



TERATEC 2014 Forum

JULY 1 & 2, 2014 - ECOLE POLYTECHNIQUE - 91 PALAISEAU - 9th EDITION

THE INTERNATIONAL MEETING FOR SIMULATION AND HIGH PERFORMANCE COMPUTING

Simulation and high performance computing: Exhibitors' novelties

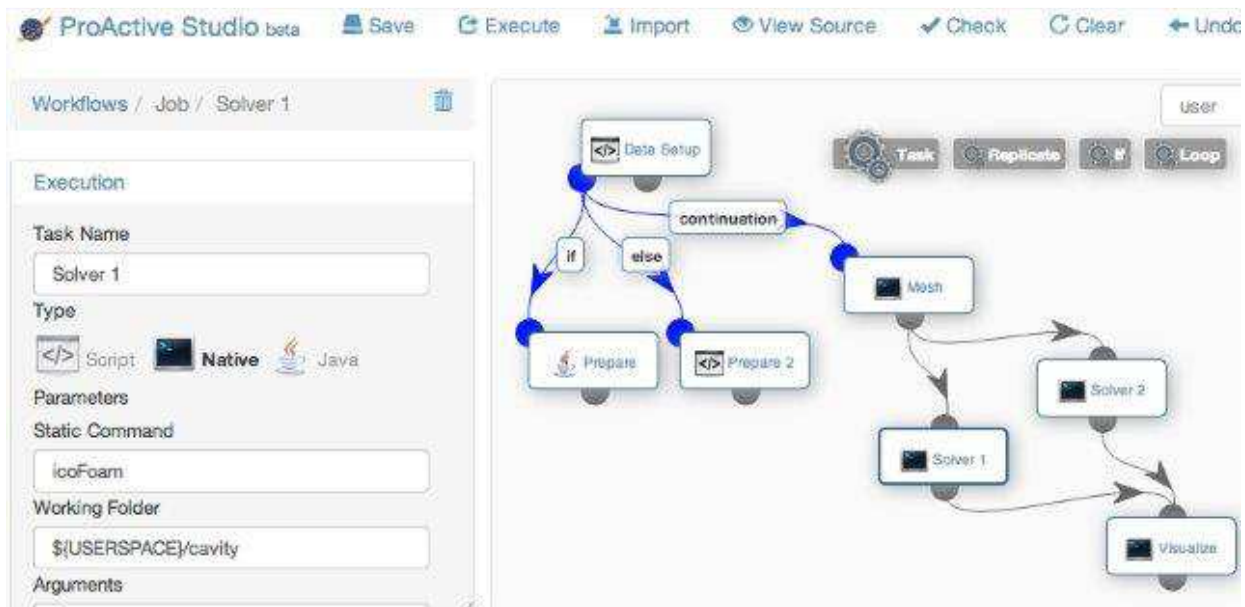
High performance digital design and simulation make a key contribution to innovation for all industry and service sectors and to increase the global quality of their products and their capacity for innovation. The exhibition part of the Forum Teratec reflects, through the offer of exhibitors, the wide variety of solutions and services in this field.

Again this year, during two days, the whole HPC industry will be gathered on 60 stands: Systems manufacturers and software vendors, integrators and distributors, service providers, academic and laboratory researchers, public and private sector developers will present their latest HPC innovations.

ACTIVEEON * Stand 15

Press Contact: Nino Zurabashvili
Tel: 09 88 77 76 65 - Mail: nino@activeeon.com

ActiveEon, an *inria* "spin-off", is a professional Open Source software company, co-developing *ProActive Parallel Suite@* with the research team OASIS (a joint team between INRIA, CNRS/I3S and the University of Nice – Sophia Antipolis) and providing a full range of services in the fields of parallel, distributed, grid and cloud computing for all our products.



Our users are present in various industries. ProActive suite of products, solutions and expertise are optimized for the following markets:

- IT & Telco ■ Banking & Insurance ■ Engineering & Manufacturing ■ Life Science & Biotech ■ Media & Entertainment ■ Academia & Government

ProActive Workflows & Scheduling

Workflows & Scheduling allows you to easily distribute and execute all company jobs and business applications, monitor activity and view jobs results. It ensures more work done with fewer resources, managing heterogeneous platforms, multiple sites. It maximizes utilization of both existing and Cloud IT infrastructures.

ProActive Distributed Matlab, Scilab, R

As modern scientific and engineering problems grow in complexity, the computation time and memory requirement increase and parallel computing becomes a necessity. ProActive integrates with de-facto standards in scientific and engineering environments such as Matlab, Scilab, R. Directly from within these familiar environments, it provides users with the capacity to distribute executions and manage data transfers on other Desktop machines, Clusters, Grids and Clouds. A single tool for accelerating all your scientific languages.

ProActive Cloud Automation

With ActiveEon's Cloud solution, automate the deployment of complex multi-VMs applications. You simplify self-service deployment, as well as automate the delivery of in production services. You set and control application elasticity with automated scaling-up and down, horizontally and vertically. Once you define the service deployment, you are able to deploy anywhere with workflow reversibility. With unified cloud management, applications execute on multi-vendor private, public and hybrid clouds.

ALLINEA SOFTWARE * Stand 13

Press Contact Jacques PHILOUZE

Tel: +44 (0)1926 623 231 - Email: jacques@allinea.com

HPC performance insight improves usage at Cardiff University

Allinea Performance Reports enables supercomputers to deliver more

Warwick, UK – A new tool from Allinea Software is allowing users and supercomputer owners to see clearly what is happening within their HPC system – and ensure optimum performance from their applications. System owners, developers and users are all interested to see what is going on 'inside' after they set a job to run. CPU hours are expensive and often in short supply – it's vital to make the best use of them by ensuring the jobs run efficiently.

Allinea Performance Reports analyzes key metrics such as the effective vectorization, I/O and MPI performance for any application, allowing users and system managers to make sure code runs as effectively and efficiently as possible.

Cardiff University has been using Allinea Performance Reports to gain more insight into their systems at both its ARCCA¹ service – which provides access and training in advanced computing to students and researchers – and at HPC Wales², its world-class supercomputing service for researchers and innovators in the public and private sectors.

The University needed a simple tool to analyze what each application was doing and how it could improve. The teams are managing a wide variety of application deployments across a diverse spectrum of research domains on both the ARCCA and HPC Wales systems. They required a tool that could handle the diverse range of codes required by these communities. Allinea Performance Reports has allowed them to unlock performance across the full spread of its infrastructure.

