Proposal Book [- HDC - All]

Proposal Data

Acronym	DIONASYS
Full name	DIONASYS: Declarative and Interoperable Overlay Networks, Applications to Systems of Systems
Duration	36
Topic	HDC
Keywords	Distributed Systems, Systems of Systems, Environmental Observation Systems, Overlay Networks, Declarative Languages, Interoperability, Heterogeneity

Coordinator contact point for the proposal

Name	Etienne Rivière
Institution	University of Neuchatel
Address	Emile-Argand 11, 2000 Neuchâtel
Country	Switzerland
Phone	0041 77 432 20 37
Email	etienne.riviere@unine.ch

Consortium Partners

C/P	Institution	Contact	Other	Country	Legal Status	
C	University of Neuchatel	Etienne Rivière	Hugues Mercier, Pascal Felber	Switzerland	Public research organisation	**Ind.Eff.:**72 **Ind.Cost:**499,382 **Ind.Bud.:**328,128
P	Université de Bordeaux	Laurent Réveillère	David Bromberg, Floréal Morandat	France	Public research organisation	**Ind.Eff.:**78 **Ind.Cost:**411,071 **Ind.Bud.:**192,242
P	University of Lancaster	Gordon Blair	Geoff Coulson	United Kingdom	Public research organisation	**Ind.Eff.:**51 **Ind.Cost:**521,065 **Ind.Bud.:**416,852
P	Universitatea Tehnica din Cluj-Napoca	Virgil Dobrota	Andrei Bogdan Rus	Romania	Public research organisation	**Ind.Eff.:**21 **Ind.Cost:**95,000 **Ind.Bud.:**95,000

Abstract:

An increasing number of different kinds of resources, including everyday objects, are interconnected to each other. Some analysts estimate that around 50 billions of devices should be interconnected at the horizon of 2020. This paves the way for new large-scale systems, and requires the need for novel architectures and design principles to support such a scale. Further, these resources may be very different at both the hardware and software layer, in terms of both functional and non-functional properties. This will lead to system of systems that federate highly heterogeneous distributed systems as already illustrated, for instance, by environmental and earth observation systems. Designing such large, interconnected and heterogeneous systems is a daunting task. A possible way to overcome the complexity of contemporary distributed systems is to leverage overlay networks and their higher level

of abstraction. The virtualization of the underlying network resources allows providing a range of reusable network services. Many types of overlay networks have been proposed and developed in the previous years for a variety of networked systems, applications and services. However, the design and development of overlays remains a complex task, especially when dynamic adaptation, large-scale interoperability and composition are required. Adding interoperability, adaptation and composition capabilities often require huge and complex re-engineering of existing overlay implementations. In the context of overlay networks, it also requires appropriate abstractions and runtime support for allowing different type of overlays and structures to be linked, cooperate and provide adaptive interoperable end-to-end services in a dynamic fashion. In this project, we propose to raise the level of abstraction provided to designers of overlays and systems-of-systems. To this end, we are using a generative language approach to overlay design and composition. We will provide the corresponding new programming models, abstractions and tools. Our aim will be reached via the use of a high-level domain-specific language, declaring what should be achieved for the structure and functions of overlays, rather than by defining low-level nodes interactions. The proposed approach will be supported by a dedicated runtime implemented in a distributed systems development and deployment framework. The project follows a prototype-driven approach. It will feature a large-scale demonstrator linking heterogeneous overlays —networked systems and sensor networks—in an integrated manner, with support for adaptive and malleable end-to-end services and functionalities.

Relevance:

The proposed project targets topic 2: Heterogeneous Distributed Computing. It considers the specific domain of environmental and earth observation systems, which are highly heterogeneous distributed systems. Such systems are actually more accurately distributed systems of systems for which major issues are large-scale interoperability, energy-efficiency, and pervasive monitoring. Building applications for such distributed large-scale systems-of-systems that federate heterogeneous distributed systems is a daunting task. The project will build upon overlay networks to abstract resources from multiple, heterogeneous distributed systems. We propose a generative approach to overlay design and implementation to simplify the design of applications for complex, diverse and interconnected distributed systems. Our approach includes a new programming model, abstractions and tools such as a domain-specific language for specifying overlays from a global-scale abstraction perspective. As part of our generative design strategy, we will provide cross-layer interoperability and optimization techniques, from the application to software-defined network layers. The results of the project will be verified by building a demonstrator of the prototype implementation deployed on sensor networks at several sites and cloud aggregation layers. All code produced in the project will be made open source, with no exception and we identify several targets for long-term exploitation of the results.

