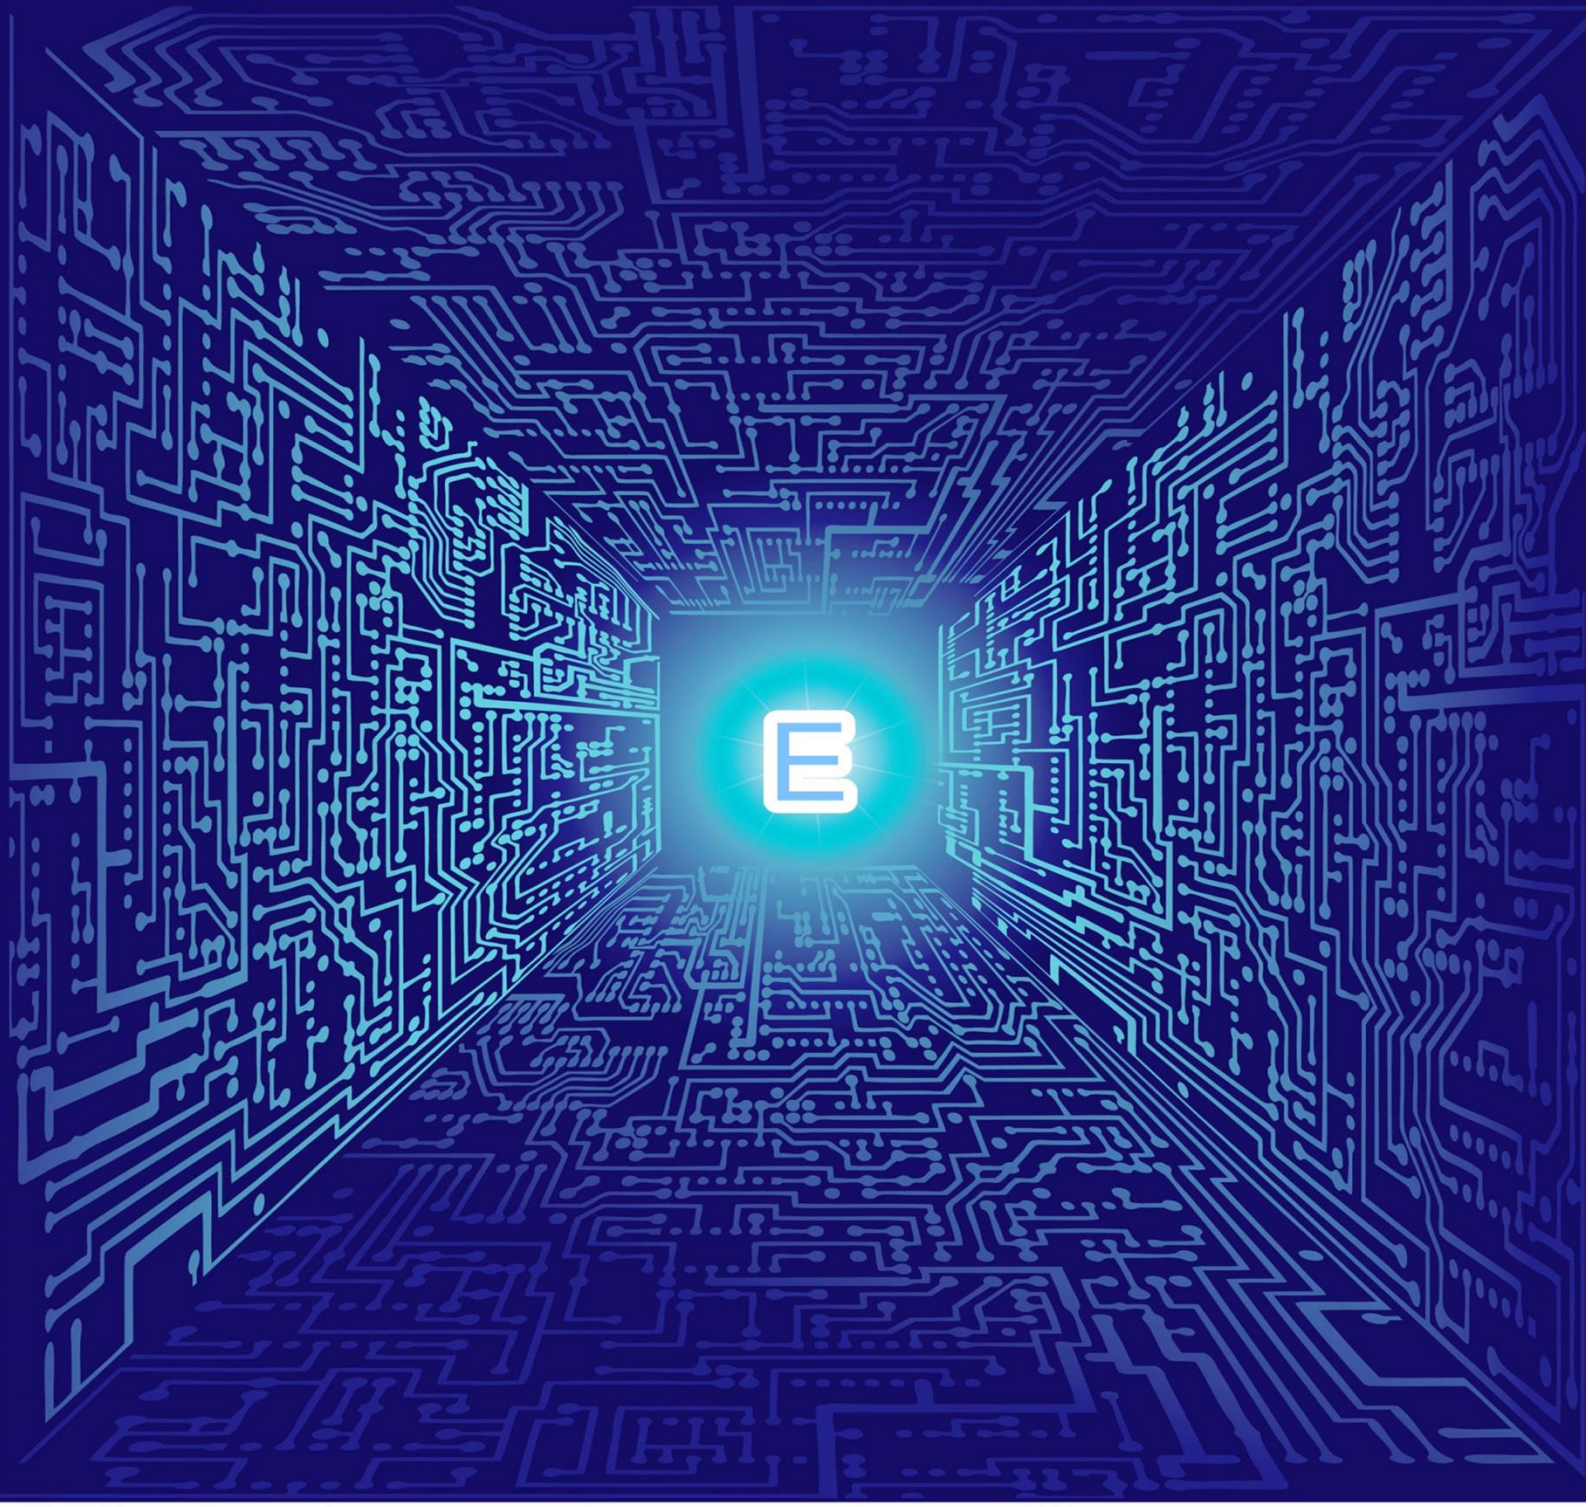

Chinese-French Workshop on Extreme Computing

PROGRAM

May 9- May 10, 2016, Guangzhou
National Supercomputer Centre in Guangzhou

Transportation

Date	Time	From	To	Vehicle
9th May	07:50-08:50	Sofitel	NSCC-GZ	*2 Buick BPVs (Plate: TBD)*
9th May	18:45-19:00	NSCC-GZ	Southern Club Hotel	*2 Buick BPVs (Plate: TBD)*
9th May	20:30-21:30	Southern Club Hotel	Sofitel	*2 Buick BPVs (Plate: TBD)*
10th May	08:30-09:30	Sofitel	NSCC-GZ	*2 Buick BPVs (Plate: TBD)*
10th May	11:30-11:45	NSCC-GZ	Southern Club Hotel	*2 Buick BPVs (Plate: TBD)*
10th May	13:30-14:30	Southern Club Hotel	Sofitel	*2 Buick BPVs (Plate: TBD)*

Meals

9th May	Lunch	Campus Canteen of Sun Yat-sen University
	Welcome Banquet	Southern Club Hotel
10th May	Lunch	Southern Club Hotel

1. ***Miss Huizhong Lu (Phone: 15989082**5) from NSCC-GZ will meet our French speakers in the hotel lobby at 7:50 on May 9th, and 8:30 on May 10th, the BPVs will have "CFWEC16" logos on them.***
2. ***Please let us know if you have any further request in terms of transportation and food, the contact person from NSCC-GZ is Miss Qian Mo (Phone: 15602391**5, E-mail: moqian@nscc-gz.cn).***