"Allinea Performance Reports has given us an insight into where workloads can be adjusted, which can be placed on more effectively on the hardware, or even recompiled to be able to perform more simulations to satisfy the demand for time on our systems," says Martyn Guest, Director of ARCCA and Technical Director of HPC Wales.

More efficient machine use on the same compute resource means more research outputs, optimal performance in terms of energy consumption, more industrial success, resulting in greater returns on investment for users like the University.

¹ www.cardiff.ac.uk/arcca

² A pan-Wales distributed network of supercomputer systems involving all the Welsh HEIs, see www.hpcwales.co.uk

"Our customers told us they need to see what is really happening in the HPC applications running on their systems and their real impact on the results," says David Lecomber, CEO of Allinea Software. "They wanted a tool that would be so easy and non-intrusive to apply that it could analyze applications without configuration or change and help them to recognize where to focus their efforts to make the biggest differences – from using ISV software to running locally developed applications. Allinea Performance Reports allows users to improve the quality of every CPU hour used, " continues Lecomber, "and each application to make the most use of the right system, the right configuration and the right usage.

Allinea Performance Reports was first showcased in November at SC13, in Colorado, USA, and won immediate fans. After an early access program that saw it deployed on a large number of HPC installations across the world; the product is now available from Allinea.

ANSYS France * Stand 02


Press contact: Sabine MAÏDA
Tel: 01 30 60 15 00 - Mail: sabine.maida@ansys.com

ANSYS: Built for HPC

As the acknowledged industry leader in engineering simulation, ANSYS has made a commitment to offer a comprehensive suite of solver and HPC advancements across the entire range of physics. Whether your focus is on structural, thermal, fluids or electromagnetic analysis -or the complex interactions of multiple physical.

ANSYS® technology enables highly scalable HPC deployment, giving you virtually unlimited capacity for high-fidelity simulation and the detail it provides. You can launch our HPC solutions within a workgroup or across a distributed enterprise — whether using local workstations, department clusters or enterprise servers — wherever your resources and people are located.

For ANSYS software to effectively leverage today's hardware, efficient execution on multiple cores is essential. ANSYS continues to release consistent, significant solution improvements, developed specifically to sustain speed and scaling on the latest HPC workstations, servers or clusters.



The various ANSYS HPC licensing options allow scalability to whatever computational level a simulation requires, from small user group options to enable entry-level parallel processing up to virtually unlimited parallel capacity. For large user groups, ANSYS facilitates multiple parallel processing simulations, highly scalable for the most challenging projects when needed. Our HPC tools intelligently distribute complex problems across multiple CPUs and GPUs, leading to the fastest, best possible solution.

- ANSYS HPC Packs
- ANSYS HPC Workgroup
- ANSYS HPC Enterprise
- ANSYS HPC Parametric Pack
- ANSYS Electronics HPC

BULL * Stand 47

Press Contact: Aurélie Negro Tel: +33 1 58 04 05 02 - Mail: aurelie.negro@bull.net
Contact on the stand: Pascale Bernier-Bruna, +33 6 74 09 38 82 - Mail: pascale.bernier-bruna@bull.net

Bull is a leader in secure mission-critical digital systems. The Group is dedicated to developing and implementing solutions where computing power and security serve to optimize its customers' information systems, to support their business. Bull operates in high added-value markets including computer simulation, Cloud computing and 'computing power plants', outsourcing and security.


Currently Bull employs around 9,000 people across more than 50 countries, with over 700 staff totally focused on R&D. In 2013, Bull recorded revenues of €1.26 billion.

As an expert in delivering ultra high-power, Bull is now one of the world leaders in Extreme Computing. Bull is recognized for the technological excellence of its bullx range of servers, its HPC applications expertise and its ability to manage large-scale projects. bullx benefits from a major, patented innovation from Bull: direct liquid cooling, which boosts energy performance by 40%.

Across the world, numerous institutions (SURFsara in the Netherlands, IT4Innovations in the Czech Republic, CSC in Finland, the Universities of Dresden, Düsseldorf, Grenoble, Reims...) and companies (Météo France, Dassault Aviation, Ceaero...) have turned to Bull to implement powerful, robust systems that are easy to manage and use, and are designed for round-the-clock operation. Every day, thanks to Bull, their researchers and engineers are pushing back the boundaries of the possible.

★★★**In preview at the Teratec Forum, Bull will present:**

- its « extreme factory »HPC private and public cloud offer
- a prototype of the new direct liquid cooled blade, incorporating the future Intel Xeon processor (code-named Haswell EP)



BUREAU 14 * Stand 5

Never compromise again: performance, scalability and reliability in one product!


Bureau 14 SAS is proud to announce that its flagship product, quasardb, the infinitely scalable database, now fully supports ACID transactions. This groundbreaking innovation has been achieved thanks to years of research and is now ready for production.

Users no longer have to choose between performance, reliability and scalability: thanks to quasardb they can now have the three! This feature will be available for beta-testing this summer to selected customers. General availability is planned for quasardb 1.2.0 scheduled for Q3 2014.

What is quasardb?

Quasardb is a distributed, high-performance, associative database designed from the ground up for the most demanding environments.

Based on decades of theoretical research and years of prototyping, quasardb stands on the shoulder of giants: it combines breakthroughs from relational databases, operating systems and network distribution to redefine the state of the art. quasardb already withstood the fire of critical environments where failure isn't an option and will change the way you look at associative databases.



CRAY * Stand 18

Contact: Federica PISANI, Cray Marketing Manager Europe
Tel : +44 117 9744048 - Cell +44 742 345 2356 - Email: pisani@cray.com

Cray Solutions for High Performance Computing

Cray offers a comprehensive portfolio of solutions for supercomputing, storage and big data, for a wide variety of requirements. Whether your organization is small or one of the world's largest research centers, Cray technologies enable users in technical and scientific computing to accelerate performance, improve efficiency and extend the capabilities of the most demanding applications. With a solution for every budget and requirement, Cray makes it easy to take advantage of the latest advances in high performance computing.

Computing

Cray offers a range of supercomputing solutions geared to meet each customer's unique need. Ranging from highly configurable and fast cluster systems to integrated, scalable and powerful supercomputers, Cray technologies address today's critical computing challenges.