Proposal Book [- HDC - All]

Proposal Data

Acronym	DIVIDEND
Full name	Distributed Heterogeneous Vertically IntegrateD ENergy Efficient Data centres
Duration	24
Topic	HDC
Keywords	vertical integration, energy crisis, data centres, energy accounting, data movement, compiler technology, specialisation, cross-layer optimisation, messaging fabric, heterogeneous systems architecture

Coordinator contact point for the proposal

Name	Michael O'Boyle
Institution	University of Edinburgh, School of Informatics
Address	Informatics Forum, 10 Crichton Street, Edinburgh, EH8 9AB
Country	United Kingdom
Phone	+44 (0) 131 650 5117
Email	mob@inf.ed.ac.uk

Consortium Partners

C/P	Institution	Contact	Other	Country	Legal Status	
C	University of Edinburgh, School of Informatics	Michael O'Boyle	Stratis Viglas, Hugh Leather, Boris Grot	United Kingdom	Public research organisation	**Ind.Eff.:**28 **Ind.Cost:**218,675 **Ind.Bud.:**174,940
P	UNIVERSITATEA POLITEHNICA TIMISOARA	Alexandru Amaricai	Marius Marcu, Oana Boncalo, Sebastian Fuicu	Romania	Public research organisation	**Ind.Eff.:**40 **Ind.Cost:**144,000 **Ind.Bud.:**144,000
P	Ecole Polytechnique Federale de Lausanne	Babak Falsafi	Edouard Bugnion	Switzerland	Public research organisation	**Ind.Eff.:**51 **Ind.Cost:**195,300 **Ind.Bud.:**195,300
P	Lancaster University, School of Computing and Communications	Zheng Wang		United Kingdom	Public research organisation	**Ind.Eff.:**30 **Ind.Cost:**193,555 **Ind.Bud.:**154,844
P	Queen's University of Belfast	Dimitrios Nikolopoulos	Hans Vandierendonck	United Kingdom	Public research organisation	**Ind.Eff.:**32 **Ind.Cost:**188,802 **Ind.Bud.:**151,042
P	Advanced Micro Devices (AMD)	Mauricio Breternitz	Wayne Burleson, Gabriel Loh	France	Large Enterprises	**Ind.Eff.:**62 **Ind.Cost:**750,000 **Ind.Bud.:**600,000

Abstract:

Our world is in the midst of a “big data” revolution, driven by the ubiquitous ability to gather, analyse, and query

datasets of unprecedented variety and size. The sheer storage volume and processing capacity required to manage these datasets has resulted in a transition away from desktop processing and toward warehouse-scale computing inside data centres. State-of-the-art data centres, employed by the likes of Google and Facebook, draw 20-30 MW of power, equivalent to 20,000 homes, with these companies needing many data centres each. The global data centre energy footprint is estimated at around 2% of the world's energy consumption and doubles every five years. Contemporary data centres have an average overhead of 90%, meaning that they consume up to 1.9 MW to deliver 1 MW of IT support; this is not cost-effective or environmentally sound. If the exponential data growth and processing capacity are to scale in the way that both the public and industry have come to rely upon, we must tackle the data centre energy crisis or face the reality of stagnated progress. With the semiconductor industry's inability to further lower operating voltages in processor and memory chips, the challenge is in developing technologies for large-scale data-centric computation with energy as a first-order design constraint. The DIVIDEND project attacks the data centre energy efficiency bottleneck through vertical integration, specialisation, and cross-layer optimisation. Our vision is to present heterogeneous data centres, combining CPUs, GPUs, and task-specific accelerators, as a unified entity to the application developer and let the runtime optimise the utilisation of the system resources during task execution. DIVIDEND embraces heterogeneity to dramatically lower the energy per task through extensive hardware specialisation while maintaining the ease of programmability of a homogeneous architecture. To lower communication latency and energy, DIVIDEND leverages SoC integration and prefers a lean point-to-point messaging fabric over complex connection-oriented network protocols. DIVIDEND addresses the programmability challenge by adapting and extending the industry-led heterogeneous systems architecture programming language and runtime initiative to account for energy awareness and data movement. DIVIDEND provides for a cross-layer energy optimisa- tion framework via a set of APIs for energy accounting and feedback between hardware, compilation, runtime, and application layers. The DIVIDEND project will usher in a new class of vertically integrated data centres and will take a first stab at resolving the energy crisis by improving the power usage effectiveness of data centres by at least 50%.