Agenda for Chinese-French Workshop on Extreme Computing (CFWEC2016)

9th-10th May, 2016, Guangzhou, China

May 9th			
9:00~9:30	Welcome Addresses	Abdo Malac & Depei Qian	French Embassy & SDCS, SYSU
09:30~10:00	Programming Paradigm and Intelligent Linear Algebra for Future Extreme Computing	Serge G. Petiton	U. Lille 1, CNRS
10:00 ~10:30	Convergence of HPC and Bigdata	Yutong Lu	NSCC-GZ
10:30-11:00	**Tour of Tianhe-2 Supercomputer System, Group Photo, Tea Break**		
11:00~11:30	HPC & Big Data: Examples and Perspectives in CNRS and Beyond	Michel Dayde	CNRS
11:30~12:00	Heterogeneous Domain Decomposition Algorithms with Applications on Tianhe-2	Chao Yang	ISCAS
12:00~12:30	On the Path to Exascale from the Application Point of View	Christophe Calvin	CEA, DRF
12:30-13:30	**Lunch at Campus Canteen, SYSU**		
14:00~14:30	Applications of CFD on Tianhe-2	Hui Yan	NSCC-GZ
14:30~15:00	BigData Processing in Dynamic Hybrid Computing Environment Using MapReduce	Haiwu He	CAS
15:00~15:30	Deep Learning Applications on Multiple GPUs	Liang Lin	SDCS, SYSU
15:30~16:00	Challenges of Application Performance	Xavier Vigouroux,	ATOS-Bull
16:00-16.30	Design a Suitable Program Frame to Minimize I/O Cost for Large Scale Simulations on Tianhe2	Jingkun Chen	NSCC-GZ
16:30~16:45	**Tea Break**		
16:45~17:15	Damaris: Jitter-Free I/O Management and In Situ Visualization of HPC Simulations using Dedicated Cores	Gabriel Antoniu	INRIA
17:15~17:45	Applications of Computational Biology and Precision Medicine on Tianhe-2	Jiahui Li	NSCC-GZ

17:45~ 18:15	HPC in Total Exploration and Production	Diego Klahr	TOTAL
18:15 ~ 18:45	New Software and Hardware Architecture Solutions for HPC and BigData	Thierry Collette	CEA
19:00~20:30	**Welcome Banquet at Southern Club Hotel**		
May 10th			
9:30~11:30	Discussion		
11:45-13:30	Lunch at Southern Club Hotel		

Participants

Abdo Malac

French Embassy scientific officer deputy

Serge G. Petiton

U. Lille 1, and CNRS-Maison de la simulation@saclay

Prof. Serge G. Petiton received the M.S. in Applied Mathematics, the Ph.D. degree in computer science, and the “Habilitation à diriger des recherches”, from Pierre and Marie Curie University, Univ. PARIS 6. He was post-doc student, registered at the graduate school, and junior researcher scientist at the YALE University, 1989-1990. He has been researcher at the “Site Experimental en Hyperparallelisme” (supported by CNRS, CEA, and French DoD) from 1991 to 1994. He also was affiliate research scientist at YALE and visiting research fellow in several US laboratories, especially in NASA-ICASE and the AHPCRC during the period 1991-1994. Since then, Serge G. Petiton is Professor at the University of Lille, Sciences and Technologies. Serge G. Petiton was, and is, P.I. of several international projects with Japan, Venezuela and Germany. Since 2012, Serge G. Petiton has a half time CNRS senior position at the “Maison de la Simulation” in Saclay.

Serge G. Petiton has been scientific director of more than 22 Ph.D.s and has authored more than 100 articles on international journals and conferences. His main current research interests are in “Parallel and Distributed Computing”, “Post-Petascale Auto/smart-tuned Dense and Sparse Linear Algebra”, and “Language and Programming Paradigm for Extreme Modern Scientific Computing”; targeting especially geoscience and big data applications.

Serge G. Petiton is a member of SIAM, ACM, IEEE, the YALE club of France and the Ivy Plus European leaders club.

Talks:

Programming Paradigm and Intelligent Linear Algebra for Future Extreme Computing

Abstract: Exascale *hypercomputers* are expected to have highly hierarchical architectures with nodes composed by *lot-of-core* processors and often accelerators. Methods have to be redesigned and new ones introduced or rehabilitated to optimize communication, data distribution, resilience and energy consumption.

The different programming levels will generate new difficult algorithm issues. New language and framework should be defined and evaluated with respect to modern state-of-the-art of scientific methods. We propose YML (yml.prism.uvsq.fr), associated with a multilevel programming paradigm. YML with its high level language permits to automate and delegate the managements of dependencies between loosely coupled clusters of processors. Besides, the tightly coupled processors inside each cluster could be programmed through a PGAS language such as XMP. Thanks to the component-oriented software architecture of YML, it is relatively easy to integrate new components such as numerical libraries, encapsulated XMP programs; and scientific end-users may give expertise to help smart tuning at runtime of numerical and others parameters.

