30th International Symposium on Computer Architecture and High Performance Computing

Ecole Normale Supérieure de Lyon, France

24-27 September 2018

This program has been optimized to reduce its environmental impact. Don't throw it on the nice streets of Lyon.

Index

Wifi Acces during the Conference :

Eduroam will be available in all rooms. For those who don't have an **Eduroam** access, a free WiFi access will be available :

WiFi SSID : invites
Login : z471015
Password : 7twxKEBJ

Conference proceedings :

http://conferences.computer.org/sbac-pad/2018,
username: sbac-pad18
password: conf18//

This program has been optimized to reduce its environmental impact. Don't throw it on the nice streets of Lyon.

Message from the SBAC-PAD 2018 General Chairs

On behalf of the organizing committee, we welcome you to the 30th International Symposium on Computer Architecture and High-Performance Computing (SBAC-PAD 2018), at the Ecole Normale Supérieure of Lyon, France. Started in 1987, SBAC-PAD has continuously presented an overview of new developments, applications, and trends in parallel and distributed computing technologies. SBAC-PAD has been open for faculty members, researchers, specialists and graduate students around the world. Striving to continue this tradition, we have received and evaluated papers across a large spectrum of the computer architecture, scheduling and high performance computing areas.

Several important players in the field have accepted to share their expertise and present invited talks at the conference: Jose E. Moreira (Thomas J. Watson Research Center, IBM Research), Yves Robert (ENS Lyon, France and University of Tennessee, Knoxville, USA), Haohuan Fu (Tsinghua University, China) and Manish Parashar (Rutgers University, USA). The first edition of an exciting workshop on High Performance Machine Learning (HPML) offers the research community with a forefront opportunity to present and discuss ideas on the intersection of machine learning and HPC. This workshop is held along with 9th edition of the Workshop on Applications for Multi-Core Architectures (WAMCA). Three tutorials are offered covering topics on secure application execution on the Cloud, application containers for HPC, and execution of distributed-computing experiments on the Grid'5000 testbed.

SBAC-PAD 2018 is a team effort comprising researchers and practitioners aimed at tackling difficult research challenges and enable their discussion on the symposium main track and collocated events. We thank everyone who contributed to making SBAC-PAD 2018 a reality. We would like to express a special gratitude to the chairs responsible for program committee, tracks, workshops, tutorials, posters, publicity, publication, website as well as the local people involved in the organization of the symposium. We also express our gratitude to the authors of all submitted papers, the program committee, and all external reviewers, who worked really hard on reviewing manuscripts and providing feedback for improvements.

Finally we would like to thank the support of Inria, Laboratoire de l'Informatique du Parallelisme (LIP), Ecole Normale Supérieure de Lyon, IEEE Computer Society and Brazilian Computer Society (SBC).

Laurent Lefevre (Inria, LIP, ENS Lyon)
Marcos Dias de Assuncao (Inria, LIP, ENS Lyon)
Alfredo Goldman (University of Sao Paulo)
General Chairs of SBAC-PAD 2018

Message from the SBAC-PAD 2018 Program Chairs

On behalf of the program committee, we welcome you to the 30th Symposium on Computer Architecture and High-Performance Computing (SBAC-PAD 2018). We are confident that the exceptional program that we have put together will attract many researchers and provide a great opportunity for exchanging experiences and research directions.

The SBAC-PAD 2018 Program Committee has endeavored to put together a stimulating technical program consisting of full and short papers. The Program Committee had the notably difficult task of selecting 30 papers out the 108 submissions (27% acceptance rate) received from 313 individuals from Brazil (28%), United States (16%), France (12%), Germany (9%) and 22 other different countries. The accepted full papers are from United States (8 papers), Brazil (7), France (6), Germany (5), and one paper from Austria, Greece, Spain, and Turkey. We also accepted 10 short papers.

The peer review was organized in five tracks, including the new Performance Evaluation track. A total of 343 reviews were submitted by 79 program committee members and 51 external reviewers. Each paper received at least 3 reviews; some submissions received up to 5 reviews. Cases with diverging recommendations were given additional attention by the Track Chairs and the Program Committee co-chairs.

We should make note here of the outstanding work performed by the Program Committee and would like to recognize the extraordinary leadership of the five Track Chairs: Kalyana Chadalavada (Intel, USA) of Computer Architecture, Wagner Meira Jr. (Federal University of Minas Gerais, Brazil) of Networking and Distributed Systems, Enrique Quintana-Ortí (Universidad Jaime I, Spain) of Parallel Applications and Algorithms, Arnaud Legrand (CNRS/Inria/Université Grenoble Alpes, France) of Performance Evaluation, and Adrien Lebre (IMT Atlantique/Inria/LS2N, France) of System Software. Their valuable work has made ours much easier.

We must emphasize that a number of people have worked hard to make SBAC-PAD 2018 a great success. We would like to thank in particular all the members of the steering and program committees, the special sessions chairs, the publicity and publication chairs, the local arrangements chair, and all the authors and reviewers.

Thank you.

