

Handbook on Summaries of European Exascale Projects

EXDCI Summaries: <https://exdci.eu/collaboration/fethpc> <https://exdci.eu/collaboration/exascale-projects> <https://exdci.eu/collaboration/coe>
http://cordis.europa.eu/search/result_en?q=contenttype=%27project%27%20AND%20/project/relations/associations/relatedCall/call/identifier=%27H2020-FETHPC-2014%27 <http://www.exascale-projects.eu/>

<h3>Software oriented projects</h3> <ul style="list-style-type: none"> AllScale: Nested Parallelism: An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism http://www.allscale.eu ANTAREX: AutoTuning and Adaptive approach for Energy efficient eXascale HPC systems http://www.antarex-project.eu CRESTA: Software co-design: Collaborative Research into Exascale Systemware, Tools & Applications http://www.cresta-project.eu EPIGRAM: Programming model: Exascale Programming Models http://www.epigram-project.eu EXA2CT: Solver, programming model: Exascale Algorithms and Advanced Computational Techniques http://www.exa2ct.eu EXAHYPE: Pde solver: An Exascale Hyperbolic PDE Engine http://exahype.eu INTERWIRE: Programming Model: Interoperability Towards Exascale (INTERWIRE) https://www.interwire-project.eu/partners READEX: Energy efficiency management: Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing http://www.readex.eu COMPAT: New mathematical & algorithmic approaches: Computing Patterns for High Performance Multiscale Computing http://www.compat-project.eu 	<h3>Domain oriented projects</h3> <ul style="list-style-type: none"> ESCAPE: Weather forecast: Energy-efficient Scalable Algorithms for weather Prediction at Exascale http://www.ecmwf.int/en/research/projects/escape EXA-FLOW: Fluid dynamics: Enabling Exascale Fluid Dynamics Simulations http://exaflow-project.eu EXCAPE: Machine learning: ExaScale Compound Activity Prediction Engine http://www.ecmwf.int/en/research/projects/escape NLAFET: Linear algebra: Parallel Numerical Linear Algebra for Future Extreme-Scale Systems http://www.nlafet.eu NESUS: Engineering: Network for Sustainable Ultrascale Computing http://www.nesus.eu
<h3>Architecture projects</h3> <ul style="list-style-type: none"> DEEP-ER: Accelerator: Dynamical Exascale Entry Platform – and it's Extended Reach http://www.deep-project.eu/www.deep-er.eu ECOSCALE: Heterogeneous computing: Energy-efficient Heterogeneous Computing at exaSCALE http://www.ecoscale.eu EXANEST: Interconnect, storage: European Exascale System Interconnect and Storage http://www.exanest.eu EXANODE: Chiplet, Interposer: European Exascale Processor Memory Node Design http://exanode.eu EXTRA: Reconfigurable architecture: Exploiting Exascale Technology with Reconfigurable Architectures https://www.extrahpc.eu GREEN FLASH: Real time control: Green Flash, energy efficient high performance computing for real-time science MANGO: Heterogeneous computing: Exploiting Manycore Architectures for Next-Generation HPC systems http://www.mango-project.eu MONTBLANC: ARM based HPC: European Approach Towards Energy Efficient HPC http://www.montblanc-project.eu NEXTGENIO: New memory hierarchy: Next Generation I/O for Exascale http://www.nextgenio.eu SAGE: Object oriented storage: SAGE http://www.sagestorage.eu 	<h3>Centres of Excellence</h3> <ul style="list-style-type: none"> AD: Material sciences: Materials design at the eXascale NOMAD: Material sciences: The Novel Materials Discovery Laboratory cam: Material sciences: An e-infrastructure for software, training and consultancy in simulation and modelling CoE GSS: Global Systems Science: Center of Excellence for Global Systems Science ESIWBC: Climate: Excellence in Simulation of Weather and Climate in Europe bioexcel: Bioscience: Centre of Excellence for Biomolecular Research E-C-E: Energy: Energy oriented Centre of Excellence for computer applications POP: Performance: Performance Optimisation and Productivity CompBioMed: Biomedicine: A Centre of Excellence in Computational Biomedicine

An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism



Main Achievements: Develop an API and programming environment for extreme-scale architectures that exploits nested recursive parallelism and which is based on a compiler and runtime system implementing multi-objective optimisations and resilience management.

Areas of Potential Collaboration: We see great potential to collaborate with compiler, runtime system, resilience and online performance monitoring developers and with application developers that can benefit by nested recursive parallelism for small to extreme-scale parallel architectures.

Extreme scale HPC systems impose significant challenges for developers aiming at obtaining applications efficiently utilising all available resources. In particular, the development of such applications is accompanied by the complex and labour-intensive task of managing parallel control flows, data dependencies and underlying hardware resources – each of these obligations constituting challenging problems on its own. The AllScale environment, the focus of this project, will provide a novel, sophisticated approach enabling the decoupling of the specification of parallelism from the associated management activities during program execution. Its foundation is a parallel programming model based on nested recursive parallelism, opening up the potential for a variety of compiler and runtime system based techniques adding to the capabilities of resulting applications. These include the (i) automated porting of application from small- to extreme scale architectures, (ii) the flexible tuning of the program execution to satisfy trade-offs among multiple objectives including execution time, energy and resource usage, (iii) the management of hardware resources and associated parameters (e.g. clock speed), (iv) the integration of resilience management measures to compensate for isolated hardware failures and (v) the possibility of online performance monitoring and analysis. All these services will be provided in an application independent, reusable fashion by a combination of sophisticated, modular, and customizable compiler and runtime system based solutions.

AllScale will boost the development productivity, portability, and runtime, energy, and resource efficiency of parallel applications targeting small to extreme scale parallel systems by leveraging the inherent advantages of nested recursive parallelism, and will be validated with applications from fluid dynamics, environmental hazard and space weather simulations provided by SME, industry and scientific partners.

AutoTuning and Adaptivity approach for Energy efficient eXascale HPC systems



Energy-efficient heterogeneous supercomputing architectures need to be coupled with a radically new software stack capable of exploiting the benefits offered by the heterogeneity at all the different levels (supercomputer, job, node) to meet the scalability and energy efficiency required by Exascale supercomputers.

ANTAREX will solve these challenging problems by proposing a disruptive holistic approach spanning all the decision layers composing the supercomputer software stack and exploiting effectively the full system capabilities (including heterogeneity and energy management). The main goal of the ANTAREX project is to provide a breakthrough approach to express application self-adaptivity at design-time and to runtime manage and autotune applications for green and heterogenous High Performance Computing (HPC) systems up to the Exascale level.

Collaborative Research Into Exascale Systemware, Tools and Application

For the past thirty years, the need for ever greater supercomputer performance has driven the development of many computing technologies which have subsequently been exploited in the mass market. Delivering an exaflop (or 10^{18} calculations per second) by the end of this decade is the challenge that the supercomputing community worldwide has set itself. The Collaborative Research into Exascale Systemware, Tools and Applications project (CRESTA) brings together four of Europe's leading supercomputing centres, with one of the world's major equipment vendors, two of Europe's leading programming tools providers and six application and problem owners to explore how the exaflop challenge can be met. CRESTA focuses on the use of six applications with exascale potential and uses them as co-design vehicles to develop: the development environment, algorithms and libraries, user tools, and the underpinning and cross-cutting technologies required to support the execution of applications at the exascale. The applications represented in CRESTA have been chosen as a representative sample from across the supercomputing domain including: biomolecular systems, fusion energy, the virtual physiological human, numerical weather prediction and engineering.