In this talk, we present this multilevel programming paradigm for exascale computing and propose our approach based on YML. We present the Block Gauss-Jordan method to invert dense matrices, and the Multiple Explicitly Restarted Arnoldi Method (MERAM) to compute eigenvalues of sparse matrices as study cases. Experimental results are obtained on Japanese “K” and “T2K” supercomputers, on the French Grid5000 platform, and on the “Hooper” supercomputer in LBNL.

We explain how such programming paradigm is well-adapted for future intelligent linear algebra required for extreme computing, as discussed on my keynote talk at HPC China 2016.

Michel Dayde

Centre National de la Recherche Scientifique (CNRS)

Michel J. Daydé received his Ph.D. from Institut National Polytechnique de Toulouse (France) in 1986 in Computer Science. From 1987 to 1995, he was a postdoctorate fellow then visiting Senior Scientist in the Parallel Algorithms Group at CERFACS. From 1988, he is Senior Lecturer and from 2000 Professor at "Ecole Nationale Supérieure d'Electrotechnique, d'Electronique, d'Informatique, d'Hydraulique et de Télécommunications" (ENSEEIHT) from Institut National Polytechnique de Toulouse presently in the Networking and Telecommunication Department.

Since 1996, he is Research Director in the "Algorithmes Parallèles et Optimisation" Group at Institut de Recherche en Informatique de Toulouse (IRIT). He was vice-chair of the Euro-Par'99 conference held in Toulouse, and Chair of the Steering Committee of the VECPAR Conference series in addition to several other conference committees. He is now Head of IRIT (Computer Science Research Institute of Toulouse) which is one the largest laboratory associated with CNRS in France.

He is Scientific Representative at the CNRS Institute for Information Sciences and Technologies (INS2I) in charge of HPC, Grid and Cloud. He is also the Director the Orientation Committee for Scientific Computing within CNRS (COCIN) and is involved in various committees related to research (e-) infrastructures at CNRS and national levels.

His current research interests are in parallel computing, computational kernels, full linear algebra, preconditioners for iterative solvers, large scale nonlinear optimization on high-performance computers, and cloud and grid computing. He has also been involved in porting industrial codes on large scale computing infrastructures.

Talk:
HPC & Big Data: Examples and Perspectives in CNRS and Beyond

Abstract: “Big Data” is considered to be the fourth pillar of science nowadays and, as High Performance Computing for modelling and numerical simulation, is crucial for science and industrial competitivity.

The volume and the complexity of data do not stop growing and, in several areas, the volume and the complexity of data challenge our capacities to explore and to analyze them.

It is not anymore possible to dissociate HPC and “Big Data” i.e. the exploitation of the data arising from digital simulations (climate, combustion, fusion, astrophysics), large instruments (LHC, ITER, LSST, LOFAR, genomic platforms), ground or space systems of observation (seismology and geodesy RESIF, Euclid, WFIRST, GAIA, imaging and interferometry) or simply multiple devices of data acquisition (broadband sequencer, sensors' networks, social networks, etc.). Processing such large amounts of data requires both significant processing capabilities and suitable methods for

data analysis at large scale.

In multiple scientific and socioeconomic domains, the volumetry and the variety of the existing data as well as the time constraints of calculation revealed new challenges. In a wide range of areas: in fundamental sciences (physics of high energy, fusion, earth and universe sciences, bio-computing, neurosciences, etc.), digital economy (business intelligence, Web, e-commerce, social networks, e-government, health, telecommunications and media), ground and air transport, financial markets, environment (climate, natural risks, energy resources, smart cities, connected house), security, industry (smart industry, customized products, design and production chain), new tools, scientific methods and new technologies are necessary.

We will illustrate some of these issues using examples coming from the IRIT research laboratory and from CNRS.

Christophe Calvin

(Commissariatà l'énergie atomique et aux énergies alternatives (CEA) /DRF)

Dr. Calvin received his Ph.D. in Applied Mathematics from INPG – Grenoble - before joining CEA in 1995. He received his "Habilitation à Diriger des Recherches" (HDR grade) in Computer Sciences in 2013.

