-Chip (ReCoSoC) (pp. 1-8). IEEE.

[21] Ogura, M., Harada, J., Kishida, M. and Yassine, A., 2019. Resource optimization of product development projects with time-varying dependency structure. Research in Engineering Design, 30(3), pp.435-452.

[22] Shi, Z. and Burns, A., 2008, April. Real-time communication analysis for on-chip networks with wormhole switching. In Second ACM/IEEE International Symposium on Networks-on-Chip (nocs 2008) (pp. 161-170). IEEE.

[23] Yun, H., Mancuso, R., Wu, Z.P. and Pellizzoni, R., 2014, April. PALLOC: DRAM bank-aware memory allocator for performance isolation on multicore platforms. In 2014 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS) (pp. 155-166). IEEE.

[24] Springer, R., Lowenthal, D.K., Rountree, B. and Freeh, V.W., 2006, March. Minimizing execution time in MPI programs on an energy-constrained, power-scalable cluster. In Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming (pp. 230-238).

[25] Yagna, K., Patil, O. and Mueller, F., 2016, June. Efficient and predictable group communication for manycore nocs. In International Conference on High Performance Computing (pp. 383-403). Springer, Cham.

[26] Pan, X., Gownivaripalli, Y.J. and Mueller, F., 2016, May. Tintmalloc: Reducing memory access divergence via controller-aware coloring. In 2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS) (pp. 363-372). IEEE.

[27] Carayannis, E.G. and Korres, G.M. eds., 2013. The Innovation Union in Europe: a socio-economic perspective on EU integration. Edward Elgar Publishing.

[28] Good practice: Mental border obstacles, online: http://www.interregeurope.eu/policylearning/good-practices/item/1955/mental-border-obstacles, 2020

[29] Skobelev, P., 2018, June. Towards autonomous AI systems for resource management: applications in industry and lessons learned. In International Conference on Practical Applications of Agents and Multi-Agent Systems (pp. 12-25). Springer, Cham.

[30] Furber, S., 2016. Interview with Steve Furber: The Designer of the ARM Chip Shares Lessons on Energy-Efficient Computing. ACM Queue, 8(2).

[31] Gartner, "Market Share: All Software Markets, Worldwide", 2018, published 2019, ID: G00385615, online: https://www.gartner.com/en/documents/3906672/market-share-all-software-markets-worldwide-2018