No one organisation, be they a hardware or software vendor or service provider can deliver the necessary range of technological innovations required to enable computing at the exascale. This is recognised through the on-going work of the International Exascale Software Project and, in Europe, the European Exascale Software Initiative. CRESTA will actively engage with European and International collaborative activities to ensure that Europe plays its full role worldwide. Over its 39 month duration the project will deliver key, exploitable technologies that will allow the co-design applications to successfully execute on multi-petaflop systems in preparation for the first exascale systems towards the end of this decade.

Computing Patterns for High Performance Multiscale Computing



Multiscale phenomena are ubiquitous and they are the key to understanding the complexity of our world. Despite the significant progress achieved through computer simulations over the last decades, we are still limited in our capability to accurately and reliably simulate hierarchies of interacting multiscale physical processes that span a wide range of time and length scales, thus quickly reaching the limits of contemporary high performance computing at the tera- and petascale. Exascale supercomputers promise to lift this limitation, and in this project we will develop multiscale computing algorithms capable of producing high-fidelity scientific results and scalable to exascale computing systems. Our main objective is to develop generic and reusable High Performance Multiscale Computing algorithms that will address the exascale challenges posed by heterogeneous architectures and will enable us to run multiscale applications with extreme data requirements while achieving scalability, robustness, resiliency, and energy efficiency. Our approach is based on generic multiscale computing patterns that allow us to implement customized algorithms to optimise load balancing, data handling, fault tolerance and energy consumption under generic exascale application scenarios. We will realise an experimental execution environment on our pan-European facility, which will be used to measure performance characteristics and develop models that can provide reliable performance predictions for emerging and future exascale architectures. The viability of our approach will be demonstrated by implementing nine grand challenge applications which are exascale-ready and pave the road to unprecedented scientific discoveries. Our ambition is to establish new standards for multiscale computing at exascale, and provision a robust and reliable software technology stack that empowers multiscale modellers to transform computer simulations into predictive science.

Dynamical Exascale Entry Platform – and it's Extended Reach

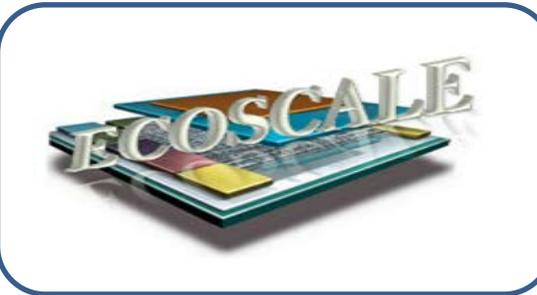


Main Achievements: The DEEP and DEEP-ER consortium has developed a novel, Exascale-enabling supercomputing architecture with a matching software stack and a set of optimized grand-challenge simulation applications.

Areas of Potential Collaboration: For DEEP and DEEP-ER it would be very valuable to get feedback from a diverse group of application scientists on how to use and to benefit from the Cluster-Booster architecture.

The proposed project DEEP-ER (DEEP-Extended Reach) addresses two significant Exascale challenges: the growing gap between I/O bandwidth and compute speed, and the need to significantly improve system resiliency. DEEP-ER will extend the Cluster-Booster architecture of the Dynamical Exascale Entry Platform (DEEP) project by a highly scalable I/O system and will implement an efficient mechanism to recover application tasks that fail due to hardware errors. The project will leverage new memory technology to provide increased performance and power efficiency. As a result, I/O parts of HPC codes will run faster and scale up better HPC applications will be able to profit from checkpointing and task restart on large systems reducing overhead seen today. Systems that use the DEEP-ER results can run more applications increasing scientific throughput, and the loss of computational work through system failures will be substantially reduced.

DEEP-ER will build a prototype with the second generation Intel® Xeon Phi processor, a uniform high-speed interconnect across Cluster and Booster, non-volatile memory on the compute nodes, and network attached memory providing high-speed shared memory access. A highly scalable and efficient I/O system based on the BeeGFS file system from Fraunhofer-ITWM will support I/O intensive applications, using optimised I/O middleware SIONlib and E10. A multi-level checkpoint scheme will exploit scalable I/O and fast, non-volatile memory close to the nodes to reduce the overhead of saving state for long-running tasks. The OmpSs based DEEP programming model will govern the creation of checkpoints and restart failed tasks from the beginning or recover saved state depending on their granularity. Seven important HPC applications will be optimised demonstrating the usability, performance and resiliency of the DEEP-ER Prototype. The applications come from

<p>Energy-efficient Heterogeneous Computing at exascale</p> 	<p>different scientific and engineering areas and represent requirements of simulation-based and data-intensive HPC codes.</p> <p>In order to reach exascale performance current HPC servers need to be improved. Simple scaling is not a feasible solution due to the increasing utility costs and power consumption limitations. Apart from improvements in implementation technology, what is needed is to refine the HPC application development as well as the architecture of the future HPC systems. ECOSCALE tackles this challenge by proposing a scalable programming environment and hardware architecture tailored to the characteristics and trends of current and future HPC applications, reducing significantly the data traffic as well as the energy consumption and delays. We first propose a novel heterogeneous energy-efficient hierarchical architecture and a hybrid MPI+OpenCL programming environment and runtime system. The proposed architecture, programming model and runtime system follows a hierarchical approach where the system is partitioned into multiple autonomous Workers (i.e. compute nodes). Workers are interconnected in a tree-like structure in order to form larger Partitioned Global Address Space (PGAS) partitions, which are further hierarchically interconnected via an MPI protocol.</p> <p>Secondly, to further increase the energy efficiency of the system as well as its resilience, the Workers will employ reconfigurable accelerators that can perform coherent memory accesses in the virtual address space utilizing an IOMMU. The ECOSCALE architecture will support shared partitioned reconfigurable resources accessed by any Worker in a PGAS partition, and, more importantly, automated hardware synthesis of these resources from an OpenCL-based programming model. We follow a co-design approach that spans a scalable HPC hardware platform, a middleware layer, a programming and a runtime environment as well as a high-level design environment for mapping applications onto the system. A proof of concept prototype and a simulator will be built in order to run two real-world HPC applications and several benchmarks.</p>
<p>Exascale Programming Models</p> 	<p>Exascale computing power will likely be reached in the next decade. While the precise system architectures are still evolving, one can safely assume that they will be largely based on deep hierarchies of multicore CPUs with similarly-deep memory hierarchies, potentially also supported by accelerators. New and disruptive programming models are needed to allow applications to run efficiently at large scale on these platforms. The Message Passing Interface (MPI) has emerged as the de-facto standard for parallel programming on current petascale machines; but Partitioned Global Address Space (PGAS) languages and libraries are increasingly being considered as alternatives or complements to MPI. However, both approaches have severe problems that will prevent them reaching exascale performance.</p> <p>The aim of this proposal is to prepare Message Passing (MP) and PGAS programming models for exascale systems by fundamentally addressing their main current limitations. We will introduce new disruptive concepts to fill the technological gap between the petascale and exascale era in two ways:</p> <ul style="list-style-type: none"> ▣ First, innovative algorithms will be used in both MP and PGAS, specifically to provide fast collective communication in both MP and PGAS, to decrease the memory consumption in MP, to enable fast synchronization in PGAS, to provide fault tolerance mechanisms in PGAS, and potential strategies for fault tolerance in MP. ▣ Second, we will combine the best features of MP and PGAS by developing an MP interface using a PGAS library as communication substrate. The concepts developed will be tested and guided by two applications in the engineering and space weather domains chosen from the suite of codes in current EC exascale projects. By providing prototype implementations for both MP and PGAS concepts we will contribute significantly to advancement in programming models and interfaces for ultra-scale computing systems, and provide stimuli for European research in this vital area.
<p>Energy-efficient Scalable Algorithms for weather Prediction at Exascale</p>	<p>"ESCAPE will develop world-class, extreme-scale computing capabilities for European operational numerical weather prediction (NWP) and future climate models. The biggest challenge for state-of-the-art NWP arises from the need to simulate complex physical phenomena within tight production schedules. Existing extreme-scale application software of weather and climate services is ill-equipped to adapt to the rapidly evolving hardware. This is exacerbated by other drivers for hardware development, with processor arrangements not necessarily optimal for weather and climate simulations. ESCAPE will redress this imbalance through innovation actions that fundamentally reform Earth-system modelling.</p> <p>ESCAPE addresses the ETP4HPC SRA 'Energy and resiliency' priority topic, developing a holistic understanding of energy-efficiency for extreme-scale applications using heterogeneous architectures, accelerators and special compute units. The three key reasons why this proposal will provide the necessary means to take a huge step forward in weather and climate modelling as well as interdisciplinary research on energy-efficient high-performance computing are:</p>