Rosa Badia (Barcelona Supercomputing Center, Spain)
Manish Parashar (Rutgers University, USA)
Lucas Mello Schnorr (Federal University of Rio Grande do Sul, Brazil)
SBAC-PAD 2018 Program Chairs

Committees of SBAC-PAD 2018

General Chairs

Laurent Lefèvre (Inria, ENS Lyon, University of Lyon, France)
Alfredo Goldman (Sao Paulo University, Brazi)
Marcos Dias de Assuncao (Inria, ENS Lyon, University of Lyon, France)

Program Chairs

Rosa Badia (Barcelona Supercomputing Center, Spain)
Manish Parashar (Rutgers University, USA)
Lucas Schnorr (Federal University of Rio Grande do Sul, Brazil)

Workshop Chairs

Jesus Carretero (University Carlos III of Madrid, Spain)
Lucia Drummond (Federal Fluminense University, Brazil)

Tutorial Chairs

Frédéric Desprez (Inria Grenoble, France)
Dilma da Silva (University of Texas, USA)

Publicity Chairs

Alba Cristina de Melo (University of Brasilia, Brazil)
Pierre Sens (LIP6, Inria Paris Rocquencourt, France)
Vaidy Sunderam (Emory University, USA)

Publication Chair

Hermes Senger (Federal University of São Carlos, Brazil)

Website Chair

Jean-Christophe Mignot (CNRS, ENS Lyon, University of Lyon, France)

Program Committee

Computer Architecture

Chair: Kalyana Chadalavada (Intel, USA)

Celso Mendes (University of Illinois at Urbana-Champaign, USA)
Cristiana Bentes (UERJ, Brazil)
Edson Borin (University of Campinas, Brazil)
Farah Fargo (University of Arizona, USA)
Felipe Maia Galvao Franca (COPPE-UFRJ, Brazil)
Jose Moreira (IBM, USA)
Kalyana Chadalavada (Intel, USA)
Naveen Cherukuri (NVIDIA, USA)
Rathish Jayabharathi (Intel, USA)
Rodolfo Azevedo (University of Campinas, Brazil)
Adwait Jog (College of William and Mary, USA)
Arkaprava Basu (Indian Institute of Science, India)
Biswabandan Panda (IIT Kanpur, India)
David Kaeli (Northeastern University, USA)

Networking and Distributed Systems

Chair: Wagner Meira, Jr, (Federal University of Minas Gerais, Brazil)

Andrey Britto (UFCG, Brazil)
Bruno Schulze (LNCC, Brazil)
Carlos Varela (Rensselaer Polytechnic Institute, USA)
Domenico Talia (University of Calabria, Italy)
Dorgival Guedes Neto (UFMG, Brazil)
Elias P. Duarte Jr. (UFPR, Brazil)
Harold Castro (Universidad de los Andes, Colombia)
Javier Garcia-Blas (University Carlos III of Madrid, Spain)
Juan Durillo (University of Innsbruck, Austria)
Nicolae Bogdan (Huawei, China)
Omer Rana (Cardiff University, UK)
Priscila Solis (UnB, Brazil)
Rizos Sakellariou (University of Manchester, UK)
Rodrigo Fonseca (Brown, USA)
Ronaldo Ferreira (UFMS, Brazil)
Sandra Gesing (Notre Dame University, USA)
Zhiyi Huang (University of Otago, New Zealand)
Ignacio Blanquer (UPV, Spain)
Renato Ferreira (UFMG, Brazil)

Parallel Applications and Algorithms

Chair: Enrique Quintana-Ortí, (Universidad Jaime I, Spain)

Emmanuel Agullo (Inria, France)
José Ignacio Aliaga (Computer Science and Engineering Dpt, University Jaime I, Spain)
Hartwig Anzt (Karlsruhe Institute of Technology, Germany)
Michael Bader (Technical University of Munich, Germany)
Paolo Bientinesi (RWTH Aachen University, Germany)
Kei Davis (Los Alamos National Laboratory, USA)
Alba Melo (University of Brasilia, Brazil)

Pablo Ezzatti Udelar (Uruguay)
Mathieu Faverge (Bordeaux INP, France)
Domingo Gimenez (University of Murcia, Spain)
Daniel S. Katz (University of Illinois Urbana-Champaign, USA)
Hatem Ltaief (KAUST, Saudi Arabia)
Piotr Luszczek (University of Tennessee Knoxville, USA)
Viktor Prasanna (University of Southern California, USA)
Olaf Schenk (Universita della Svizzera Italiana, Switzerland)
Leonel Sousa (Lisbon University, Portugal)
Guillermo Taboada (University of A Coruña, Spain)
Pedro Valero-Lara (The University of Manchester, UK)
Jaroslaw Zola (State University of New York, USA)
Jairo Panetta (INPE, Brazil)

Performance Evaluation

Chair: Arnaud Legrand (CNRS/Inria/Université Grenoble Alpes, France)

Laura Carrington (San Diego Supercomputer Center/University of California, USA)
Cesar De Rose (PUCRS, Brazil)
Alfredo Goldman (University of São Paulo, Brazil)
Andreas Knupfer (TU Dresden, Germany)
Swann Perarnau (Argonne National Laboratory, USA)
Martin Schulz (Technical University of Munich, Germany)
Frederic Suter (CC IN2P3 / CNRS, France)
Sascha Hunold (TUW, Austria)
Rafael Ferreira da Silva (USC USA)
Judit Gimenez (BSC, Spain)
Nelson Amaral (U. Alberta, Canada)

System Software

Chair: Adrien Lebre (IMT Atlantique/Inria/LS2N, France)

Rafael Asenjo (Dpt. Computer Architecture, Univ. Malaga. Spain)
Paolo Bellavista (University of Bologna, Italy)
Siegfried Benkner (University of Vienna, Austria)
Carlo Bertolli (IBM TJ Watson Research Center, United States)
Francisco Brasileiro (UFCG, Brazil)
Helene Coullon (Inria, France)
Michael Gerndt (Technical University of Munich, Germany)
Antonio J. Peña (Barcelona Supercomputing Center -BSC, Spain)
Manuel Prieto-Matias (Complutense University of Madrid, Spain)
Mauricio Pillon (UDESC, Brazil)
Radu Prodan (University of Klagenfurt, Austria)
Mario Südholt (IMT Atlantique, France)
Cedric Tedeschi (Inria, France)
Ayal Zaks (Intel, Israel)