Relevance:

The big data revolution that is currently transforming the modern world is threatened by the semi-conductor energy crisis. While today's data centres use thousands of homogeneous servers interconnected with commodity networking technology, the need to improve both performance and energy-efficiency argue for heterogeneity and specialisation at both processing and communication layers. The associated challenges include the development of specialised hardware and interconnect technologies, programmability and runtime management of a heterogeneous system, and a comprehensive monitoring framework orchestrating computation and communication for peak efficiency. Programming models and tools: We will extend the industry-led heterogeneous systems architecture (HSA) with higher levels of abstraction to account for a multitude of diverse processing elements and new ways of non-intrusively exposing optimisation opportunities to the lower levels. Data movement and management: We will develop a point-to-point messaging fabric for low overhead data movement within a data centre. Monitoring and optimisation techniques: We will make energy monitoring and accounting a first-class citizen and enable energy-performance optimisations by integrating into tools machine learning and dynamic code reconstruction techniques. We will provide a vertical solution that holistically integrates programming language, runtime, compilation, interconnect, architecture, and monitoring layers for Distributed Heterogeneous Vertically IntegrateD ENergy Efficient Data centres (DIVIDEND).

Proposal Book [- HDC - All]

Proposal Data

Acronym	HPDCJ
Full name	Heterogenous parallel and distributed computing with Java
Duration	36
Topic	HDC
Keywords	parallel computing, heterogenous computimg, Java, GPU, PGAS

Coordinator contact point for the proposal

Name	Piotr Bala
Institution	University of Warsaw / ICM
Address	Krakowskie Przedmiescie 26, 00-927 Warszawa, POLAND
Country	Poland
Phone	+48 22 8749 400
Email	bala@icm.edu.pl

Consortium Partners

C/P	Institution	Contact	Other	Country	Legal Status	
C	University of Warsaw / ICM	Piotr Bala		Poland	Public research organisation	**Ind.Eff.:**72 **Ind.Cost:**242,354 **Ind.Bud.:**242,354
P	IBM Research Lab	Costas Bekas	Allessandro Curoni	Switzerland	Private research organisation	**Ind.Eff.:**40 **Ind.Cost:**630,680 **Ind.Bud.:**324,096
P	Queens University of Belfast / School of Electronics, Electrical Engineering and Computer Science	Dimitrios Nikolopoulos	Hans Vandierendonck	United Kingdom	Public research organisation	**Ind.Eff.:**32 **Ind.Cost:**349,559 **Ind.Bud.:**279,647
P	Bilkent University	Gedik Bugra		Turkey	Public research organisation	**Ind.Eff.:**60 **Ind.Cost:**124,292 **Ind.Bud.:**124,292

Abstract:

Our proposal focuses on the ease of use and programmability of Java for distributed heterogeneous computing in order to make it exploitable by the huge user base of mainstream computing. Based on the previous work (PCJ library http://pcj.icm.edu.pl), we will introduce and transparently expose parallelism in Java, with minimal change to the specifics of the language thus allowing programmers to focus on the application. We have demonstrated power and scalability of the PCJ library for the parallel systems and we will extend it for the cases where communication cost and latency could be higher. We will extend existing solution with the capability of running on the heterogeneous systems including GPU and mobile devices. The user will obtain possibility to execute

computational intensive parts of the application on the multiple GPUs. Since our solution is based on Java it can be easily run on the mobile devices. Within project we will extend library capabilities with the optimised communication and scheduling mechanism necessary to use fully such devices. We will utilize potential of parallel Java library to process distribute data. The existing solution benefits from the parallel I/O performed by the multiple JVMs. We will use this solution to optimize process of data distribution and storage including streaming od the large data sets. We will address dependability and resilience by adding fault tolerance mechanisms to the parallel Java library including fault detection and rescheduling of the application execution. The mechanism will extend capabilities of the existing PCJ library and will be transparent to the users. We will show the applicability of our framework for distributed heterogeneous systems by a set of selected, key applications including data-intensive Big Data applications. Our potential success will create solution for Java programming that will be attractive to a wide mainstream user base and will thus have a game-changing influence on the European computing industry. We assembled a carefully selected team with complementary focuses and the right degree of overlap. Most of the partners have worked in close collaboration in previous (EU) projects with remarkable success. We believe this to become a key pilot project that can open the way for future research which will have a profound impact on mainstream computing.

Relevance:

The proposal is dealing with the Java as a language for parallel and distributed simulations. Java is not yet widely used in the area of a high performance parallel computing, but is considered as future emerging technology. There is no doubt that introducing Java in this area will have significant scientific and technical impacts in the mid and long term. Since project involves participants from the different countries with the different research focus the proposed research is international and multidisciplinary. We address theme Programming models and tools by proposing here new programming models, abstractions and tools for software development. The usage of Java ensures abstraction from physical devices and connectivity. Proposed solution allows to obtain high performance across platforms. It also includes verification and resource management. We address theme Data movement and management by proposing software and tools which benefit from the Java functionality which allows for easy development of streaming and placement of data across platforms, data reduction and inference. In the project we will further develop these functionalities. The theme Dependability and resilience is addressed by the design and development of fault-tolerance and computational reliability into PCJ library. It will be also extended to support heterogeneous distributed systems in the secure way.