- 1) Defining and encapsulating the fundamental algorithmic building blocks ("Weather & Climate Dwarfs") underlying weather and climate services. This is the pre-requisite for any subsequent co-design, optimization, and adaptation efforts.
- 2) Combining ground-breaking frontier research on algorithm development for use in extreme-scale, high-performance computing applications, minimizing time- and cost-to-solution.
- 3) Synthesizing the complementary skills of all project partners. This includes ECMWF, the world leader in global NWP together with leading European regional forecasting consortia, teaming up with excellent university research and experienced high-performance computing centres, two world-leading hardware companies, and one European start-up SME, providing entirely new knowledge and technology to the field."

Exascale Algorithms and Advanced Computational Techniques



Numerical simulation is a crucial part of science and industry in Europe. The advancement of simulation as a discipline relies on increasingly compute intensive models that require more computational resources to run. This is the driver for the evolution to exascale. Due to limits in the increase in single processor performance, exascale machines will rely on massive parallelism on and off chip, with a complex hierarchy of resources. The large number of components and the machine complexity introduce severe problems for reliability and programmability. The former of these will require novel fault-aware algorithms and support software. In addition, the scale of the numerical models exacerbates the difficulties by making the use of more complex simulation algorithms necessary, for numerical stability reasons. A key example of this is increased reliance on solvers. Such solvers require global communication, which impacts scalability, and are often used with preconditioners, increasing complexity again. Unless there is a major rethink of the design of solver algorithms, their components and software structure, a large class of important numerical simulations will not scale beyond petascale. This in turn will hold back the development of European science and industry which will fail to reap the benefits from exascale.

The EXA2CT project brings together experts at the cutting edge of the development of solvers, related algorithmic techniques, and HPC software architects for programming models and communication. It will take a revolutionary approach to exascale solvers and programming models, rather than the incremental approach of other projects. We will produce modular open source proto-applications that demonstrate the algorithms and programming techniques developed in the project, to help boot-strap the creation of genuine exascale codes.

Enabling Exascale Fluid Dynamics Simulations



Main Achievements: Algorithmic improvements for increased performance, scalability and exascale readiness of major high order, open source computational fluid dynamics codes.

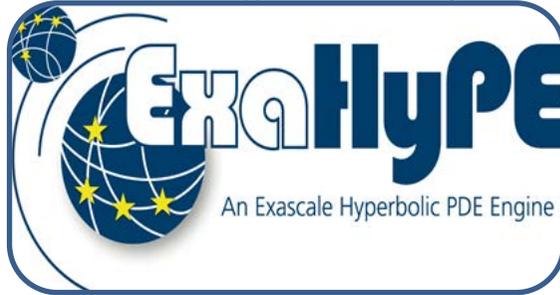
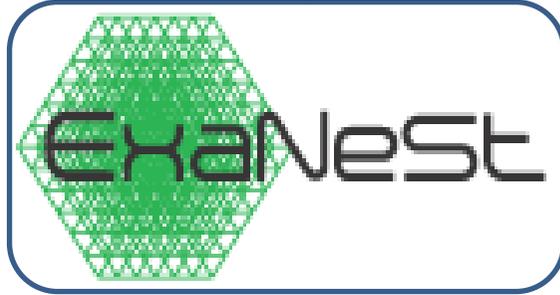
Areas of Potential Collaboration: Extreme scale computational fluid dynamics simulations; Algorithmics improvements towards exascale.

We are surrounded by moving fluids (gases and liquids), be it during breathing or the blood flowing in arteries; the flow around cars, ships, and airplanes; the changes in cloud formations or the plankton transport in oceans; even the formation of stars and galaxies are closely modeled as phenomena in fluid dynamics. Fluid Dynamics (FD) simulations provide a powerful tool for the analysis of such fluid flows and are an essential element of many industrial and academic problems.

The complexities and nature of fluid flows, often combined with problems set in open domains, implies that the resources needed to computationally model problems of industrial and academic relevance is virtually unbounded. FD simulations therefore are a natural driver for exascale computing and have the potential for substantial societal impact, like reduced energy consumption, alternative sources of energy, improved health care, and improved climate models.

The main goal of this project is to address algorithmic challenges to enable the use of accurate simulation models in exascale environments. Driven by problems of practical engineering interest we focus on important simulation aspects including:

- error control and adaptive mesh refinement in complex computational domains,
- resilience and fault tolerance in complex simulations
- heterogeneous modeling
- evaluation of energy efficiency in solver design
- parallel input/output and in-situ compression for extreme data.

	<p>The algorithms developed by the project will be prototyped in major open-source simulation packages in a co-design fashion, exploiting software engineering techniques for exascale. We are building directly on the results of previous exascale projects (CRESTA, EPiGRAM, etc.) and will exploit advanced and novel parallelism features required for emerging exascale architectures. The results will be validated in a number of pilot applications of concrete practical importance in close collaboration with industrial partners.</p>
<p>An Exascale Hyperbolic PDE Engine</p> 	<p>Main Achievements: We have implemented the ADER-DG numerical approach on adaptive spacetree meshes within a first prototype of our engine for solving hyperbolic systems of PDE with high performance on future exascale supercomputers.</p> <p>Areas of Potential Collaboration: ExaHyPE welcomes potential users of the hyperbolic PDE engine to discuss and prototypically implement possible applications, but also seeks collaboration with providers of emerging supercomputing technology, including actual supercomputers in the same way as programming environments.</p> <p>We are surrounded by moving fluids (gases and liquids), be it during breathing or the blood flowing in arteries; the flow around cars, ships, and airplanes; the changes in cloud formations or the plankton transport in oceans; even the formation of stars and galaxies are closely modeled as phenomena in fluid dynamics. Fluid Dynamics (FD) simulations provide a powerful tool for the analysis of such fluid flows and are an essential element of many industrial and academic problems.</p> <p>The complexities and nature of fluid flows, often combined with problems set in open domains, implies that the resources needed to computationally model problems of industrial and academic relevance is virtually unbounded. FD simulations therefore are a natural driver for exascale computing and have the potential for substantial societal impact, like reduced energy consumption, alternative sources of energy, improved health care, and improved climate models.</p> <p>The main goal of this project is to address algorithmic challenges to enable the use of accurate simulation models in exascale environments. Driven by problems of practical engineering interest we focus on important simulation aspects including:</p> <ul style="list-style-type: none"> • error control and adaptive mesh refinement in complex computational domains, • resilience and fault tolerance in complex simulations • heterogeneous modeling • evaluation of energy efficiency in solver design • parallel input/output and in-situ compression for extreme data. <p>The algorithms developed by the project will be prototyped in major open-source simulation packages in a co-design fashion, exploiting software engineering techniques for exascale. We are building directly on the results of previous exascale projects (CRESTA, EPiGRAM, etc.) and will exploit advanced and novel parallelism features required for emerging exascale architectures. The results will be validated in a number of pilot applications of concrete practical importance in close collaboration with industrial partners.</p>
<p>European Exascale System Interconnect and Storage</p> 	<p>ExaNeSt will develop, evaluate, and prototype the physical platform and architectural solution for a unified Communication and Storage Interconnect and the physical rack and environmental structures required to deliver European Exascale Systems. The consortium brings technology, skills, and knowledge across the entire value chain from computing IP to packaging and system deployment; and from operating systems, storage, and communication to HPC with big data management, algorithms, applications, and frameworks. Building on a decade of advanced R&D, ExaNeSt will deliver the solution that can support exascale deployment in the follow-up industrial commercialization phases. Using direction from the ETP4HPC roadmap and soon-available high density and efficiency compute, we will model, simulate, and validate through prototype, a system with:</p> <ol style="list-style-type: none"> 1. High throughput, low latency connectivity, suitable for exascale-level compute, their storage, and I/O, with congestion mitigation, QoS guarantees, and resilience. 2. Support for distributed storage located with the compute elements providing low latency that non-volatile memories require, while reducing energy, complexity, and costs. 3. Support for task-to-data sw locality models to ensure minimum data communication energy overheads and property maintenance in databases. 4. Hyper-density system integration scheme that will develop a modular, commercial, European-sourced advanced cooling system for exascale