SBAC-PAD 2018 Conference Program

Monday, September 24, 2018	
08:00-17:00	Ecole Normale Supérieure of Lyon - Site Buisson - Main Entrance: **Registration**
	Workshops & Tutorials Day
Location	D8 001 Room
09:00-10:00	Joint Keynote 1: **"Serendipity: How Supercomputing Technology is Enabling a Revolution in Artificial Intelligence ", José E. Moreira, IBM Research** *Keynote chair: Laurent Lefevre* **Abstract :** With the availability of both large compute power and large data sets, we have witnessed a revolution in machine learning technology, which has become a mainstream tool for both business and scientific applications. This revolution is likely to accelerate, as even more compute power is brought to bear, and deliver many of the promises of artificial intelligence. In this talk we will investigate how far the impacts of machine learning can go. We will cover the new Summit supercomputer, which brings unprecedented compute capabilities to both traditional high performance computing and artificial intelligence problems, analyzing the similarities as well as the differences in those two fields. We will also speculate about the future of machine learning and, in particular, its possible limitations. We will conclude with a discussion of one of the most important scientific questions of our time: Is

	consciousness computable?  José E. Moreira is a Distinguished Research Staff Member in the Scalable Systems Department at the Thomas J. Watson Research Center. He received a B.S. degree in physics and B.S. and M.S. degrees in electrical engineering from the University of Sao Paulo, Brazil, in 1987, 1988 and 1990, respectively. He also received a Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 1995. Since joining IBM at the Thomas J. Watson Research Center, he has worked on a variety of high-performance computing projects. He was system software architect for the Blue Gene/L supercomputer and chief architect of the Commercial Scale Out project. He currently leads the IBM Research work on the architecture of Power processor. He is an author or coauthor of over 100 technical papers and 10 patents. Dr. Moreira is a member of the IEEE (Institute of Electrical and Electronics Engineers) and a Distinguished Scientist of the ACM (Association for Computing Machinery).
10:00-10:30	**Coffee & Tea & Juice Break**

Location	Room Buisson 1 D8 001	Room Buisson 2 D8 003	Room Buisson 3 D8 006	Room Buisson 4 D8 007
10:30-13:00	**HPLM Workshop**	WAMCA Workshop	Tutorial 1: **Secure execution in the cloud using Intel SGX**	Tutorial 2: **Singularity: an HPC application container**
13:00-14:00	**Lunch at the Buisson Cafeteria**			
14:00-16:00	**HPLM Workshop**	WAMCA Workshop	Tutorial 3: **Distributed computing experiments using Grid'5000 / SILECS testbed**	
16:00-16:30	**Coffee & Tea & Juice Break**			
16:30-18:00	**HPLM Workshop**	WAMCA Workshop	Tutorial 3: **Distributed computing experiments using Grid'5000 / SILECS testbed**	

Tuesday, September 25, 2018

08:00-17:00	Ecole Normale Supérieure of Lyon - Site Buisson - Main Entrance: **Registration**
09:00-09:15	Room Buisson 1 - D8 001 : **SBAC-PAD 2018 - Day 1 - Opening Session** *Laurent Lefèvre, Marcos Dias de Assuncao, Alfredo Goldman, Lucas Schnorr*
09:15-10:15	**Keynote 2: "Scheduling Matters", Yves Robert, ENS Lyon, France and Univ. Tenn. Knoxville, USA** *Keynote chair: Alfredo Goldman* **Abstract :** This talk will review a few scheduling algorithms to solve simple computational problems on large-scale platforms. Faults, energy/power shortage, I/O contention, the constraints are numerous and challenging. The talk will provide a few answers and discuss open research directions. 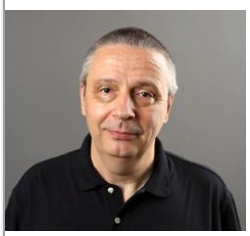 Yves Robert received the PhD degree from Institut National Polytechnique de Grenoble. He is currently a full professor in the Computer Science Laboratory LIP at ENS Lyon. He is the author of 7 books, 150 papers published in international journals, and 240 papers published in international conferences. He is the editor of 11 book proceedings and 13 journal special issues. He has advised 30 PhD students. His main research interests are scheduling techniques and resilient algorithms for large-scale platforms. Yves Robert served on many editorial boards, including IEEE TPDS, JPDC and ACM TOPC. He is a Fellow of the IEEE. He was elected a Senior Member of Institut Universitaire de France in 2007 and renewed in 2012. He was awarded the 2014 IEEE TCSC Award for Excellence in Scalable Computing, and the 2016 IEEE TCPP Outstanding Service Award. He holds a Visiting Scientist position at the University of Tennessee Knoxville since 2011.
10:15-10:45	**Coffee & Tea & Juice Break**
	Conference
Location	D8 001
10:45-12:30	**Session 1: Computer Architecture and Compilers** *Session chair: Philippe Olivier and Alexandre Navaux (UFRGS)* **MLNoC: A Machine Learning Based Approach to Network-on-Chip Design:** Nishant Rao, Akshay Ramachandran and Amish Shah