Storage

Cray is your trusted expert in complete and open storage solutions for big data and high performance computing. Cray helps you utilize the right storage, build on open systems, and get results faster with scalable solutions built on best-in-class storage systems — offering choice for flexibility and complete integration.



Analytics

Cray technologies are built to solve big data problems. The YarcData® Urika® appliance was created to transform big data into meaningful information. The turnkey Cray® CS300™ cluster supercomputer for Hadoop provides performance to address the most demanding Hadoop requirements.



EMG2/BITWARE * Stand 53

Press contact: Anthony BESSEAU
Tel: 01 69 59 14 31 - Mail: anthony.besseau@emg2.com

Bittware – Terabox : A highly Flexible and Low Power HPC solution

BittWare, a leader in high-end signal and data processing boards and solutions, has recently made a remarkable entry on the HPC world with their new **TERABOX**, a High Performance Reconfigurable Computing Platform, based on latest FPGA technology.



For less than a kW, the TERABOX provides up to 20 TeraFLOPS of total processing power, along with 6.5 Terabits/sec of memory bandwidth and 1.28 Terabits/sec of I/O (up to 128 x 10GigE or 32 x 40GigE links) all in a very compact 5U rackmount solution. This low power, highly flexible system, fully tested and configured, is ideal for demanding power processing and high density network applications and comes instantly ready for HPC application development.

Bittware is also announcing their **newest product family based on ARRIA 10**, ALTERA's highest density and power-efficient FPGA, with Transceivers up to 28Gbps, hardened floating-point DSP delivering up to 1.5 TFLOPs blocks and dual-core ARM® Cortex™-A9 MPCore™). BittWare's ARRIA 10 family bring the most powerful and usable FPGA boards and systems to date in a variety of formats with features including flexible memory configurations, sophisticated clocking and timing options, QSFP28 cages that support 100Gbps (including 100GigE) optical transceivers, FPGA Mezzanine Card (FMC).

Bittware is also part of ALTERA's **OpenCl initiative** to enable designers to code their systems and algorithms in a high-level C-based framework, cutting development cycles from many months to a few weeks or even days.

ENGINSOFT FRANCE * Stand 20

Press Contact: Laure RAYMOND
Tel: 01.41.22.99.30 - Mail: l.raymond@enginsoft.com



In both industry and research, High Performance Computing and Simulation encourage innovative product and service development. By using HPC, industries have found a way to accelerate time for modeling and analyzing large data volumes. These time savings lead to cost reduction during each step of the product or the process life cycle (design, optimization, validation) and increase the competitiveness of businesses.

On the occasion of the Forum Teratec 2014, find out more about EnginSoft France and its simulation solutions using HPC in order to design, model and analyze complex phenomena such as combustion, aero-acoustics... :

- **modeFRONTIER®** is a multi-disciplinary and multi-objective optimization and design environment, written to allow easy coupling to almost any computer aided engineering (CAE) tool, whether commercial or in-house.
- **Flowmaster®** provides a 1D solution to simulate thermo-fluid systems quickly, simply and accurately; reducing cost and lead times in product development and maintenance
- **CharLES®** is a 3D CFD code able to perform combustion, aero acoustics and multiphase calculations using LES method (Large Eddy Simulation).

News

EnginSoft is a member of ETP4HPC, the European Platform for HPC. EnginSoft France has acquired the distribution rights of FENSAP-ICE®, the premier in-flight icing simulation comprehensive “system”, used worldwide at major aerospace companies.



ETP4HPC * Stand 10

An Industry-Led Forum founded by Stakeholders of HPC Technology

ETP4HPC European Technology Platform for High Performance Computing is the voice for HPC technology in Europe and works closely with the European Commission DG-CONNECT to implement a strategy for developing European HPC all along its value chain, from technologies to usages. ETP4HPC has signed a contractual Public Private Partnership (cPPP) with the EC in December 2013 with this objective.

With more than 50 members from industry, academia and the service sector, ETP4HPC is an association open to all HPC stakeholders sharing these objectives. In addition to maintaining a Strategic Research Agenda and contributing to Horizon 2020 Work Programmes in HPC, ETP4HPC has different activities and work groups on topics such as HPC SMEs, Education and Training in HPC, Intellectual Property.



Find out more at <http://www.etp4hpc.eu> and/or contact office@etp4hp.eu.

EUROTECH * Stand 11

Contact: Giovanbattista Mattiussi, Marketing Manager HPC
direct: +39 0433 485467 - mobile: +39 345 7153193 - Mail: giovanbattista.mattiussi@eurotech.com

Eurotech is a publicly listed global company based in Italy, with subsidiaries in Europe, North America and Asia. The Eurotech Group develops and markets miniaturized computers and high performance computing systems.

The Eurotech HPC division has more than 15 years of experience in designing and manufacturing HPC systems, with a commitment to continuous innovation through R&D nurtured in house and in collaboration with the most important research centres in Europe. Eurotech HPC has deployed and maintained HPC solutions to a variety of customers, delivering value combining HPC technology and experience to meet customer needs

Eurotech supercomputer line, Aurora, offers a range of advanced, hot water cooled, ultra dense, green HPC systems. Aurora supercomputers are designed to excel in energy efficiency, footprint, reliability and performance. With direct hot water cooling, the Aurora HPC systems are based on innovative, modular and scalable architectures built on standard commercial components, to maximize functionality, performance and efficiency.

Eurotech is the key engineering partner in FP7 research projects, like DEEP (dynamic exascale entry platform), QPACE2, Euroexa, Euroserver and others. This continuous dedication to R&D allows Eurotech solutions to look into the future, challenging exascale-level problems. As funding and board member of the ETP4HPC, Eurotech wants to contribute to a stronger European HPC value chain.

LABORATOIRE EXASCALE COMPUTING RESEARCH * Stand 24

Contact: Marie-Christine SAWLEY
Email: marie-christine.sawley@intel.com

R&D for methodologies and application software for Exascale

The Exascale Computing Research lab was founded in 2010 by CEA, GENCI, Intel and UVSQ. It is part of research networks dedicated to the Exascale, working on the frontiers of extreme scaling capabilities. In particular, the lab is a member of the Intel EMEA Exascale Labs in France, Belgium, Germany and Spain. Optimizing the software from the application layer down to the system layer remains the biggest challenge for the year ahead in preparing for systems with millions of cores. In order to address this challenge, we need to work on a number of fronts: on one hand, developing advanced tools to analyse at fine grain what is happening on the computational unit and on the communication network; on the other hand, working on HPC applications to remove bottlenecks to higher scalability and better efficiency. The ECR lab works on both of these.