	<p>in ~200 racks while maintaining reliability and cost of ownership.</p> <p>5. The platform management scheme for big-data I/O to this resilient, unified distributed storage compute architecture.</p> <p>6. Demonstrate the applicability of the platform for the complete spectrum of Big Data applications, e.g. from HPC simulations to Business Intelligence support.</p> <p>All aspects will be steered and validated with the first-hand experience of HPC applications and experts, through kernel turning and subsequent data management and application analysis.</p>
<p>European Exascale Processor Memory Node Design</p> 	<p>Main Achievements: ExaNoDe designs a high energy efficient and highly integrated heterogeneous compute node enabling Pre-Exascale level computing, mixing low-power processors, co-processors and advanced hardware integration technologies with the novel UNIMEM Global Address Space memory system.</p> <p>Areas of Potential Collaboration: By aiming at creating an extremely dense and high-energy efficient compute node, ExaNoDe could create links with either other research projects or companies active in hardware design and manufacturing in both USA, Japan and other countries.</p> <p>ExaNoDe will investigate, develop and pilot (technology readiness level 7) a highly efficient, highly integrated, multi-way, high-performance, heterogeneous compute element aimed towards exascale computing and demonstrated using hardware-emulated interconnect. It will build on multiple European initiatives for scalable computing, utilizing low-power processors and advanced nanotechnologies. ExaNoDe will draw heavily on the Unimem memory and system design paradigm defined within the EUROSERVER FP7 project, providing low-latency, high-bandwidth and resilient memory access, scalable to Exabyte levels.</p> <p>The ExaNoDe compute element aims towards exascale compute goals through:</p> <ul style="list-style-type: none"> • Integration of the most advanced low-power processors and accelerators across scalar, SIMD, GPGPU and FPGA processing elements supported by research and innovation in the deployment of associated nanotechnologies and in the mechanical requirements to enable the development of a high-density, high-performance integrated compute element with advanced thermal characteristics and connectivity to the next generation of system interconnect and storage; • Undertaking essential research to ensure the ExaNoDe compute element provides necessary support of HPC applications including I/O and storage virtualization techniques, operating system and semantically aware runtime capabilities and PGAS, OpenMP and MPI paradigms; • The development of an instantiation of a hardware emulation of interconnect to enable the evaluation of Unimem for the deployment of multiple compute elements and the evaluation, tuning and analysis of HPC mini-apps. <p>Each aspect of ExaNoDE is aligned with the goals of the ETP4HPC. The work will be steered by first-hand experience and analysis of high-performance applications, their requirements and the tuning of their kernels.</p>
<p>Exascale Compound Activity Prediction Engine</p> 	<p>Main Achievements: The ExCAPE project has prepared machine learning data sets for the pharmaceutical industry, made them accessible on the Czech supercomputing cluster, and applied machine learning techniques using the cluster.</p> <p>Areas of Potential Collaboration: ExCAPE would like international collaboration in the areas of large (10M to 100M training point) multi-task, sparse-evidence and sparse feature vector machine learning algorithms that can be run on supercomputer hardware, especially in the area of compound activity prediction for the pharmaceutical industry.</p> <p>Scalable machine learning of complex models on extreme data will be an important industrial application of exascale computers. In this project, we take the example of predicting compound bioactivity for the pharmaceutical industry, an important sector for Europe for employment, income, and solving the problems of an ageing society. Small scale approaches to machine learning have already been trialed and show great promise to reduce empirical testing costs by acting as a virtual screen to filter out tests unlikely to work. However, it is not yet possible to use all available data to make the best possible models, as algorithms (and their implementations) capable of learning the best models do not scale to such sizes and heterogeneity of input data. There are also further challenges including imbalanced data, confidence estimation, data standards model quality and feature diversity.</p> <p>The ExCAPE project aims to solve these problems by producing state of the art scalable algorithms and implementations thereof suitable for running on future Exascale machines. These approaches will scale programs for complex pharmaceutical workloads to input data sets at</p>

	<p>industry scale. The programs will be targeted at exascale platforms by using a mix of HPC programming techniques, advanced platform simulation for tuning and suitable accelerators.</p>
<p>Exploiting exascale Technology with Reconfigurable Architectures</p> 	<p>To handle the stringent performance requirements of future exascale High Performance Computing (HPC) applications, HPC systems need ultra-efficient heterogeneous compute nodes. To reduce power and increase performance, such compute nodes will require reconfiguration as an intrinsic feature, so that specific HPC application features can be optimally accelerated at all times, even if they regularly change over time. In the EXTRA project, we create a new and flexible exploration platform for developing reconfigurable architectures, design tools and HPC applications with run-time reconfiguration built-in from the start. The idea is to enable the efficient co-design and joint optimization of architecture, tools, applications, and reconfiguration technology in order to prepare for the necessary HPC hardware nodes of the future. The project EXTRA covers the complete chain from architecture up to the application:</p> <ul style="list-style-type: none"> • More coarse-grain reconfigurable architectures that allow reconfiguration on higher functionality levels and therefore provide much faster reconfiguration than at the bit level. • The development of just-in time synthesis tools that are optimized for fast (but still efficient) re-synthesis of application phases to new, specialized implementations through reconfiguration. • The optimization of applications that maximally exploit reconfiguration. • Suggestions for improvements to reconfigurable technologies to enable the proposed reconfiguration of the architectures. <p>In conclusion, EXTRA focuses on the fundamental building blocks for run-time reconfigurable exascale HPC systems: new reconfigurable architectures with very low reconfiguration overhead, new tools that truly take reconfiguration as a design concept, and applications that are tuned to maximally exploit run-time reconfiguration techniques. Our goal is to provide the European platform for run-time reconfiguration to maintain Europe’s competitive edge and leadership in run-time reconfigurable computing.</p>
<p>Green Flash, energy efficient high performance computing for real-time science</p> 	<p>The main goal of Green Flash is to design and build a prototype for a Real-Time Controller (RTC) targeting the European Extremely Large Telescope (E-ELT) Adaptive Optics (AO) instrumentation. The E-ELT is a 39m diameter telescope to see first light in the early 2020s. To build this critical component of the telescope operations, the astronomical community is facing technical challenges, emerging from the combination of high data transfer bandwidth, low latency and high throughput requirements, similar to the identified critical barriers on the road to Exascale. With Green Flash, we will propose technical solutions, assess these enabling technologies through prototyping and assemble a full scale demonstrator to be validated with a simulator and tested on sky.</p> <p>With this R&D program we aim at feeding the E-ELT AO systems preliminary design studies, led by the selected first-light instruments consortia, with technological validations supporting the designs of their RTC modules. Our strategy is based on a strong interaction between academic and industrial partners. Components specifications and system requirements are derived from the AO application. Industrial partners lead the development of enabling technologies aiming at innovative tailored solutions with potential wide application range. The academic partners provide the missing links in the ecosystem, targeting their application with mainstream solutions. This increases both the value and market opportunities of the developed products. A prototype harboring all the features is used to assess the performance. It also provides the proof of concept for a resilient modular solution to equip a large scale European scientific facility, while containing the development cost by providing opportunities for return on investment.</p>
<p>Programming Model Interoperability towards Exascale</p>	<p>Main Achievements: INTERTWinE has established a Europe-wide programme of advanced training on parallel and interoperable programming for extreme scale, supported by Best Practice Guides that distill the expertise of leading European computational scientists.</p> <p>Areas of Potential Collaboration: INTERTWinE has established, and is addressing, the need for a model-agnostic interoperability forum at which representatives for key parallel APIs and scientific libraries can work together to solve the exascale programming problem.</p> <p>This project addresses the problem of programming model design and implementation for the Exascale. The first Exascale computers will be very highly parallel systems, consisting of a hierarchy of architectural levels. To program such systems effectively and portably, programming APIs with efficient and robust implementations must be ready in the appropriate timescale. A single, “silver bullet” API which addresses all the</p>