	ADeLe: Rapid Architectural Simulation for Approximate Hardware: Isaias Bittencourt Felzmann, Matheus M. Susin, Liana Duenha, Rodolfo Azevedo and Lucas Wanner **From Java to FPGA: an Experience with the Intel HARP System:** Pedro Caldeira, Jerônimo Penha, Lucas Bragança, José Nacif, Ricardo Ferreira, Renato Ferreira and Fernando Quintao **Online Detection of Spectre Attacks Using Microarchitectural Traces from Performance Counters:** Congmiao Li and Jean-Luc Gaudiot (short paper) **DOACROSS Parallelization based on Component Annotation and Loop-carried Probability:** Luis Felipe Mattos, César Divino, Juan Salamanca, Joao Paulo Carvalho, Marcio Machado Pereira and Guido Araujo (short paper)
12:30-14:00	**Lunch at the Restarant ENS Descarte**
14:00-16:00	**Session 2: Scheduling** *Session chair: Arnaud Legrand (CNRS)* **Scheduling independent stochastic tasks under deadline and budget constraints:** Louis-Claude Canon, Aurélie Kong Win Chang, Yves Robert and Frédéric Vivien **Adaptive scheduling of collocated applications using a task-based runtime system:** Jiri Dokulil and Siegfried Benkner **A Batch Task Migration Approach for Decentralized Global Rescheduling:** Vinicius Freitas, Alexandre Santana, M´arcio Castro and Laércio L. Pilla **Exploring Power Budget Scheduling Opportunities and Trade-offs for AMR-based Applications:** Yubo Qin, Ivan Rodero, Pradeep Subedi, Manish Parashar and Sandro Rigo **EASE: Energy Efficiency and ProportionalityAware Virtual Machine Scheduling:** Congfeng Jiang, Yumei Wang, Dongyang Ou, Yeliang Qiu, Youhuizi Li, Jian Wan, Bing Luo, Weisong Shi and Christophe Cerin (short paper)
16:00-16:30	**Coffee & Tea & Juice Break**
16:30-18:00	**Session 3: Energy in the Cloud, Network** *Session chair: Ivan Rodero (Rutgers University)* **Energy-Efficient IaaS-PaaS Co-design for Flexible Cloud Deployment of Scientific Applications:** David Guyon, Anne-Cécile Orgerie and Christine Morin **Frequency Selection Approach for Energy Aware Cloud Database:** Chaopeng Guo and Jean-Marc Pierson **Network-aware energy-efficient virtual machine management in distributed Cloud infrastructures with on-site photovoltaic production:** Benjamin Camus, Fanny Dufossé, Anne Blavette, Martin Quinson and Anne-Cécile Orgerie **A Novel Broker-Based Hierarchical Authentication Scheme in Proxy Mobile IPv6 Networks,** Jang Su Hwan and Jeong Jongpil (short paper)
18:30-20:30	**Welcome Reception and Cocktail at the Descarte Restaurant and Buisson Garden**

Wednesday, September 26, 2018

08:00-17:00	Ecole Normale Supérieure of Lyon - Site Buisson - Main Entrance: **Registration**
09:00-09:15	Room Buisson 1 - D8 001 : **SBAC-PAD 2018 - Day 2 -Opening Session** *Laurent Lefèvre, Marcos Dias de Assuncao, Alfredo Goldman, Lucas Schnorr*
09:15-10:15	**Keynote 3: "Extreme-Scale Earthquake Simulation on Sunway TaihuLight", Haohuan Fu, Tsinghua University, China** *Keynote chair: Lucas Schnorr* **Abstract** : This talk would first introduce and discuss the design philosophy about the Sunway TaihuLight system, and then describe our recent efforts on performing earthquake simulations on such a large-scale system. Our work in 2017 accomplished a complete redesign of AWP-ODC for Sunway architectures, achieves over 15% of the system's peak, better than the 11.8% achieved by a similar software running on Titan, whose byte to flop ratio is 5 times better than TaihuLight. The extreme cases demonstrate a sustained performance of over 18.9 Pflops, enabling the simulation of Tangshan earthquake as an 18-Hz scenario with an 8-meter resolution. Our recent work further improves the simulation framework with capabilities to describe complex surface topography, and to drive building damage prediction and landslide simulation, which are demonstrated with a case study of the Wenchuan earthquake with accurate surface topography and improved coda wave effects.  Haohuan Fu is the deputy director of the National Supercomputing Center in Wuxi, leading the research and development division. He is also an associate professor in the Ministry of Education Key Laboratory for Earth System Modeling, and Department of Earth System Science in Tsinghua University, where he leads the research group of High Performance Geo-Computing (HPGC). Fu has a PhD in computing from Imperial College London. Since