MAQAO with its associated tools form a suite co-developed by the Exascale computing research, the Universities of Versailles St Quentin and Bordeaux, designed for performance analysis of HPC applications. In 2013, the complete version of MAQAO for the Xeon PHI was published as Open Source¹. A number of tutorials were organised for the HPC community, in particular under the VI-HPS organisation, which has already announced forthcoming workshops⁵. A second significant step was taken with the Open Source release of CTI, the Codelet Tuning Infrastructure². CTI consists of a series of tools and a repository of codelets – a few lines of codes representing the most critical parts- extracted from the complete application with its runtime environment, targeting parametric reruns to explore compilation options or different architectures. The work has been presented in a number of communications and tutorials, including during Teratec 2013.

Central to the lab, the group Programming and Execution Models works to increase efficiency of data movement and messages transport between computing nodes. In 2012 and 2013, the group developed an innovative method for MPI applications to reduce waiting times due to message exchange on the communication network. This method, based on Collaborative Polling allows the efficient interleaving of computing and communication phases, without adding any overhead due to the utilization of threads for the communication. The gains in performance for on a number of relevant codes obtained through an Infiniband implementation were presented in a scientific communication³.

Hand in hand with the development of the tools, the tasks of analyzing and optimizing HPC applications from the industrial or academic partners continued. Based upon a strong collaborative approach between the developer and the lab, the goal is to build expertise to optimize performance and efficiency of such applications and to prepare them for higher scalability. In 2014, the most salient results obtained with the molecular dynamics code Polaris(MD) of the CEA Life Science were presented at the EASC 2014 conference in Stockholm. We will present the rich application portfolio on domains such as combustion, turbulence, free surface simulation and thermonuclear fusion.

The lab was represented in large events of the HPC community: on the Intel stand in Leipzig (ISC14) and Denver (SC13); Two members of the lab, INTEL and UVSQ, along with a third French partner, INRIA, are members of the consortium of the EXA2CT project⁴. This project, financed in the Framework Program 7, is coordinated by IMEC in Belgium. The goal is to create a library of proto applications integrating novel developments in exascale algorithms and programming models, and to demonstrate scaling capacities up to T0 PRACE class by 2016 for the selected codes, most of them with high industrial usage.

The lab has moved its operation and staff to the TERATEC technopole in June 2014.

¹ http://maqao.bordeaux.inria.fr/index.php?option=com_content&view=article&id=2

² <http://code.google.com/p/codelet-tuning-infrastructure/>

³ <http://rd.springer.com/article/10.1007%2Fs00607-013-0327-z>

⁴ www.exa2ct.eu

⁵ <http://www.vi-hps.org/training/tws/tw14.html>

FUJITSU * Stand 40

Contact : Thomas Empeigne, Assistant Chef de Marché Infrastructure

Tel : +33 1 41 97 90 13 - Mail : Thomas.Empeigne@ts.fujitsu.com

Demand for supercomputing capabilities is rising the world over, driven primarily by the need for effective, reliable solutions to increasingly complex social, environmental and business challenges. This is also pushing high-end computational modeling and simulation capabilities beyond the R&D labs and into the commercial and manufacturing world.

We are positioned at the forefront of the supercomputing space with [30 years' experience](#) in the successful development of high-performance systems. Cooperating with leading research bodies, we believe that supercomputing has the potential to solve the world's most pressing dilemmas

Computer simulation is an essential technology that is instrumental in solving many of today's most puzzling and complex problems. It enables organizations to address a large variety of topics from research and development to product design and optimization.

Fujitsu's PRIMEHPC FX10 supercomputer provides the ability to address these high magnitude problems by delivering over 23 petaflops, a quantum leap in processing performance.

Ultra-high Speed and Ultra-large Scale Supercomputer

Problems previously constrained or impossible to solve due to performance limits are now able to be handled. This is due to the PRIMEHPC FX10's maximum peak performance of 23.2 Petaflops and memory that scales up to 6 PB with a 98,304 node configuration.

Green Credentials as well as High Performance Mean Power Savings

In today's quest for a greener

world the compromise between high performance and environmental footprint is a major issue. At the heart of PRIMEHPC FX10 are SPARC64™ IXfx processors that deliver ultrahigh performance of 236.5 Gigaflops and superb power efficiency of over 2 Gigaflops per watt.

Application Performance and Simple Development

SPARC64™ IXfx processor includes extensions for HPC applications known as HPC-ACE. This plus wide memory bandwidth, high performance Tofu interconnect, advanced compilers and libraries, enable applications to achieve the best performance ever. In addition, the time and effort to adapt to massively parallel processing is reduced through the use of VISIMPACT, which simplifies the implementation of hybrid parallel applications combining MPI and thread parallelism.

High Reliability and Operability in Large Systems

Incorporating RAS functions, proven on mainframe and high-end SPARC64 servers, SPARC64™ IXfx processor delivers higher reliability and operability. The flexible 6D Mesh/Torus architecture of the Tofu interconnect also contributes to overall reliability. The result is outstanding operation: enhanced by the advanced set of system management, monitoring, and job management software, and the highly scalable distributed file system.



GENCI * Stand 19

Contact Presse: Laetitia BAUDIN
Tel: +33 6 16 27 68 73 - Mail: laetitia.baudin@genci.fr

GENCI: Democratizing the use of numerical simulation and HPC in all the fields of science and industry

Thanks to the commitment of the leading players in academic research and the support of the French government, GENCI has been carrying out three main missions:

- Implementing the French national strategy for HPC for the benefit of the scientific community in relation with the three national computing centers for academic research.
Thanks to the investments made by GENCI since 2007, 600 scientific teams access to powerful resources each year. The renewal of Cines' computing resources, at the end of 2014, will increase the French computing power to 3.5 Pflops.
- Contributing to the setup of an integrated HPC ecosystem at the European level.
Representing France, GENCI is involved in PRACE, which is the European research infrastructure, pooling together a network of world-class supercomputers available for European scientists and industrials.
- Promoting numerical simulation and HPC within academic research and industry, and through a specific initiative toward SMEs.