architectural levels does not exist and seems very unlikely to emerge soon enough. We must therefore expect that using combinations of different APIs at different system levels will be the only practical solution in the short to medium term. Although there remains room for improvement in individual programming models and their implementations, the main challenges lie in interoperability between APIs. It is this interoperability, both at the specification level and at the implementation level, which this project seeks to address and to further the state of the art. INTERTWinE brings together the principal European organisations driving the evolution of programming models and their implementations. The project will focus on seven key programming APIs: MPI, GASPI, OpenMP, OmpSs, StarPU, QUARK and PaRSEC, each of which has a project partner with extensive experience in API design and implementation. Interoperability requirements and evaluation of implementations will be driven by a set of kernels and applications, each of which has a project partner with a major role in their development. The project will implement a co- design cycle, by feeding back advances in API design and implementation into the applications and kernels, thereby driving new requirements and hence further advances.

Exploring Many-core Architectures for Next-Generation HPC systems



MANGO targets to achieve extreme resource efficiency in future QoS-sensitive HPC through ambitious cross-boundary architecture exploration for performance/power/predictability (PPP) based on the definition of new-generation high-performance, power-efficient, heterogeneous architectures with native mechanisms for isolation and quality-of-service, and an innovative two-phase passive cooling system. Its disruptive approach will involve many interrelated mechanisms at various architectural levels, including heterogeneous computing cores, memory architectures, interconnects, run-time resource management, power monitoring and cooling, to the programming models. The system architecture will be inherently heterogeneous as an enabler for efficiency and application-based customization, where general-purpose compute nodes (GN) are intertwined with heterogeneous acceleration nodes (HN), linked by an across-boundary homogeneous interconnect. It will provide guarantees for predictability, bandwidth and latency for the whole HN node infrastructure, allowing dynamic adaptation to applications. MANGO will develop a toolset for PPP and explore holistic pro-active thermal and power management for energy optimization including chip, board and rack cooling levels, creating a hitherto inexistent link between HW and SW effects at all layers. Project will build an effective large-scale emulation platform. The architecture will be validated through noticeable examples of application with QoS and high-performance requirements. Ultimately, the combined interplay of the multi-level innovative solutions brought by MANGO will result in a new positioning in the PPP space, ensuring sustainable performance as high as 100 PFLOPS for the realistic levels of power consumption (<15MWatt) delivered to QoS-sensitive applications in large-scale capacity computing scenarios providing essential building blocks at the architectural level enabling the full realization of the ETP4HPC strategic research agenda.

European Approach Towards Energy Efficient HPC



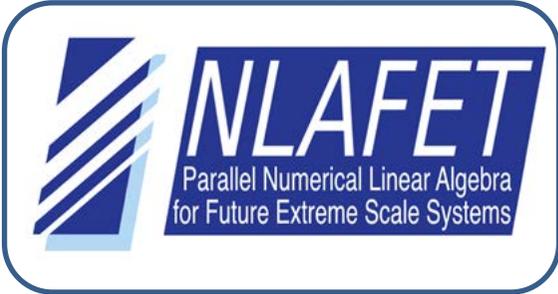
Main Achievements: As the project offers a fairly complete ARM-based ecosystem of hardware/software infrastructure for HPC, international collaborations can involve tests of production codes on the Mont-Blanc platforms, development and test of tools for power monitoring, joint resiliency studies.

Areas of Potential Collaboration: Also, collaborations can be at dissemination level, with the involvement of the project in events, workshops, training or conferences.

This project addresses the problem of programming model design and implementation for the Exascale. The first Exascale computers will be very highly parallel systems, consisting of a hierarchy of architectural levels. To program such systems effectively and portably, programming APIs with efficient and robust implementations must be ready in the appropriate timescale. A single, “silver bullet” API which addresses all the architectural levels does not exist and seems very unlikely to emerge soon enough. We must therefore expect that using combinations of different APIs at different system levels will be the only practical solution in the short to medium term. Although there remains room for improvement in individual programming models and their implementations, the main challenges lie in interoperability between APIs. It is this interoperability, both at the specification level and at the implementation level, which this project seeks to address and to further the state of the art. INTERTWinE brings together the principal European organisations driving the evolution of programming models and their implementations. The project will focus on seven key programming APIs: MPI, GASPI, OpenMP, OmpSs, StarPU, QUARK and PaRSEC, each of which has a project partner with extensive experience in API design and implementation. Interoperability requirements, and evaluation of implementations will be driven by a set of kernels and applications, each of which has a project partner with a major role in their development.