	joining Tsinghua in 2011, Dr. Fu has been working towards the goal of providing both the most efficient simulation platforms and the most intelligent data management and analysis platforms for geoscience applications. His research has, for example, led to efficient designs of atmospheric dynamic solvers for both Tianhe-1A, Tianhe-2, Sunway TaihuLight supercomputers, and the reconfigurable computing platforms. The work based on the Sunway TaihuLight supercomputer manages to scale a fully-implicit solver to over 10 million cores, which won the Gordon Bell Prize of SC16.
10:15-10:45	**Coffee & Tea & Juice Break**
	Conference
Location	D8 001
10:45-12:30	**Session 4: Applications** *Session chair: Guido Araujo (UNICAMP)* **Designing a Parallel Memory-Aware Lattice Boltzmann Algorithm on Manycore Systems:** Yuankun Fu, Feng Li, Fengguang Song and Luoding Zhu **A new efficient parallel algorithm for minimum spanning tree:** Jucele Vasconcellos, Edson Caceres, Henrique Mongelli and Siang Song **Exploring Self-Adaptivity towards Performance and Energy for Time-stepping Methods:** Natalia Kalinnik, Robert Kiesel, Thomas Rauber, Marcel Richter and Gudula Rünger **Predicting the Reliability Behavior of HPC Applications:** Daniel Oliveira, Francis Birck Moreira, Paolo Rech and Philippe Navaux
12:30-13:30	**Lunch at the Restarant ENS Descarte**
13:30-15:30	**Session 5: GPU based computing** *Session chair: Eduardo Rodrigues (IBM Research)* **Variable-size batched condition number calculation on GPUs:** Hartwig Anzt, Jack Dongarra, Goran Flegar and Thomas Grutzmacher **Towards a Single-host Many-GPU System:** Ming-Hung Chen, Ihsin Chung, Bulent Abali and Paul Crumley **Exploiting Limited Access Distance for Kernel Fusion Across the Stages of Explicit One-Step Methods on GPUs:** Tim Werner and Matthias Korch **Balancing load of GPU subsystems to accelerate image reconstruction in parallel beam tomography:** Suren Chilingaryan, Evelina Ametova, Andreas Kopmann and Alessandro Mirone **Performance Prediction of GPU-based Deep Learning Applications:** Eugenio Gianniti, Li Zhang and Danilo Ardagna (short paper)
15:30-16:30	**Poster Session during the Coffee & Tea & Juice Break** **Hybrid MPI+OpenMP Implementation of eXtended Discrete Element Method** Abdoul Wahid Mainassara Chekaraou, Alban Rousset, Xavier Besseron and Bernhard Peters **Privacy Preserving Data Outsourcing via Secure Order-Preserving Secret Splitting Scheme:** Somayeh Sobati-M. **Deep Reinforcement Learning Based Approach for Energy-Aware Job Scheduling in Grid Computing:** Lucas Casagrande, Kleiton Pereira, Renato Tanaka, Charles Miers, Guilherme Koslovski and Mauricio Pillon **Analysis and Characterization of Control Network Traffic in OpenStack Based IaaS Clouds:** Tiago Reinert, Adnei Donatti, Guilherme Koslovski, Maurício Pillon and Charles Miers **Performance and Energy Consumption of Parallel Programming Interfaces in Multicore Architectures: A Case Study:** Adriano Garcia, Claudio Schepke, Alessandro Girardi, Sherlon Almeida Da Silva **Automatic Parallelization for Shared Memory Scientific Multiprocessing: A State-of-the-art:** Re'Em Harel, Idan Mosseri, Harel Levin, Matan Rusanovsky, Gal Oren
16:30-18:00	**Session 6: Programming Paradigms and Memory** *Session chair: Nelson J. Amaral (UAlberta)* **Polyhedral Dataflow Programming: a Case Study:** Romain Fontaine, Laure Gonnord and Lionel Morel **Enabling Efficient Job Dispatching in Accelerator-extended Heterogeneous Systems with Unified Address Space:** Georgios Kornaros **Phase-Based Data Placement Scheme for Heterogeneous Memory Systems:** Mohammad Laghari, Najeeb Ahmad and Didem Unat **Exploiting Compute Caches for Memory Bound Vector Operations:** Joao Vieira, Paolo Ienne, Nuno Roma, Gabriel Falcao and Pedro Tomas (short paper)
	Free Evening

Thursday, September 27, 2018

08:00-17:00	Ecole Normale Supérieure of Lyon - Site Buisson - Main Entrance: **Registration**
09:00-09:15	Room Buisson 1 - D8 001: **SBAC-PAD 2018 - Day 3 - Opening Session** *Laurent Lefèvre, Marcos Dias de Assuncao, Alfredo Goldman, Lucas Schnorr*
09:15-10:15	**Keynote 4: "Big Data at Extreme-Scales: Addressing Computational Challenges in the 21st Century", Manish Parashar, Rutgers University, USA** *Keynote chair: Laurent Lefevre* **Abstract** :Data-related challenges are quickly dominating computational and data-enabled sciences and are limiting the potential impact of scientific application workflows enabled by current and emerging extreme scale, high-performance, distributed computing environments. These data-intensive application workflows involve dynamic coordination, interactions and data coupling between multiple application processes that run at scale on different resources, and with services for monitoring, analysis and visualization and archiving, and present challenges due to increasing data volumes and complex data-coupling patterns, system energy constraints, increasing failure rates, etc. In this talk I will explore some of these challenges and investigate how solutions based on data sharing abstractions, managed data pipelines, data-staging service, and in-situ / in-transit data placement and processing can be used to help address them. This research is part of the DataSpaces project at the Rutgers Discovery Informatics Institute.  Manish Parashar is Distinguished Professor of Computer Science at Rutgers University. He is also the founding Director of the Rutgers Discovery Informatics Institute (RDI2). His research interests are in the broad areas of Parallel and Distributed Computing and Computational and Data-Enabled Science and Engineering. Manish is the founding chair of the IEEE Technical Consortium on High Performance Computing (TCHPC), Editor-in-Chief of the IEEE Transactions on Parallel and Distributed Systems, and serves on the editorial boards and organizing committees of a large number of journals and international conferences and workshops. He has over 350 publications, has deployed several software systems that are widely used, and has received a number of awards for his research and leadership. Manish is Fellow of AAAS, Fellow of IEEE/IEEE Computer Society and ACM Distinguished Scientist.
10:15-10:45	**Coffee & Tea & Juice Break**
	Conference
Location	D8 001
10:45-12:30	**Session 7: Data Analytics, Locality and I/O** *Session chair: Bruno Raffin (INRIA)* **Exploring the Potential of Next Generation Software-Defined In-Memory Frameworks:** Shouwei Chen and Ivan Rodero **Towards Green Scientific Data Compression Through High-Level I/O Interfaces:** Yevhen Alforov, Anastasiia Novikova, Michael Kuhn, Julian Kunkel and Thomas Ludwig **Improving Data Locality in P2P-based Fog Computing Platforms:** Luiz Angelo Steffenel **echofs: A Scheduler-guided Temporary Filesystem to leverage Node-local NVMs:** Alberto Miranda, Ramon Nou and Toni Cortes (short paper) **A Jaccard Weights Kernel Leveraging Independent Thread Scheduling on GPUs:** Hartwig Anzt and Jack Dongarra (short paper)
12:30-14:00	**Lunch at the Restarant ENS Descarte**
14:00-16:00	**Session 8: Performance Prediction and Evaluation** Session chair: Hartwig Anzt (University of Tennessee) **Multicore Performance Engineering of Sparse Triangular Solves Using a Modified Roofline Model:** Markus Wittmann, Georg Hager, Radim Janaľık, Martin Lanser, Axel Klawonn, Oliver Rheinbach, Olaf Schenk and Gerhard Wellein **Predicting the Performance Impact of Increasing Memory Bandwidth for Scientific Workflows:** Nelson Mimura Gonzalez, Jose Brunheroto, Fausto Artico, Yoonho Park, Tereza Carvalho, Charles Christian Miers, Mauricio Aronne Pillon and Guilherme Piegas Koslovski **Mainstream vs. Emerging HPC: Metrics, Trade-offs and Lessons Learned:** Milan Radulovic, Kazi Asifuzzaman, Darko Zivanovic, Nikola Rajovic, Guillaume Colin de Verdi`ere, Dirk Pleiter, Manolis Marazakis, Nikolaos Kallimanis, Paul Carpenter, Petar Radojkovic and Eduard Ayguadé **Assessing Time Predictability Features of ARM big.LITTLE Multicores:** Gabriel Fernandez, Francisco J Cazorla, Jaume Abella and Sylvain Girbal (short paper) **Adaptive Partitioning for Iterated Sequences of Irregular OpenCL Kernels:** Pierre Huchant, Marie-Christine Counilh