He started his career as software architect for a 3D parallel CFD code in CEA Grenoble. He moved to CEA Saclay to lead a project with EDF and AREVA concerning the development of new generation of simulation code for neutron transport. In 2006, he led a laboratory of more 20 people devoted to the development of computer codes for Reactor Physics. From 2012 to 2014 is acted as "Chargé de Mission" in charge of setting a roadmap for the use and development of HPC in the Nuclear Energy Division Simulation Program. From 2014 to 2015 he led two projects for the development and the validation of thermal-hydraulic codes, uncertainties processing and HPC. Since end of 2015, Dr. Calvin coordinates Numerical Simulation and HPC R&D for the Fundamental Research Division of CEA.

Dr. Calvin has been involved in the European PRACE project since its beginning. He has co-led FP3C French-Japanese research project for post-petascale programming for 4 years, and he is the technical coordinator of CEA/DEN-DOE HPC collaboration. Since 2009 he is a CEA senior expert for "Mathematics, Computer Science – software technologies – architecture and algorithmic – High Performance Computing applied to Reactor Physics and Thermal hydraulic". Dr. Calvin is a member of the Maison de la Simulation Scientific Council and co-holder of the chair of excellence of "Maison de la Simulation" on Computational Science and Scientific Computing.

His main research activities concern the study and the development of highly scalable numerical algorithms for nuclear engineering applications. Dr. Calvin has published more than 30 papers in Journal and International conferences.

Talk:

On the Path to Exascale from the Application Point of View

Abstract: Even if the detailed features of the 1st exascale system are not known, we already have main of them. These main features have a deep impact on today large scientific applications in order to obtain significant performances on these future exascale systems. Numerical simulation is definitively a key tool at Fundamental Research Division of CEA to conduct state of the art research in many different fields of fundamental physics, astrophysics, climate modelling, high energy physics, fusion energy, medical imaging, genomics … These simulations are conducted using large scientific applications which have to be adapted to be efficient on future exascale systems.

In this talk, we will illustrate using some examples the today use of HPC (High Performance

Computing) and TC (High Throughput Computing) at CEA/DRF and some R&D axes to prepare applications to be efficient on future exascale systems.

Xavier Vigouroux

ATOS-Bull

Dr Xavier Vigouroux [M], after a Ph.D. from Ecole normale Superieure de Lyon in Distributed computing, worked for several major companies in different positions. He has now been working for bull for 9 years. He led the HPC benchmarking team for the first five years, then in charge of the "Education and Research" market for HPC at Bull, he is now managing de "Center for Excellence in Parallel Programming" of Bull

Talk:

Challenges of Application Performance

Abstract: Constraints coming from power consumption is putting a lof of pressure on all the actors implied in the HPC. As a consequence, from the silicon up to the datacenter, through the application writer, impacts can be already detected. In this presentation, Xavier VIGOUROUX will expose the consequences of that new constraints on an integrator as Bull (group atos) : System designs, software stack, applications ....

Gabriel Antoniu

INRIA

Gabriel Antoniu is a Senior Research Scientist at Inria/IRISA, Rennes. He leads the KerData research team, focusing on storage and I/O management for Big Data processing on scalable infrastructures (clouds, HPC systems). He received his Ph.D. degree in Computer Science in 2001 from ENS Lyon. He leads several international projects in partnership with Microsoft Research, IBM, Argonne National Lab, the University of Illinois at Urbana Champaign. He served as Program Chair for the IEEE Cluster 2014 conference and regularly serves as a PC member of major conferences in the area of HPC, cloud computing and Big Data (SC, HPDC, CCGRID, Cluster, Big Data, etc.). He has acted as advisor for 17 PhD theses and has co-authored over 100 international publications in the above-mentioned areas.

Talk:

Damaris: Jitter-Free I/O Management and In Situ Visualization of HPC Simulations Using Dedicated Cores

Abstract: Large scale simulations running on leadership class supercomputers generate massive amounts of data for subsequent analysis and visualization. As the performance of storage systems shows its limits, an alternative consists in embedding visualization and analysis algorithms within the simulation code (in situ visualization). In this context, we present the benefits of using Damaris, a middleware for I/O forwarding and post-processing using dedicated cores to offload in situ visualization while sharing resources with the running simulation. Damaris fully hides the I/O variability as well as all I/O-related costs, which makes simulation performance predictable; it increases the sustained write throughput by a factor of up to 15 compared with standard I/O approaches; it allows almost perfect scalability of the simulation up to over 9,000 cores, as opposed to state-of-the-art approaches that fail to scale; through its extension, Damaris/Viz, it enables a seamless connection to the VisIt visualization software to perform in situ analysis and visualization in a way that does not impact the performance of the simulation, nor its variability. Damaris/Viz was evaluated with the CM1 atmospheric simulation on Grid'5000 and on NCSA's Blue Waters with up to 6400 cores, and with the Nek5000 CFD simulation.