	<p>The project will implement a co- design cycle, by feeding back advances in API design and implementation into the applications and kernels, thereby driving new requirements and hence further advances.</p> <p>Among the most relevant results of the Mont-Blanc project are:</p> <p>a) the deployment of the world's first ARM-based HPC cluster, a prototype that largely contributed to validate the concept of using ARM technology for HPC, and which, together with various ARM-based prototypes, allowed the project to substantially contribute to the development of the HPC system software stack for ARM-based platforms</p> <p>b) a strong involvement in the realization of a complete software ecosystem allowing the execution of production HPC applications on ARM based clusters.</p>
<p>Network for Sustainable Ultrascale Computing</p> 	<p>Main Achievements: Development of a major research network with more than 80 institutions belonging to 45 countries.</p> <p>Areas of Potential Collaboration: Programming environment, Data-intensive computing, HPC and Energy HPC and Big Data</p> <p>NESUS is the acronym of the COST Action IC1305. The Action is a research network composed by almost 80 institutions from 45 countries (most EU, Russia, USA, India, Australia and Colombia) created to research on the evolution of today's large-scale system, not only in HPC world, but also in HTC, and the challenges arising that could stop adoption of those systems in the future.</p> <p>The project has created a strong network of researchers that are cooperating in shared publications, workshops and applications related to the topics of the action. A first research roadmap will be issued by the end of 2016. It also promotes a winter school every year to share experiences and to train early researchers.</p> <p>NESUS strongly promotes international cooperation of the Action's member, as cooperation is one of its goals. The Action itself is a cooperation tool, but it could be enhanced by means of researchers exchanges, development of joint applications, participation in the winter school, and organization of scientific events related to the Action.</p>
<p>Next Generation I/O for Exascale</p> 	<p>Main Achievements: To date, our main achievements are the definition of a new HPC architecture (hardware and software) that can make use of NVDIMMs; the project has now entered its implementation phase.</p> <p>Areas of Potential Collaboration: Partners that have IO requirements that currently prevent them from using traditional HPC systems.</p> <p>The overall objective of the Next Generation I/O Project (NEXTGenIO) is to design and prototype a new, scalable, high-performance, energy efficient computing platform designed to address the challenge of delivering scalable I/O performance to applications at the Exascale. It will achieve this using highly innovative, non-volatile, dual in-line memory modules (NV-DIMMs). These hardware and systemware developments will be coupled to a co-design approach driven by the needs of some of today's most demanding HPC applications. By meeting this overall objective, NEXTGenIO will solve a key part of the Exascale challenge and enable HPC and Big Data applications to overcome the limitations of today's HPC I/O subsystems.</p> <p>Today most high-end HPC systems employ data storage separate from the main system and the I/O subsystem often struggles to deal with the degree of parallelism present. As we move into the domain of extreme parallelism at the Exascale we need to address I/O if such systems are to deliver appropriate performance and efficiency for their application user communities.</p> <p>The NEXTGenIO project will explore the use of NV-DIMMs and associated systemware developments through a co-design process with three 'end-user' partners: a high-end academic HPC service provider, a numerical weather forecasting service provider and a commercial on-demand HPC service provider. These partners will develop a set of I/O workload simulators to allow quantitative improvements in I/O performance to be directly measured on the new system in a variety of research configurations. Systemware software developed in the project will include performance analysis tools, improved job schedulers that take into account data locality and energy efficiency, optimised programming models, and APIs and drivers for optimal use of the new I/O hierarchy.</p> <p>The project will deliver immediately exploitable hardware and software results and show how to deliver high performance I/O at the Exascale. To date, our main achievements are the definition of a new HPC architecture (hardware and software) that can make use of NVDIMMs; the project has now entered its implementation phase.</p>

Parallel Numerical Linear Algebra for Future Extreme-Scale Systems



The NLA FET proposal is a direct response to the demands for new mathematical and algorithmic approaches for applications on extreme scale systems, as identified in the FETHPC work programme and call. This project will enable a radical improvement in the performance and scalability of a wide range of real-world applications relying on linear algebra software, by developing novel architecture-aware algorithms and software libraries, and the supporting runtime capabilities to achieve scalable performance and resilience on heterogeneous architectures. The focus is on a critical set of fundamental linear algebra operations including direct and iterative solvers for dense and sparse linear systems of equations and eigenvalue problems. Achieving this requires a co-design effort due to the characteristics and overwhelming complexity and immense scale of such systems. Recognized experts in algorithm design and theory, parallelism, and auto-tuning will work together to explore and negotiate the necessary tradeoffs.

The main research objectives are:

- (i) development of novel algorithms that expose as much parallelism as possible, exploit heterogeneity, avoid communication bottlenecks, respond to escalating fault rates, and help meet emerging power constraints;
- (ii) exploration of advanced scheduling strategies and runtime systems focusing on the extreme scale and strong scalability in multi/many-core and hybrid environments;
- (iii) design and evaluation of novel strategies and software support for both offline and online auto-tuning.

The validation and dissemination of results will be done by integrating new software solutions into challenging scientific applications in materials science, power systems, study of energy solutions, and data analysis in astrophysics. The deliverables also include a sustainable set of methods and tools for cross-cutting issues such as scheduling, auto-tuning, and algorithm-based fault tolerance packaged into open-source library modules.

Numerical Methods and Tools for Key Exascale



Exaflop computing will make a considerable impact on several areas of engineering and applied sciences, where current high-end computing capabilities are deemed grossly insufficient. The overall aim of the NUMEXAS project is to develop, implement and validate the next generation of numerical methods to be run under exascale computing architectures. This will be done by implementing a new paradigm for the development of advanced numerical methods to really exploit the intrinsic capabilities of the future exascale computing infrastructures. The project covers RTD activities along the complete simulation pipeline: parallel pre-processing and grid generation, new numerical methods for parallel structured/unstructured multidisciplinary field solvers of high order, optimum design parallel solvers considering uncertainties and parallel in-solver visualization and feature extraction. The new numerical methods and software will be validated for a selected number of exascale size problems in engineering and applied sciences in state-of-the-art high performance computing platforms.

The main outcome of NUMEXAS will be a new set of numerical methods and codes that will allow industry, government and academia to solve exascale-class problems in engineering and applied sciences in the next generation of exaflop computers with the efficiency and ease of use as today's state-of-the-art codes. The consortium has a well-balanced distribution of institutions specialized in the development of numerical methods to solve grand challenge engineering and scientific problems (CIMNE, LUH-IKM and NTUA) and institutions hosting HPC facilities and supercomputing infrastructures (CESCA and LUH-HLRN). The partnership is completed with QUANTECH, an SME specialized in the development and marketing of simulation software for industrial forming processes.

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

High Performance Computing (HPC) has become a major instrument for many scientific and industrial fields to generate new insights and product developments. There is a continuous demand for growing compute power, leading to a constant increase in system size and complexity. Efficiently utilizing the resources provided on Exascale systems will be a challenging task, potentially causing a large amount of underutilized resources and wasted energy. Parameters for adjusting the system to application requirements exist both on the hardware and on the system software level but are mostly unused today. Moreover, accelerators and co-processors offer a significant performance improvement at the cost of increased overhead, e.g., for data-transfers.

While HPC applications are usually highly compute intensive, they also exhibit a large degree of dynamic behavior, e.g., the alternation between communication phases and compute kernels. Manually detecting and leveraging this dynamism to improve energy-efficiency is a tedious task that is commonly neglected by developers. However, using an automatic optimization approach, application dynamism can be detected at design-time and used to generate optimized system configurations. A light-weight run-time system will then detect this dynamic behavior in

READEX

Runtime Exploitation of Application Dynamism
for Energy-efficient eXascale computing

production and switch parameter configurations if beneficial for the performance and energy-efficiency of the application. The READEX project will develop an integrated tool-suite and the READEX Programming Paradigm to exploit application domain knowledge, together achieving an improvement in energy-efficiency of up to 22.5%.

Driven by a consortium of European experts from academia, HPC resource providers, and industry, the READEX project will develop a tool-aided methodology to exploit the dynamic behavior of applications to achieve improved energy-efficiency and performance. The developed tool-suite will be efficient and scalable to support current and future extreme scale systems.

Percipient Storage for Exascale Data Centric Computing



Main Achievements: The project has completed the first phase of requirements gathering activity from the use cases for building the SAGE storage platform and completed the architecture and design of the key Software & Hardware components of SAGE in the first year.

Areas of Potential Collaboration: Further/continued collaboration on the co-design of the storage platform (incl. its API) and supporting ecosystem for exascale and beyond by better understanding future extreme computing/data workflows as needed by the international community.