	and Denis Barthou (short paper)
16:00-16:30	**Coffee & Tea & Juice Break**
16:30-18:00	**Session 9: IoT, Fog, Edge, and Cloud Computing** Session chair: Anne-Cécile Orgerie (CNRS) **Partitioning convolutional neural networks for inference on constrained Internet-of-Things devices:** Fabiola Martins Campos de Oliveira and Edson Borin **Runtime Management of Data Quality for Scientific Observatories Using Edge and In-Transit Resources:** Ali Reza Zamani, Daniel Balouek-Thomert, J. J. Villalobos, Ivan Rodero and Manish Parashar **A Fault-Tolerant Agent-based Architecture for Transient Servers in Fog Computing:** Jose Pergentino de Araujo Neto, Célia G. Ralha and Donald M. Pianto
19:00-22:30	 **Banquet Reception on the Hermes boat** **3 Hours Cruise along the Rhône and Sâone Rivers** **Best Paper Award Ceremony** **Annoucement of SBAC-PAD 2019** (see page 15 for details)

Workshops

1st High Performance Machine Learning (HPML) Workshop (https://hpml2018.github.io)

We are delighted to present in these proceedings the papers of the first edition of the **H**igh **P**erformance **M**achine **L**earning Workshop (HPML 2018). HPML 2018 is intended to bring together the Machine Learning (ML), Artificial Intelligence (AI) and High Performance Computing (HPC) communities. In recent years, much progress has been made in ML and AI in general. This progress required heavy use of high performance computers and accelerators. Moreover, ML and AI have become a "killer application" for HPC and, consequently, driven much research in this area as well. These facts point to an important cross-fertilization that this workshop intends to nourish.

HPML 2018 is this year co-located with the 30th International Symposium on Computer Architecture and High Performance Computing - SBAC-PAD in Lyon, France. We received 31 submissions from which we selected 14 papers in a peer-reviewed process. Most papers received three or more reviews. The paper topics show a broad interest in the interplay of HPC, ML and AI. Moreover, the authors of the selected papers are affiliated to a good mix of universities (12), research centers (4) and companies (3), from a total of 9 countries. In addition to the paper presentations, HPML 2018 program included the keynote speech by Jose Moreira (IBM Research): "Serendipity: How Supercomputing Technology is Enabling a Revolution in Artificial Intelligence".

Moreover, a number of papers from HPML 2018 will be selected and their authors invited to submit an extension to the *Special Issue on Advances on Parallel and High Performance Computing for AI Applications*. This is a special issue of the high impact *Journal of Parallel and Distributed Computing* - Elsevier JPDC.

We want to thank the authors who responded to our call for papers, the program committee and the numerous reviewers for the effort to evaluate the submitted papers, ensuring a very high quality program. We want also to thank SBAC-PAD committees for the help organizing this workshop.

Finally, we hope these proceedings will serve as reference for the work of other researchers interested in Machine Learning, Artificial Intelligence a High Performance Computing.

Eduardo Rodrigues (IBM Research, Brazil)
Jairo Panetta (Instituto Tecnologico de Aeronautica, ITA, Brazil)
Bruno Raffin (INRIA, France)
Organizing Committee

Organizing Committee

Eduardo Rocha Rodrigues, IBM Research - edrodri (at) br (dot) ibm (dot) com
Jairo Panetta, Instituto Tecnologico de Aeronautica, ITA, Brazil
Bruno Raffin, INRIA, France

Program Committee

Abhishek Gupta, Schlumberger, USA
Albert N. Kahira, Barcelona Supercomputing Center, Spain
Aline Paes, Universidade Federal Fluminense, UFF, Brazil
Andrea Schwertner Charao, Universidade Federal de Santa Maria, UFSM, Brazil
Bruno Silva, IBM Research, Brazil
Celso Mendes, Instituto Nacional de Pesquisas Espaciais, INPE, Brazil
Daniel Salles Chevitarese, IBM Research, Brazil
Dingwen Tao, University of Alabama, USA
Edmilson Morais, IBM Research, Brazil
Eduardo Vasconcellos, Universidade Federal Fluminense, UFF, Brazil
Emmanuel Jeannot, INRIA, France
Fabio Cozman, Universidade de Sao Paulo, USP, Brazil
Francis Birck Moreira, Universidade Federal do Rio Grande do Sul, UFRGS, Brazil
François Tessier, Argonne National Laboratory, USA
Gina Maira Barbosa de Oliveira, Universidade Federal de Uberlandia, UFU, Brazil
Guillaume Aupy, INRIA, France
Haroldo Fraga de Campos Velho, Instituto Nacional de Pesquisas Espaciais, INPE, Brazil
Ian Masliah, Laboratoire d'Informatique de Paris, LIP, France
Jose Celaya, Schlumberger, USA
Leonardo Bautista Gomez, Barcelona Supercomputing Center, Spain
Lucas Mello Schnorr, Universidade Federal do Rio Grande do Sul, UFRGS, Brazil
Luiz Gustavo Almeida Martins, Universidade Federal de Uberlandia, UFU, Brazil
Marc Casas, Barcelona Supercomputing Center, Spain
Marceli Zanon-Boito, Laboratoire d'Informatique de Grenoble, LIG, France
Marco Netto, IBM Research, Brazil
Mauricio Araya, Shell Oil USA, USA
Oguz Kaya, INRIA, France
Olivier Beaumont, INRIA, France
Olivier Coulaud, INRIA, France
Pedro Mario Cruz, Nvidia, Brazil
Renato Cunha, IBM Research, Brazil
Suhas Suresha, Schlumberger, USA
Xin Liang, University of California Riverside, USA