Diego Klahr

TOTAL

Diego Klahr has been working in Total Exploration and Production since 2011 as an HPC infrastructure advisor where he manages the choice and the deployment of one of the biggest HPC system in the industry. Before this, he worked in IFPEN (French oil and gas research center) for 10 years where he was in charge of scientific computing engineering. He began as a consultant in HPC software development and optimization.

Talk:
HPC in Total Exploration and Production

Abstract: The extraordinary challenge that the Oil and Gas industry must face for hydrocarbon exploration and production requires the development of leading edge technologies. Seismic reflection and reservoir simulation are two of these essential and strategic technologies for O&G.

Seismic Depth imaging progress depends mainly on 3 factors: Seismic data acquisition recording more accurate and more complex wave-fields, optimized workflow from processing to geosciences integration and new numerical algorithms accommodating more physics and new imaging technology.

Production optimization needs faster and more accurate reservoir simulations including avoiding upscaling between geological and dynamic model and an improved reservoir to surface coupling.

High Performance Computing is a key component of this evolution. The recent emergence of PetaFlop systems and the perspective of Exaflop systems by end of this decade offer to our industry very interesting advancement and challenges to move ahead. However, developing those new advanced methods and anticipating the corresponding computing resources require clearly to identify and understand their impact in terms of high performance computing algorithms, data science management, data centre architecture and optimization.

Haiwu He

Chinese Academy of Sciences (CAS)

Prof. Haiwu HE received his M. Sc. and the Ph. D. degrees in computing from University of Sciences and Technologies of Lille, France, respectively in 2002 and 2005. He was a postdoctoral researcher at INRIA Saclay, France in 2007. He was a research engineer expert at INRIA Rhone-Alpes in Lyon, France from 2008 to 2014. Currently, he is 100 Talents Professor in CNIC (Compter Network Information Center), CAS. He is "Chunhui Schloar" of Ministry of Education of China from 2013. He has published about 30 refereed international journal and conference papers. His research interest covers BigData, Cloud computing, HPC, Machine Learing etc.

Talk:

BigData Processing in Dynamic Hybrid Computing Environment Using MapReduce

Abstract： A novel MapReduce computation model in hybrid computing environment called HybridMR is proposed in the talk. Using this model, high performance cluster nodes and heterogeneous desktop PCs in Internet or Intranet can be integrated to form a hybrid computing environment. In this way, the computation and storage capability of large-scale desktop PCs can be fully utilized to process large-scale datasets. HybridMR relies on a hybrid distributed file system called HybridDFS, and a time-out method has been used in HybridDFS to prevent volatility of desktop PCs, and file replication mechanism is used to realize reliable storage. A new node priority-based fair scheduling (NPBFS) algorithm has been developed in HybridMR to achieve both data storage balance and job assignment balance by assigning each node a priority through quantifying CPU speed, memory size and I/O bandwidth. Performance evaluation results show that the proposed hybrid computation model not only achieves reliable MapReduce computation, reduces task response time and improves the performance of MapReduce, but also reduces the computation cost and achieves a more environment friendly computing mode.

Thierry Collette

Commissariatà l'énergie atomique et aux énergies alternatives (CEA)

Since 2011, Thierry Collette is CEA Leti &List VP Division in charge of technological development for embedded systems and integrated components. The division deals with the hardware and software design and integration for solutions addressing the IoT (Internet of Things), the Transport, the Health, the Energy and the Security, cooperates (Industrial transfer) with more than 60 companies and publishes more than 200 scientific papers a year and 50 patents a year.

Previously, he was the deputy director of CEA List (www-list.cea.fr) in charge of programs and strategy. From 2004 to 2009 he managed the Architectures and Design Unit at CEA, unit of 100 researchers, focused on embedded computing architectures, embedded reliability and computer vision. Before that, he was the head of embedded computer laboratory at CEA and has designed several parallel and reconfigurable embedded computers. He has obtained an Electrical Engineering Degree in 1988 and a Ph.D in Microelectronics of the University of Grenoble in 1992.

He contributed to the creation of 5 startups of the CEA, ActiCM in 2000 (bought by CRAFORM - www.creaform3d.com/fr), Kalray in 2008 (www.kalray.eu), Arcure in 2009 (www.arcure.net), Kronosafe in 2011 (krono-safe.com) and WinMs in 2012 (http://www.win-ms.com/).

He wrote, as author and co-author, several papers in conferences and journals on technologies for embedded parallel and reconfigurable computing and holds several patents too. He teaches computer architectures in master degree at Ecole Centrale of Paris and University of Paris XI. He is expert CEA senior and had evaluated several international and national projects and is member of evaluation committee of the French National Research Agency.