Altogether with Bpifrance and Inria, GENCI is leading the 'HPC for SME Initiative' for assisting SMEs to integrate numerical simulation in their business models and increase their competitiveness.

KALRAY * Stand 12

Press Contact: Corine LAMAGDELEINE
Tel: +33 6 64 05 69 30 - Email: corine.lamagdeleine@kalray.eu

KALRAY will demonstrate the first PCIe accelerator based on 4 MPPA-256 processors and its target HPC applications (SGEMM & signal processing) at the TERATEC FORUM 2014 and present its MPPA ACCESSCORE and MPPA DEVELOPER products.

MPPA-256 is a low power many-core processor that integrates 256 processing engine (PE) cores and 32 resource management (RM) cores on a single 28nm CMOS chip. All the cores are fully programmable MIMD cores with a dedicated IEEE-754 SP and DP FPU.

MPPA ACCESSCORE is a complete, standard-based, many-core development environment that includes several programming models with FORTRAN, C, C++ with MPI, OpenMP and OpenCL and a dedicated real-time trace system.

MPPA DEVELOPER is a "ready to use" development platform to evaluate, develop and optimize applications on KALRAY's MPPA MANYCORE processors. MPPA DEVELOPER delivers the unique processing power of the 256 cores of the MPPA 256 associated with MPPA ACCESSCORE SDK ideally suited for compute intensive applications.

Thanks to the very high efficiency of the MPPA-256 processor (more than 20GFLOPS/W), the PCIe accelerator can be integrated with a very high computing density: depending of the application needs, a single accelerator provides more than 1000 MIMD cores and a chassis with 8 accelerator provides more than 8000 cores with direct, scalable communication between them

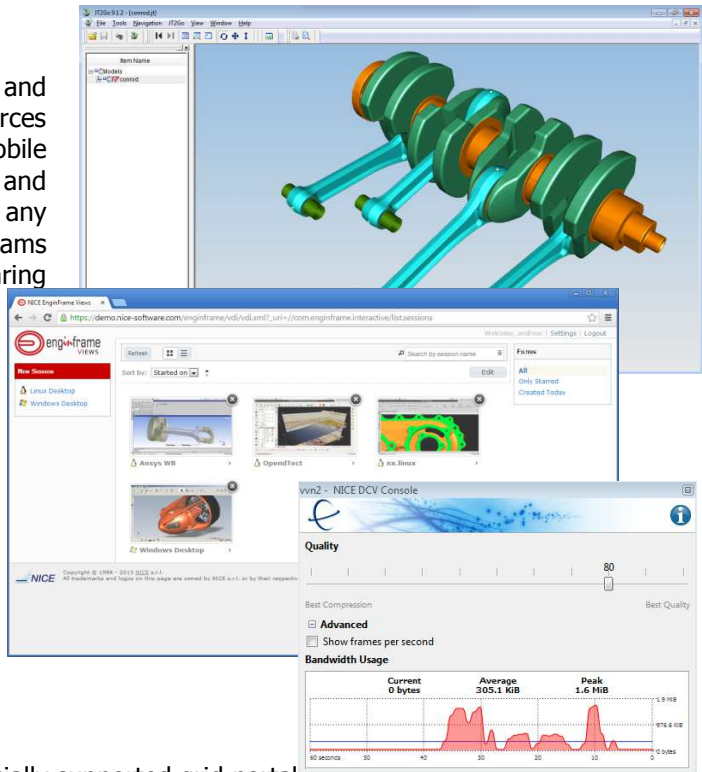
NICE SOFTWARE * Stand 32

Contact: Stefano AMICO
Tel: 00390141901516 - Mail: stefano.amico@nice-software.com

Efficient and secure remote access to graphics-intensive 3D applications in physical and virtual environments

HPC and Visualization in the Technical Cloud

NICE helps centralizing, optimizing and consolidating HPC and visualization resources while empowering distributed and mobile Engineering workforces to run batch and interactive applications anywhere, with any client. NICE also allows engineering teams around the world to collaborate by sharing application sessions over standard networks.



Key features of NICE Products

- 2D/3D Interactive session management
- Integration with NVIDIA® Grid API
- Integrated service scheduler
- Access multiple clusters from a single portal
- Compatible with OpenGL® and DirectX®
- Include file transfer automation
- Simplified installation
- Improved codecs' performance

NICE EnginFrame is an advanced, commercially supported grid portal that provides access to grid-enabled infrastructures, HPC clusters, data, licenses, and interactive applications. It can be accessed by any authorized user with a standard web browser.

Link: <http://www.enginframe.com> | **flyer:** <http://www.nice-software.com/html/NICE-EnginFrame.pdf>

NICE Desktop Cloud Visualization (DCV), NICE's remote 3D "virtual workstation" offering, is an advanced remote 3D access technology that allows technical end users to access robust 3D modeling tools through a web-based portal

Link: <http://www.nice-dcv.com> | **flyer:** <http://www.nice-software.com/html/NICE-DCV.pdf>

NOESIS SOLUTIONS * Stand 43

Press Contact: Bertrand MONNIER
Tel: +33 608 659 857 - Mail: bertrand.monnier@noesisolutions.com

Noesis Solutions is a simulation innovation partner to manufacturers in engineering-intense industries. Specialized in simulation process integration and numerical design optimization, its flagship software Optimus helps customers transform their simulation based engineering processes to adopt an 'Engineer by Objective' development strategy.

Optimus automates simulation based design processes, integrating any CAD or CAE software. Conducting Design of Experiments & Surrogate Modeling, Optimus provides the metrics that deepen engineering insight and accelerate the decision process. Using smart & reliable optimization methods, Optimus identifies the best designs in the shortest time frame. Customers using this approach report design time savings averaging over 30%, while achieving 10% or more product performance improvements.



NUMSCALE * Stand 14

Press contact: Charles PRETOT
Tel : 01 69 15 73 26 - Mail: charles.pretot@numscale.com

NumScale, start-up based in the Saclay plateau science and technology cluster, south of Paris, offers innovative solutions for the development and optimization of software that allow developers to make use of the full capacity of modern computing systems, in particular parallel architectures, so as to achieve maximum performance.