Program

Monday, September 24	
09:00 10:00	**KEYNOTE - Serendipity: How Supercomputing Technology is Enabling a Revolution in Artificial Intelligence,** José E. Moreira, IBM Research
10:00 10:30	**Coffee & Tea & Juice Break**
10:30 10:55	**Large Scale Language Modeling: Converging on 40GB of Text in Four Hours,** Raul Puri, Robert Kirby, Nikolai Yakovenko, Bryan Catanzaro, Nvidia USA
10:55 11:20	**Accelerating deep neural network training for action recognition on a cluster of GPUs,** Guojing Cong, Giacomo Domeniconi, Joshua Shapiro, Fan Zhou, Barry Chen, IBM TJ Watson, Georgia Tech, Lawrence Livermore National Laboratory
11:20 11:45	**An argument in favor of strong scaling for deep neural networks with small datasets,** Renato Cunha, Eduardo Rodrigues, Matheus Palhares Viana and Dario Augusto Borges Oliveira, IBM Research
11:45 12:10	**Deep Learning on Large-scale Multicore Clusters,** Kazumasa Sakiyama, Shinpei Kato, Yutaka Ishikawa, Atsushi Hori and Abraham Monrroy, The University of Tokyo, Riken, Nagoya University
12:10 12:35	**On the Resilience of RTL NN Accelerators: Fault Characterization and Mitigation,** Behzad Salami, Osman Unsal and Adrian Cristal Kestelman, Barcelona Supercomputing Center
12:35 13:00	**t-SNE-CUDA: GPU-Accelerated t-SNE and its Applications to Modern Data,** David Chan, Roshan Rao, Forrest Huang and John Canny, University of California, Berkeley
13:00 14:00	**Lunch at the Buisson Cafeteria**
14:00 14:25	**HyperSpace: Distributed Bayesian Hyperparameter Optimization,** M. Todd Young, Jacob Hinkle, Arvind Ramanathan, Ramakrishnan Kannan, Oak Ridge National Lab
14:25 14:50	**A Machine Learning Approach for Parameter Screening in Earthquake Simulation,** Marisol Monterrubio-Velasco, Jose Carlos Carrasco-Jiménez, Octavio Castillo-Reyes, Fernando Cucchietti, Josep De la Puente, Barcelona Supercomputing Center

14:50 15:15	**A Case Study on Optimizing Accurate Half Precision Average,** Kenny Peou, Joel Falcou and Alan Kelly, Numscale, Université Paris-Saclay
15:15 15:40	**Optimization of a sparse grid-based data mining kernel for architectures using AVX-512,** Paul Cristian Sârbu and Hans-Joachim Bungartz, Technical University of Munich
15:40 16:05	**Energy Efficient Parallel K-Means Clustering for an Intel® hybrid Multi-Chip Package,** Matheus A. Souza, Lucas A. Maciel, Pedro H. Penna and Henrique C. Freitas, PUC Minas
16:05 16:30	**Coffee & Tea & Juice Break**
16:30 16:55	**Performance Comparison of a Parallel Recommender Algorithm across three Hadoop-based Frameworks,** Christina Diedhiou, Bryan Carpenter, Aamir Shafi, Soumabha Sarkar, Ramazan Esmeli, Ryan Gadsdon, University of Portsmouth, Imam Abdulrahman Bin Faisal University
16:55 17:20	**Effect Of Network Topology On The Performance Of ADMM-based SVMs,** Shirin Tavara and Alexander Schliep, University of Boras, University of Gothenburg
17:20 17:45	**High Performance Ensembles of Online Sequential Extreme Learning Machine for Regression and Time Series Forecasting,** Luís F. L. Grim and André L. S. Gradvohl, University of Campinas
17:45 18:00	**Closing remarks**

9th Workshop on Applications for Multi-Core Architectures (WAMCA)

http://www.cri.ensmp.fr/conf/wamca2018/

We are pleased to meet our community for the ninetieth edition of the **W**orkshop on **A**pplications for **M**ulti-**C**ore **A**rchitectures (WAMCA), held together with the International Symposium on Computer Architectures and High Performance Computing at Lyon, France. Since the advent of the multicore concept, significant advances have been made on many of the related aspects. Multi-core has become a natural element in the HPC ecosystem, and its pervasiveness within the landscape of common computing devices has made it a standard. The available on-chip computing power is increasing through more CPU-cores and more efficient memory systems. However, designing programs that keep scaling well on new generation multicore processors needs skilful investigations. WAMCA aims at providing an opportunity to address all related technical challenges and more, including reports on specific case studies and prospective studies.

This year we received **12 submissions**, from which we were able to select **9 as regular papers** for presentation in the workshop and **3 as posters** to be included in the main track (SBAC) posters session. The topics ranged from applications such as *bioinformatics, particles physics, linear algebra, stencil computations, image processing, graph partitioning*. Almost all papers received at least three reviews each and those with good average scores were selected. We would like to thank the members of the program committee and ad-hoc reviewers for their outstanding work in a short timeframe. We also need to address our special thanks the general chairs of SBAC-PAD for their help to increase the number of submissions and their cooperation in a convenient hosting of WAMCA.