Worldwide data volumes are exploding and islands of storage remote from compute will not scale. We will demonstrate the first instance of intelligent data storage, uniting data processing and storage as two sides of the same rich computational model. This will enable sophisticated, intention-aware data processing to be integrated within a storage systems infrastructure, combined with the potential for Exabyte scale deployment in future generations of extreme scale HPC systems.

Enabling only the salient data to flow in and out of compute nodes, from a sea of devices spanning next generation solid state to low performance disc we enable a vision of a new model of highly efficient and effective HPC and Big Data demonstrated through the SAGE project. Objectives:

- Provide a next-generation multi-tiered object-based data storage system (hardware and enabling software) supporting future-generation multi-tier persistent storage media supporting integral computational capability, within a hierarchy.
- Significantly improve overall scientific output through advancements in systemic data access performance and drastically reduced data movements.
- Provides a roadmap of technologies supporting data access for both Exascale/Exabyte and High Performance Data Analytics.
- Provide programming models, access methods and support tools validating their usability, including 'Big-Data' access and analysis methods
- Co-Designing and validating on a smaller representative system with earth sciences, meteorology, clean energy, and physics communities
- Projecting suitability for extreme scaling through simulation based on evaluation results.

Call Alignment: We address storage data access with optimised systems for converged Big Data and HPC use, in a co-design process with scientific partners and applications from many domains. System effectiveness and power efficiency are dramatically improved through minimized data transfer, with extreme scaling and resilience.

The project has completed the first phase of requirements gathering activity from the use cases for building the SAGE storage platform and completed the architecture and design of the key Software & Hardware components of SAGE in the first year.

Further/continued collaboration on the codesign of the storage platform (incl. its API) and supporting ecosystem for exascale and beyond by better understanding future extreme computing/data workflows as needed by the international community.

Handbook on Summaries of Centres of Excellence

EXDCI Summaries: <https://exdci.eu/collaboration/coe>

Materials design at the eXascale



Materials are crucial to scientific and technological advances and industrial competitiveness, and to tackle key societal challenges - from energy and environment to health care, information and communications, manufacturing, safety and transportation. The current accuracy and predictive power of materials' simulations allow a paradigm shift for computational design and discovery, in which massive computing efforts can be launched to identify novel materials with improved properties and performance; behaviour of ever-increasing complexity can be addressed; sharing of data and work-flows accelerates synergies and empowers the science of big-data; and services can be provided in the form of data, codes, expertise, turnkey solutions, and a liquid market of computational resources. Europe has the human resources, track record and infrastructure to be worldwide leader in this field, and we want to create a CoE in materials' modelling, simulations, and design to endow our researchers and innovators with powerful new instruments to address the key scientific, industrial and societal challenges that require novel materials. This CoE will be a user-focused, thematic effort supporting the needs and the vision of all our core communities: domain scientists, software scientists and vendors, end-users in industry and in academic research, and high-performance computing centres. The proposal is structured along two core actions: (1) Community codes, their capabilities and reliability; provenance, preservation and sharing of data and work-flows; the ecosystem that integrates capabilities; and hardware support and transition to exascale architectures. (2) Integrating, training, and providing services to our core communities, while developing and implementing a model for sustainability, with the core benefit of propelling materials simulations in the practice of scientific research and industrial innovation.

The Novel Materials Discovery Laboratory



Main Achievements: In the first year of the NOMAD Laboratory CoE, methods were established to convert heterogenous data from the major computational materials science codes to a homogeneous, code-independent format and tools were developed to intuitively access, query and visualize this data in ways that will be useful for both academic and industrial users.

Areas of Potential Collaboration: The NOMAD Laboratory CoE is well placed to collaborate with other international initiatives in computational materials science, including those led by members of our Scientific Advisory Committee in the US (MICCoM) and Japan (CMI2).

Essentially every new commercial product, be they smart phones, solar cells, batteries, transport technology, artificial hips, etc., depends on improved or even novel materials. Computational materials science is increasingly influential as a method to identify such critical materials for both R&D. Enormous amounts of data, precious but heterogeneous and difficult to access or utilise, are already stored in repositories scattered across Europe. The NoMaD CoE will open new HPC opportunities by enabling access to this data and delivering powerful new tools to search, retrieve and manage it. NoMaD will foster sharing of all relevant data, building on the unique CECAM, Psi-k and ETSF communities, putting Europe ahead of materials science in other continents. Unprecedented, already initialised networking with researchers, with industry, with students and with other stakeholders will guarantee relevance and end-user value. NoMaD will become a crucial tool for atomistic simulations and multi-scale modelling in the physical, materials, and quantum-chemical sciences. This field is characterised by a healthy but heterogeneous eco-system of many different codes that are used at all HPC centers worldwide, with millions of CPU hours spent every day, some of them at petascale performance. NoMaD will integrate the leading codes and make their results comparable by converting (and compressing) existing inputs and outputs into a common format, thus making these valuable data accessible to academia and industry:

	<p>NoMaD will develop “big-data analytics” for materials science. This will require novel algorithms, e.g., for statistical learning based on the created materials encyclopedia, offering complex searches and novel visualisations. These challenges exploit the essential resources of our HPC partners. Without the infrastructure and services provided by the NoMaD CoE, much of the information created with the above mentioned petascale (towards exascale) computations would be wasted.</p>
<p>An e-infrastructure for software, training and consultancy in simulation and modelling</p> 	<p>Main Achievements: E-CAM is only in its first year, but has already had its first significant engagement with Industrial partners through a successful scoping workshop in Mainz and has started its training activities.</p> <p>Areas of Potential Collaboration: E-CAM seeks to collaborate in three main areas: advanced training; algorithmic innovation and software; and in industrial outreach.</p> <p>E-CAM will create, develop and sustain a European infrastructure for computational science applied to simulation and modelling of materials and of biological processes of industrial and societal importance. Building on the already significant network of 15 CECAM centres across Europe and the PRACE initiative, it will create a distributed, sustainable centre for simulation and modelling at and across the atomic, molecular and continuum scales.</p> <p>The ambitious goals of E-CAM will be achieved through three complementary instruments:</p> <ol style="list-style-type: none"> 1. development, testing, maintenance, and dissemination of robust software modules targeted at end-user needs; 2. advanced training of current and future academic and industrial researchers able to exploit these capabilities; 3. multidisciplinary, coordinated, top-level applied consultancy to industrial end-users (both large multinationals and SMEs). <p>The creation and development of this infrastructure will also impact academic research by creating a training opportunity for over 300 researchers in computational science as applied to their domain expertise.</p> <p>It will also provide a structure for the optimisation and long-term maintenance of important codes and provide a route for their exploitation. Based on the requests from its industrial end-users, E-CAM will deliver new software in a broad field by creating over 200 new, robust software modules. The modules will be written to run with maximum efficiency on hardware with different architectures, available at four PRACE centres and at the Hartree Centre for HPC in Industry. The modules will form the core of a software library (the E-CAM library) that will continue to grow and provide benefit well beyond the funding period of the project.</p> <p>E-CAM has a 60 month duration, involves 48 staff years of effort, has a total budget of €5,836,897 and is requesting funding from the EC of €4,836,897, commensurate with achieving its ambitious goals.</p>
<p>Energy oriented Centre of Excellence for computer applications</p> 	<p>Areas of Potential Collaboration: Several applications within the EoCoE are and will continue to be developed within international collaborations. Adding new features to applications and enable application for a broader range of architectures.</p> <p>The aim of the present proposal is to establish an Energy Oriented Centre of Excellence for computing applications, (EoCoE). EoCoE (pronounced “Echo”) will use the prodigious potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply. To achieve this goal, we believe that the present revolution in hardware technology calls for a similar paradigm change in the way application codes are designed. EoCoE will assist the energy transition via targeted support to four renewable energy pillars: Meteo, Materials, Water and Fusion, each with a heavy reliance on numerical modelling. These four pillars will be anchored within a strong transversal multidisciplinary basis providing high-end expertise in applied mathematics and HPC. EoCoE is structured around a central Franco-German hub coordinating a pan-European network, gathering a total of 8 countries and 23 teams. Its partners are strongly engaged in both the HPC and energy fields; a prerequisite for the long-term sustainability of EoCoE and also ensuring that it is deeply integrated in the overall European strategy for HPC. The primary goal of EoCoE is to create a new, long lasting and sustainable community around computational energy science. At the same time, EoCoE is committed to deliver high-impact results within the first three years. It will resolve current bottlenecks in application codes, leading to new modelling capabilities and scientific advances among the four user communities; it will develop cutting-edge mathematical and numerical methods, and tools to foster the usage of Exascale computing. Dedicated services for laboratories and industries will be established to leverage this expertise and to foster an ecosystem around HPC for energy. EoCoE will give birth to new collaborations and working methods and will encourage widely spread best practices.</p>