The NumScale solutions bring together different domains of expertise from academia so as to optimize software performance on all aspects:

- **Numerical analysis and algorithmics**, to provide the best state-of-the-art resolution methods and solve various domain-specific problems (simulation, signal processing, computer vision...)
- **Processor architecture**, to make the most of CPUs, GPUs, accelerators and DSPs, from supercomputers to embedded systems,
- **Programming languages and techniques**, in particular C, C++, Python, MATLAB and Fortran, to engineer complex software systems that are both efficient and scalable.

★★★**In preview** The company, drawing on its experience acquired working with its clients and the best universities, offers in May 2014 new products and services to broaden the reach of its solutions.

In this context, a new innovative service for application profiling will be demonstrated in its beta version at TERATEC, enabling to assist all developers in the diagnostic and performance optimization of their software.

NVIDIA * Stand 7

Press Contact: Stephane QUENTIN, Senior PR Manager France & Belgium
Tel: +33 (0) 1.556.38493 - Mob: +33 6 82 56 83 22 - Mail: squentin@nvidia.com

NVIDIA will be featuring advances in applications and scientific discovery made with GPU-accelerated computing. NVIDIA Tesla K-series GPU Accelerators are based on the NVIDIA Kepler compute architecture and powered by CUDA, the world's most pervasive parallel computing model.

The Tesla K40 GPU Accelerators can run big scientific models on its 12GB of GPU accelerator memory, capable of processing 2x larger datasets and ideal for big data analytics. It also outperforms CPUs by up to 10x with its GPU Boost feature, converting power headroom into user-controlled performance boost.

NVIDIA Tesla K-series GPU Accelerators are used for seismic processing, biochemistry Simulations, weather and climate modeling, image, video and signal processing, computational finance, computational physics, CAE, CFD, and data analytics.

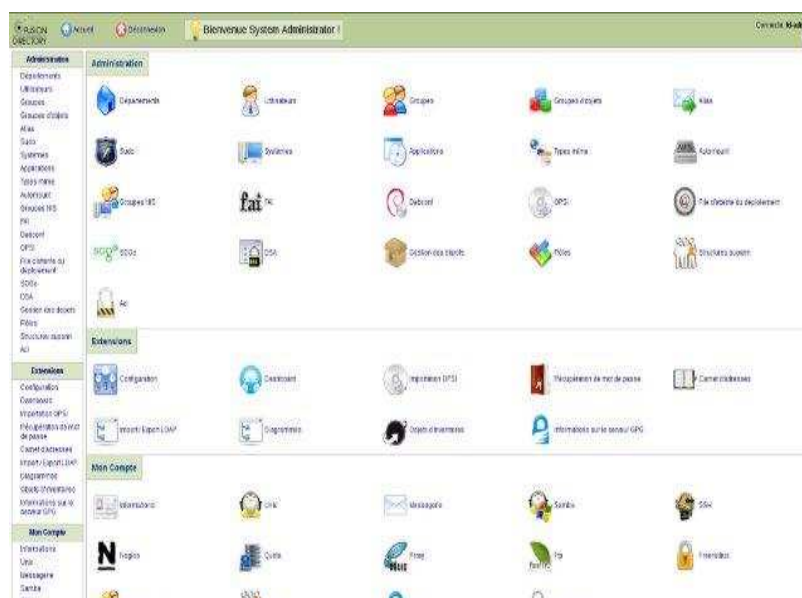


At the NVIDIA booth, there will be opportunities to discuss the latest achievements using GPU technologies in accelerated computing, as well as taking a look at how GPUs are being used to advance scientific discoveries.

A new addition for developers is the new Jetson TK1 Development Kit.



The NVIDIA Jetson TK1 development kit is a full-featured platform for Tegra K1 embedded applications. It allows you to unleash the power of 192 CUDA cores to develop solutions in computer vision, robotics, medicine, security, and automotive. Many of the interesting embedded computing domains that Jetson TK1 aims to support rely on computer vision. Jetson TK1 comes with NVIDIA's new VisionWorks Computer Vision Toolkit pre-installed. VisionWorks is an SDK that provides a rich set of algorithms optimized for NVIDIA CUDA-capable GPUs and SOCs such as Tegra K1, giving you the power to realize CV applications quickly on a scalable and flexible platform.



OPENTEXT * Stand 58

Press contact: Marion MOLINA

Tel: 01 58 18 32 46 / 06 82 92 94 61 - Mail: mmolina@hbcomcorp.com

Exceed onDemand

Exceed onDemand is a thin-client based solution that runs on IBM, HP, and Oracle UNIX/Linux servers. It can dependably and securely connect users to remote UNIX applications over any network using any device.

Exceed onDemand is a thin client-based remote UNIX application access solution that provides you with dependable managed application access from a wide variety of platforms and offers pixel perfect drawing, low cost scalability and trusted security access over any network connection.

Backed by 25 years of experience and a team of industry solution experts, OpenText Exceed onDemand is serving a range of industries, including financial services, electronics and semiconductors, oil and gas, energy and utilities, telecom, computer services, transportation, aerospace and defense, automotive, manufacturing, retail and healthcare.

QUANTUM * Stand 03

Press contact: Julieta DIAZ

Tel: +33 1 41 43 48 63 Mail: julieta.diaz@quantum.com

Quantum's Scale-out Storage Solutions Dedicated to the Scientific Community

From seismic analysis to meteorological data and satellite imaging, scientific and research companies need to process larger and larger sets of data quickly and retain it securely. But they also need to reduce storage costs and protect revenue-generating digital assets. That's why organizations around the world rely on Quantum StorNext 5 data sharing and archiving solutions that will be showcased during Teratec 2014.

Store, Manage and Archive Your Growing Volume of Scientific Data with Quantum StorNext 5

- **Store** – **Quantum StorNext** is the real-time storage enabler that helps scientific and research organizations, like CERN in Switzerland, achieve high-performance digital content ingest.

- **Share** – **Quantum StorNext** provides best-in-class shared file system capabilities to manage and share scientific data across multiple platforms and multiple tiers of storage (disk, object storage, tape and cloud)

- **Archive** – Quantum's **Lattus** next generation object storage platform and **StorNext AEL Archive** solutions help you preserve your scientific and research data from hundreds of TBs to several PBs for decades and more at the best possible costs.

During Teratec 2014, come and discover the new release of StorNext 5, its high-performance file sharing and archiving solutions and the new appliances: next generation object storage platform, high-performance disk storage systems and our large portfolio of LTO tape libraries for long term retention.