Talk:
New Software and Hardware Architecture Solutions for HPC and BigData

Abstract: The exponential growth of information required for fine grain simulations and visualization or generated by humans and objects (BigData), creates an explosion of digital data and the need for fast data processing and analysis while keeping a constant energy envelope and cost of ownership. This will cause a paradigm shift for HPC and BigData infrastructures to move from a tight focus on performance towards energy-efficiency and total cost of ownership. The components of future HPC and BigData infrastructures and their integration into a full system should be reconsidered. This will be made possible with the emergence of new domain specific highly parallelizable software and hardware, and more generic solutions to increase the compute density and energy efficiency, all these solutions minimizing the cost of data movement and optimizing the development cost of new Integrated Circuit design and manufacturing.

Depei Qian

School of Data and Computer Science, Sun Yat-sen University

Depei Qian, professor at Sun Yat-sen University and Beihang University, dean of the School of Data and Computer Science of Sun Yat-sen University.

He has been working on computer architecture and computer networks for many years. His current research interests include high performance computer architecture and implementation technologies, distributed computing, network management and network performance measurement. He has published over 300 papers in journals and conferences.

Since 1996 he has been the member of the expert group and expert committee of the National High-tech Research & Development Program (the 863 program) in information technology. He was the chief scientist of three 863 key projects on high performance computing since 2002. Currently, he is the chief scientist of the 863 key project on high productivity computer and application service environment.

Chao Yang

Institute of Software, Chinese Academy Sciences (ISCAS)

Chao Yang is a full professor at the Laboratory of Parallel Software and Computational Sciences (LPSCS), Institute of Software, Chinese Academy Sciences (ISCAS). He received his PhD in computer science from ISCAS in 2007. His research interests include numerical analysis and modeling, large-scale scientific computing, and parallel numerical software.

Talk:

Heterogeneous Domain Decomposition Algorithms with Applications on Tianhe-2

Abstract: Nowadays, heterogeneous architecture has become a competitive choice to build supercomputing systems that run at desired computational rates and fit within an affordable power envelope. Such systems are equipped with both multi-core CPUs and many-core accelerators; examples include Tianhe-2, which is comprised of both Intel Xeon CPUs and Intel Xeon Phi coprocessors, and Titan, which consists of both AMD Opteron CPUs and NVIDIA Tesla K20x GPUs. The development of highly scalable algorithms to adapt with various heterogeneous architectures is becoming an urgent demand. In this talk, I will introduce some of our recent efforts in designing highly scalable domain decomposition algorithms for domestic heterogeneous supercomputers in China. These algorithms, though may come from different application scenarios, share a common feature that special considerations are kept in mind to maintain the load balance among different processing units, reduce the data movement between heterogeneous resources, and perform vectorization and/or parallelization in the many-core accelerator. The proposed algorithms have been applied in accelerating global atmospheric simulations and the High-Performance Conjugate Gradients (HPCG) benchmark and have shown to be scalable to millions of processor cores on leading-edge heterogeneous supercomputers such as the Tianhe-2.

Liang Lin

School of Data and Computer Science, SYSU

Liang Lin is a Professor with the School of Data and Computer Science, Sun Yat-sen University (SYSU), China. He received the B.S. and Ph.D. degrees from the Beijing Institute of Technology (BIT), Beijing, China, in 1999 and 2008, respectively. From 2006 to 2007, he was a joint Ph.D. student with the Department of Statistics, University of California, Los Angeles (UCLA). His Ph.D. dissertation was nominated by the China National Excellent PhD Thesis Award in 2010. He was a Post-Doctoral Research Fellow with the Center for Vision, Cognition, Learning, and Art of UCLA.

Prof. Lin's research focuses on new models, algorithms and systems for intelligent processing and understanding of visual data. He has authorized or co-authorized more than 100 papers in top tier academic journals and conferences. He has served as an associate editor for IEEE Trans. Human-Machine Systems, Neurocomputing and The Visual Computer, and a guest editor for Pattern Recognition. He was supported by several promotive programs or funds for his works, such as "Program for New Century Excellent Talents" of Ministry of Education (China) in 2012, and Guangdong NSFs for Distinguished Young Scholars in 2013. He received the Best Paper Runners-Up Award in ACM NPAR 2010, Google Faculty Award in 2012, and Best Student Paper Award in IEEE ICME 2014.

Talk:

Deep Learning Applications on Multiple GPUs

Abstract: Over the last couple of years, deep learning techniques have made tremendous progress in many fields such as computer vision, natural language processing, and robotics. Deep learning is a family of methods that uses deep neural networks to learn different levels of feature representation. These emerging techniques have effectively improved the system performances to achieve the goals against real challenges, and most importantly, these deep learning methods fully take advantage of the large amount of data (e.g., collected via internet and social media) and powerful computing resources. On the other hand, the GPU-based high performance computing has been fast growing and widely applied in machine learning and data mining. This talk will review the recent advances of deep learning applications, particularly for image recognition and understanding, and introduce several common deep learning implementations based on multiple GPUs.