Excellence in Simulation of Weather and Climate in Europe



ESiWACE will substantially improve efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms by supporting the end-to-end workflow of global Earth system modelling in HPC environment. This will be obtained by improving and supporting (1) scalability of models, tools and data management on state-of-the-art supercomputer systems (2) Usability of models and tools throughout the European HPC eco-system, and (3) the Exploitability of the huge amount of resulting data.

We will develop solutions for cross-cutting HPC challenges particular to the weather and climate domain. This will range from the development of specific software products to the deployment of user facing services for both, computing and storage.

ESiWACE leverages two established European networks, namely (1) the European Network for Earth System modelling, representing the European climate modelling community and (2) the world leading European Centre for Medium-Range Weather Forecasts. The governance structure that defines the services to be provided will be driven by the European weather and climate science community.

Weather and climate computing have always been one of the key drivers for HPC development, with domain specific scientific and technical requirements that stretch the capability and capacity of existing software and hardware to the limits. By developing solutions for Europe and at European scale, ESiWACE will directly impact on the competitiveness of the European HPC industry by engendering new products, providing opportunities for exploitation beyond the project itself, and by enhancing the skills base of staff in both industry and academia.

ESiWACE will be at once thematic, as it focuses on the HPC application domain of climate and weather modeling, transversal, as it covers several aspects of computational science, and challenge-driven, as climate and weather predictability represents a major societal issue.

Center of Excellence for Global Systems Science



Main Achievements: CoeGSS addresses an entirely new community for the HPC domain.

Areas of Potential Collaboration: Collaborations are of particular interest within the Global Systems Science domain (political stakeholders, industrial stakeholders, decision makers, ...).

Global Systems Science – GSS – is an emerging research field focused on the risks and opportunities involved in global coordination problems. Examples of global systems include the internet, financial markets, intellectual property rights, global energy use and others.

Developing evidence and understanding in view of such systems and of related policies is rapidly becoming a vital challenge for modern societies. It requires capabilities for transdisciplinary work that cannot be mastered without massive use of ICT. By the nature of the problem, the relevant datasets are mostly very big, including data streams from social medi. To make things more complicated, the relevant algorithms do require the power of high-performance computing. High Performance Data Analysis (HPDA) is the key to success for GSS!

A key contribution of the Center of Excellence for Global Systems Science – COEGSS – will be the development of an HPC-based framework to generate customized synthetic populations for GSS applications. By blending GSS and HPC, we will be able to provide decision makers and civil society with real-time assessments of global risks and opportunities as well as with essential background knowledge about them. This will enable the HPC industry to supply hard- and software for applications well beyond the issues to which HPC has been dedicated so far.

Centre of Excellence for Biomolecular Research



Life Science research has become increasingly digital, and this development is accelerating rapidly. Biomolecular modelling techniques such as homology modelling, docking, and molecular simulation have advanced tremendously due to world-leading European research, resulting in extreme demands for better computational performance and throughput as these tools are used in applied research and industrial development. This research has direct influence on our daily life in areas such as health and medical applications, the development of new drugs, efficient drug delivery, biotechnology, environment, agriculture and food industry. Life Science is one of the largest and fastest growing communities in need of high-end computing, and it is a critically important industrial sector for Europe. However, compared to some other disciplines, the use of e-Infrastructure is still relatively new - many advanced techniques are not applied commercially due to limited experience. It requires significant support to:

- Make e-Infrastructure useable by researchers who are not computing experts,
- Improve the performance and applicability of key life science applications,
- Handle large amounts of data in computational workflows.

BioExcel proposes to tackle these challenges by establishing a dedicated CoE for Biomolecular Research, covering structural and functional studies of the building blocks of living organisms - proteins, DNA, saccharides, membranes, solvents and small molecules like drug compounds -

	<p>all areas where with large academic and industrial users bases in Europe. Specifically, BioExcel will</p> <ul style="list-style-type: none"> • Improve the efficiency and scalability of important software packages for biomolecular research; • Improve the usability of ICT technologies for biomolecular researchers in academia and industry; • Promote best practices and train end users in making good use of both software and e-Infrastructure. • Develop appropriate governance structures and business plans for a sustainable CoE. <p>Improved the performance, scalability and functionality of major codes for biomolecular modelling – GROMACS, HADDOCK and CPMD; and devised efficient workflows environments for automated large-scale modelling, simulation and analysis with associated data integration. Biomolecular modelling and simulations; development of workflow environments for biomolecular studies; exascale approaches for HPC and HTC; co-design; data management</p>
<p>Computational Biomedicine</p> 	<p>Objectives are to: coalesce the burgeoning HPC user community within the biomedical research field, promote innovation in the field of computational biomedical modelling and simulation, train future generations of scientists within the field of computational biomedicine, use best practice Software Carpentry tools and techniques to develop and sustain existing community codes, engage with a range of industries across the entire healthcare value chain, and engage closely with medical professionals to promote the tools, techniques as well as access mechanisms developed within our Centre.</p>
<p>Performance Optimisation and Productivity</p> 	<p>High performance Computing is becoming a fundamental tool for the progress of science and engineering and as such for economic competitiveness. The growing complexity of parallel computers is leading to a situation where code owners and users are not aware of the detailed issues affecting the performance of their applications. The result is often an inefficient use of the infrastructures. Even when the need to get further performance and efficiency is perceived, code developers may not have sufficient insight on its detailed causes for addressing the problem properly. This may lead to blind attempts to restructure codes and consequent lack of efficiency.</p> <p>The objective of POP is to operate a Centre of Excellence in Computing Applications in the area of Performance Optimisation and Productivity. POP will offer the service of precisely assessing the performance of computing application of any sort, from a few hundred to many thousand processors. Also, POP will show its Customers the issues affecting the performance of their code and the most optimal way to alleviate them. POP will target code owners and users from all domains, including infrastructure operators, academic and industrial users. The estimated population of such applications in Europe is 1500 and within the project lifetime POP has the ambition of serving over 150 such codes. The Added Value of POP's services is the savings generated in the operation and use of a code, which will result in a significant Return on Investment (fixing a code costs less than running it below its optimal levels) by employing best-in-class services and release capacity for resolving other priority issues.</p> <p>POP will be a best-in-class centre. By bringing together the European world-class expertise in the area and combining excellent academic resources with a practical, hand-on approach, it will improve the access to computing applications, thus allowing European researchers and industry to be more competitive.</p>