SCILAB ENTERPRISE * Stand 57


Press contact: Emilie DAVID

Tel: 01 80 77 04 60 - Mail: emilie.david@scilab-enterprises.com

Pushing HPC in Scilab : 2 major innovations to improve the use of High Performance Computing with Scilab have been developed.

- **First, the implementations of the MPI** (Message Passing Interface) API for parallel computation. Scilab/MPI manages most Scilab types and is included in new Scilab 5.5.0. It is interesting to use Scilab/MPI in most of general cases.

- **Secondly, the sciGPGPU module for using the GPU** (Graphics Processor Unit) of the computer, in particular for managing large datasets. It can be automatically installed from Scilab using ATOMS module manager, manages OpenCL and CUDA kernels, provides a number of GPU based Scilab functions and maps many CUBLAS / CLBLAS and CUFFT / CLFFT functions. These innovations, developed in H4H project, improve the using of Scilab Software and have been used by Repsol and Dassault Aviation.



Scilab is free and open source software for numerical computation providing a powerful computing environment for engineering and scientific applications. A large number of functionalities is included in Scilab and it is used in a wide spectrum of areas: space, aeronautics, automotive, energy and so on.

Scilab is published by Scilab Enterprises that provide advice, support and services around Scilab software to companies. Scilab Enterprises is committed to providing expertise to industrial companies in the implementation of optimal solutions in terms of technological innovation with computation and simulation tools.

SOGETI HIGH TECH * Stand 32

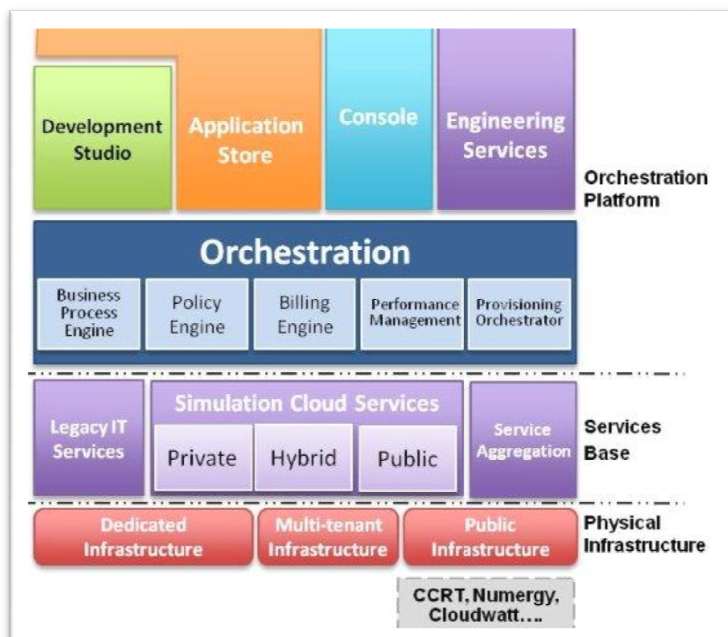
Press Contact: Philippe MÉLÉARD, Directeur Marketing Communication

Tel : +33 1 40 93 73 23 - Email : philippe.meleard@sogeti.com - www.sogeti-hightech.fr

Sogeti High Tech is an international group in Engineering and Technology Consulting Services. Major player in Simulation, for more than 20 years, Sogeti High Tech accompanies market leaders from Aeronautic, Energy, Railway and health sectors in their competitiveness improvement.

Sogeti High Tech's Simulation offering fit naturally into an environment of High Performance Calculation for its main components:

Virtual Testing enables to sharply decrease the number of real tests in favor of simulated tests which reduce conception costs and accelerate development cycle by anticipating the validation phase.



The **HPC** component for scientific softwares ensures an optimal use of HW resources by the software while ensuring their sustainability and scalability.

As reference information systems integrator dedicated to simulation, Sogeti High Tech with its **Simulation Store** project, propose to users a transparent and secured access to numerical simulation key services and offer to scientific code users an online experience that increase their innovation capacity with current technological standards.

Through Simulation Store, we redefine access and use conditions of numerical simulation, for now on reserved to firms enabled to free themselves from cost (equipment, software, training).

SYSFERA * Stand 14

Press Contact: David LOUREIRO

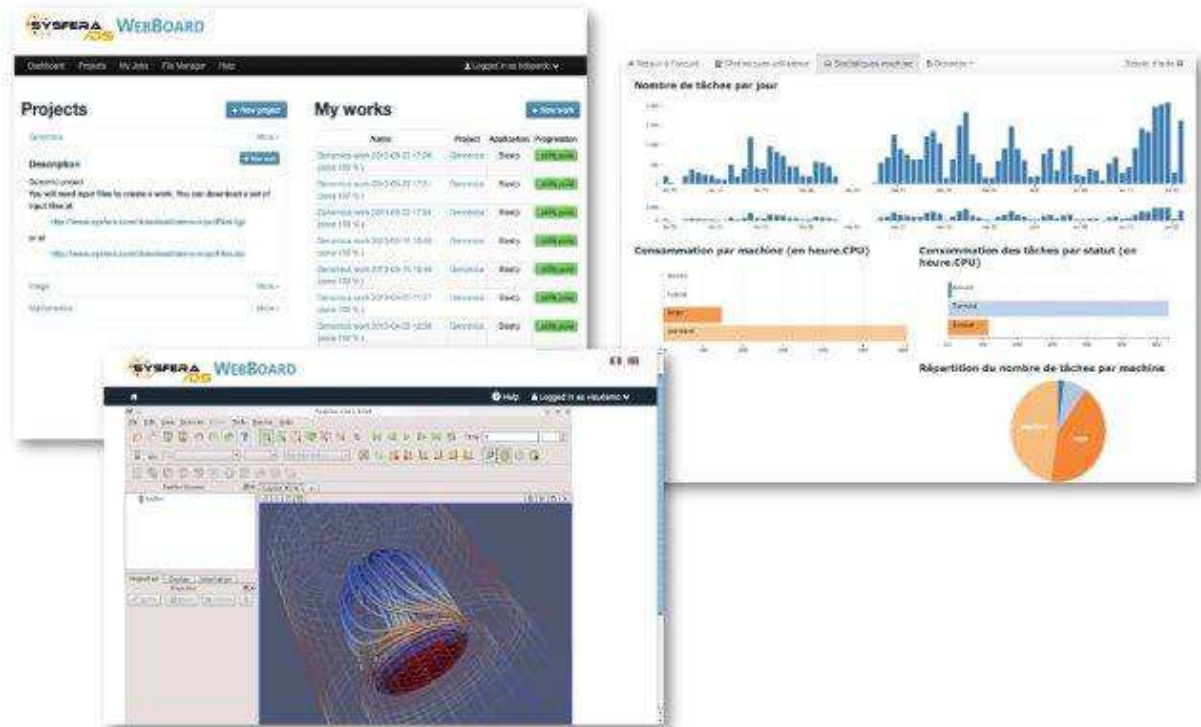
Tel: 04 81 76 28 90 - Mail: david.loureiro@sysfera.com

SysFera is a fast growing independent software vendor building innovative software that enable supercomputing centers to make their applications and resources available and billable with fewer support staff and help large companies to optimize their HPC ecosystem to provide a better service to users.