Yutong Lu

National Supercomputer Center in Guangzhou, SYSU

Professor Yutong Lu is the Director of National Supercomputer Center in Guangzhou, and she is professor of School of Computer Science, National University of Defense Technology (NUDT), Changsha and Sun Yat-san University in Guangzhou. She got her B.S, M.S, and PhD degrees from the NUDT. Her extensive research and development experience has spanned several generations of domestic supercomputers in China. Prof. Lu is deputy chief designer of Tianhe Project. Her continuing research interests include parallel operating systems (OS), high-speed communication, global file system, and advanced programming environment.

Talk:

Convergence of HPC and Bigdata

Abstract: Nowadays, advanced computing and visualizing tools are facilitating scientists and engineers to perform virtual experiments and analyze large-scale datasets. Computing-driven and Bigdata-driven scientific discovery has become a necessary approach in global environment, life science, nano-materials, high energy physics and other fields. Furthermore, the fast increasing computing requirements from economic and social development also call for the birth of the Exascale system. This talk will discuss the convergence of the HPC and Bigdata on Tianhe2 system.

Hui Yan

National Supercomputer Center in Guangzhou, SYSU

Hui Yan is currently a high-performance computation engineer of National Supercomputer Center in Guangzhou. He obtained his PhD in Mechanical Engineering from Florida Institute of Technology, where his main research is on computational material science and parallel computing. Since joining the NSCC-GZ in 2015, he has been focusing on the application of computational fluid dynamic (CFD) on Tianhe-2.

Talk:

Applications of CFD on Tianhe-2

Abstract: Computational fluid dynamics has been applied in a wide range of science and engineering field, which overcomes the limitation of real experiments. Various software and algorithms have been developed. National Supercomputer Center in Guangzhou（NSCC-GZ）has been committed to promote and optimize the application of CFD on Tianhe-2. In this talk, a whole work flow based on OpenFOAM is proposed. It is applied to analyze hydrodynamic involved in shipbuilding industry. Besides, an automatic mesh partition tool is introduced to enhance the performance of FDS on Tianhe-2, which is utilized to simulate fire-driven fluid flow inside buildings.

Jingkun Chen

National Supercomputer Center in Guangzhou, SYSU

Jingkun Chen obtained his PhD in Condensed Matter Physics from Sun Yat-sen University, where his main research focuses on computational material science. Since joining the NSCC-GZ in 2014, he has been working as an HPC application engineer and is the group leader of the NSCC-GZ's Astronomy, Geoscience and Environment Engineering Application Platform.

Talk:

Design a Suitable Program Frame to Minimize I/O Cost for Large Scale Simulations on Tianhe2

Abstract: I/O is an important part of large scale evolution simulation applications such as weather forecast, particle diffusion and CFD. With the expansion of scale, the time of I/O increases as the number of cores raises. One solution to this problem is to overlap the I/O and iterative computation as much as possible, hence special consideration is needed during the design or optimization of the program frame. Actually, the problem of I/O contains three components, communication, read-write speed and memory size, therefore, a solution of I/O also needs to obtain the balances between physical limits. In this talk, I will mainly discuss my work of optimizing the program frame for better I/O performance on Tianhe2. After a number of tests, I found that a suitable frame can make the cost of I/O ignorable with even up to 10000 cores, and can improve the performance by 7% to 33%.

Jiahui Li

National Supercomputer Center in Guangzhou, SYSU

Dr. Jiahui Li is currently working as an HPC Application Engineer in the National Supercomputer Center in Guangzhou. He obtained his PhD from ZheJiang University, where his main research are Plasma Physics and Parallel Computing. Since joining the NSCC-GZ in 2014, he has been focusing on HPC applications of computational biology and precision medicine on Tianhe-2.

Talk:
Applications of Computational Biology and Precision Medicine on Tianhe-2

Abstract: A number of applications on Tianhe-2 are focused on fields of bioinformatics and drug discovery. NSCC-GZ has devoted a lot of work into these areas, some of which will be introduced in this talk. One is deploying the well-known web-based platform GALAXY for biomedical research on Tianhe-2. The front end of the GALAXY is deployed on the Cloud platform of Tianhe-2, while the back end is running on its HPC platform. We have achieved some satisfying outcomes but efforts are still needed. Another is virtual screening on Tianhe-2. To implement massive molecular docking, we are trying to combine Autodock vina and MongoDB, while the Intel Xeon Phi coprocessor is also used to accelerate the screening.