***** in preview : SysFera is proud to announce at Ter@Tec the new release SysFera-DS v4.5.**

The release focuses mainly on a native and improved support of GPU for visualization on HPC computing resources and Amazon Web Services platform, administration and support simplifying the installation phase but also infrastructure management ensuring an full installation in less than 2 hours. This release thus allows SysFera partners to efficiently perform the distribution and the support of the solution outside of France.

Focusing primarily on ease of use and integration with existing system tools, SysFera's web portal makes it simple and fast to run both non-interactive and interactive graphical HPC applications on traditional servers, clusters, or cloud environments – all via a simple web browser. The product requires no modification by application vendors, no changes to the HPC system security, and no local software installation. The software allows non-computer scientists to become productive more quickly focusing on their core expertise; Project managers to spend less time administering applications, users, projets and resources; and IT departments to bill customers more accurately and with less effort.



UNIVERSITAT POLITECNICA DE VALENCIA * Stand 23

Contact name: José FLICH, Associate Professor

Phone: +34 963877007 - jflich@disca.upv.es

Multicore systems are the enablers for High-Performance Computing of tomorrow. Multicore design and validation is a tough process which may severely slow down adoption of new and advanced architectures. Simulation processes of HPC systems inherit large design loops which impede the proper advance of research.

UPV provides an open platform for future multicore design and validation with FPGA-based modern tools and platforms, enabling fast and accurate implementations, reaching solutions with hundreds of cores. Partitioning and isolation of concurrent applications are enabled through the PEAK architecture demonstrated in the forum. These capabilities, exhibited at the on-chip network level leverage the management of a hundred-core system and enable optimum chip usage through truly capacity computing.

UPV research in multicore systems and HPC interconnects provides valid solutions for early adoption in new designs, spanning from interconnect designs and reaching memory hierarchy implementations. UPV has been involved and is currently coordinating several EU projects related with multicore designs. vertical project aims to provide full virtualization capabilities to embedded systems. UPVs contribution lies on the design of partition capabilities for multicore systems, providing security aspects.

XYRATEX * Stand 53

Press contact: Wendy MADIEROS

Tel: +1 510-687-5253 - Email: wendy_madieros@xyratex.com

Xyratex, a Seagate Company, is the largest OEM provider of enterprise data storage products including the award-winning ClusterStor™ family of high performance scale-out storage solutions. ClusterStor's software architecture tightly integrates the Lustre® filesystem with Xyratex's proven enterprise quality storage providing the first truly engineered HPC storage solution.

The ClusterStor family delivers unmatched availability, ease of use, price-performance and overall TCO leadership for HPC, Big Data and Cloud.

For applications where users need storage that scales both in capacity and performance ClusterStor supports high performance data processing of relatively small data quantities to some of the largest supercomputers in the world.

Please visit our stand #53 and we will be delighted to discuss how ClusterStor can also help your organization accelerate discoveries or visit www.xyratex.com to learn more.

ooo



Exhibitors'list (dated June 2nd 2014)

(In bold, newcomers)

ACTIVEEON ■ ADVANCITY ■ ALINEOS ■ ALLINEA SOFTWARE ■ ALTAIR ENGINEERING ■ AMD ■ ANSYS ■ AS+ Groupe
EOLEN ■ BARCO ■ BULL ■ **BUREAU 14** ■ CAP DIGITAL ■ CARRI SYSTEMS ■ CCI DE L'ESSONNE ■ CEA ■
CLUSTERVERSION ■ **COMSOL France** ■ **CRAY** ■ CS SI ■ **CYBELETECH** (Pépinière Teratec) ■ DDN DATADIRECT
NETWORKS ■ DELL ■ EAGOCOM (Pépinière Teratec) ■ **EMG2 / BITTWARE** ■ ENGINSOFT ■ ESI GROUP ■ **ETP4HPC**
■ EUROTECH ■ EXASCALE COMPUTING RESEARCH LAB ■ **EXTREME NETWORKS** ■ FUJITSU ■ GENCI ■ HP ■ IBM ■
IFPEN - IFP ENERGIES NOUVELLES ■ INRIA ■ INTEL ■ KALRAY ■ MATHWORKS ■ MEDICEN ■ NAFEMS ■ NICE
SOFTWARE ■ NOESIS SOLUTIONS ■ **NUMSCALE** ■ NUMTECH (Pépinière Teratec) ■ NVIDIA ■ OPENSIDES ■
OPENTEXT ■ OXALYA OVH ■ PANASAS ■ PEPINIERE TERATEC ■ QUANTUM ■ ROGUE WAVE SOFTWARE ■ SCILAB
ENTERPRISES ■ SGI ■ SILKAN ■ SOGETI HIGH TECH ■ SYSFERA ■ SYSTEMATIC ■ TERATEC ■ TOTALINUX ■
TRANSTEC ■ **UNIVERSITAT POLITECNICA DE VALENCIA** ■ **XYRATEX**

NEXT RENDEZ-VOUS: Teratec Forum 2014



July 1st and 2nd, 2014 - Ecole Polytechnique, Palaiseau, France

To find out more on [TERATEC Forum](#)

PRESS Contact:

Colette REY - C&REY COMMUNICATION
Tel: +33 (0)9 51 70 20 57
Mob: +33 (0)6 14 73 97 43
Mail: colette.rey@c-reycom.com

TERATEC Contact:

Jean-Pascal JEGU
Tel: +33 (0)9 70 65 02 10
Mob: + 33 (0)6 11 44 49 59
Mail: jean-pascal.jegu@teratec.fr