In this edition, we were able to secure the participation in one special issue in a journal, the *International Journal of High Performance Computing and Networking (IJHPC)*, and other alternatives are under consideration (thanks to the willingness of Dr. Gabriele Mencagli from the University of Pisa for the alternative suggestions).

Finally, we would like to thank the entire community for keeping our workshop as a relevant forum for disseminations and discussions related to HPC topics. Enjoy WAMCA 2018!

Claude Tadonki (MINES ParisTech, France)
Cristiana Bentes (State University of Rio de Janeiro, Brazil)
Guido Araujo (State University of Campinas, Brazil)
Lucia Drummond (Fluminense Federal University, Brazil)
Mauricio Pilla (Federal University of Pelotas, Brazil)
Phillippe Navaux (Federal University of Rio Grande do Sul, Brazil)
Ricardo Farias (Federal University of Rio de Janeiro, Brazil)

General Chair

Claude Tadonki (MINES ParisTech - PSL)

Workshop Chairs

Cristiana Bentes (State University of Rio de Janeiro)
Gabriele Mencagli (University of Pisa - Italy)
Lucia Drummond (Fluminense Federal University)
Mauricio Pilla (Federal University of Pelotas)

Program Chairs

Guido Araujo (State University of Campinas)
Phillippe Navaux (Federal University of Rio Grande do Sul)
Ricardo Farias (Federal University of Rio de Janeiro)

Program Committee

Alexandre Gonçalves (IFRJ - Brazil)
Andre Rauber Du Bois (UFPEL - Brazil)
Corinne Ancourt (Mines ParisTech - France)
Costin Iancu (Lawrence Berkeley National Lab - USA)

Denis Barthou (INRIA - France)
Diego Brandao (CEFET-RJ - Brazil)
Diego Dutra (UFRJ - Brazil)
Eyder Rios (UESPI - Brazil)
Guido Araujo (UNICAMP - Brazil)
Guilherme Cox (Rutgers University - US)
Henrique Freitas (PUC MG - Brazil)
Igor Machado (UERJ - Brazil)
Juliana Silva Zamith (UFRRJ - Brazil)
Khaled Ibrahim (Lawrence Berkeley National Lab - USA)
Konstantin Petrov (INRIA - France)
Leandro Marzulo (UERJ - Brazil)
Leonel Sousa (INESC-ID - Portugal)
Mauricio Pilla (UFPEL - Brazil)
Marcelo Zamith (UFRRJ - Brazil)
Marcin Paprzycki (IBS PAN and WSM - Poland)
Mitsuhisa Sato (RIKEN AICS / University of Tsukuba - Japan)
Rodolfo Azevedo (UNICAMP - Brazil)
Sandro Rigo (UNICAMP - Brazil)

Program

Session I : Architecture and Performance Analysis	
11:00 - 11:25	**Design Space Exploration of Energy Efficient NoC- and Cache-based Many-Core Architectures**, Matheus A. Souza, Henrique C. Freitas and Jean-François Mehaut, Pontifícia Universidade Católica (PUC) de Minas Gerais - Université Grenoble Alpes
11:30 - 11:55	**Highly Scalable Stencil-based Matrix-free Stochastic Estimator for the Diagonal of the Inverse**, Fabio Verbosio, Juraj Kardos, Mauro Bianco and Olaf Schenk, Università della Svizzera italiana - ETH/CSCS
12:00 - 12:25	**A Scalability and Sensitivity Study of Parallel Geometric Algorithms for Graph Partitioning**, Shad Kirmani, Hongyang Sun and Padma Raghavan, eBay Inc. - Vanderbilt University
13:00 - 14:00	Lunch at the Buisson Cafeteria
Session II OpenMP Parallelization	
14:00 - 14:25	**Automatic Ray-Tracer Cloud Offloading in OpenMP**, Matheus Diamantino, Hervé Yviquel and Guido Araujo, Universidade Estadual de Campinas (UNICAMP)
14:30 - 14:55	**Evaluation of an OpenMP Parallelization of Lucas-Kanade on a NUMA-Manycore**, Olfa Haggui, Claude Tadonki, Fatma Sayadi, and Bouraoui Ouni, Sousse National School of Engineering - Mines ParisTech - Université de Monastir
15:00 - 15:25	**Automated GPU Grid Geometry Selection for OpenMP Kernels**, Taylor Lloyd, Artem Chikin, Sanket Kedia, Dhruv Jain and José Nelson Amaral, University of Alberta - IIT Kharagpur
15:30 - 15:55	**OpenMP Implementation of eXtended Discrete Element Method**, Wahid Mainassara Chekaraou, Alban Rousset, Xavier Besseron, Sebastien Varrette and Bernhard Peters, LuXDEM Research Centre at University of Luxembourg
16:00 - 16:30	Coffee & Tea & Juice Break
Session III Hybrid Parallelization	
16:30 - 16:55	**Impacts of Three Soft-Fault Models on Hybrid Parallel Asynchronous Iterative Methods**, Erik Jensen, Evan Coleman and Masha Sosonkina, Old Dominion University
17:00 - 17:25	**Scaling and optimizing the GYSELA code on a cluster of many-core processors**, Guillaume Latu, Yuuichi Asahi, Julien Bigot, Tamas Bela-Feher and Virginie Grandgirard, CEA (Cadarache & Maison de la Simulation) - Rokkasho Fusion Institute - Max Planck Institute

Tutorials

Tutorial 1: Secure Execution in the Cloud Using Intel SGX

Speaker: Marcelo Pasin (University of Neuchatel, Switzerland)

Overview: Cloud computing is an appealing support for application deployment today. Its inherent resource sharing allows for economies of scale. Easy access and ubiquitous availability makes it very convenient. But the shared and remote nature of cloud resources render them tempting for attacks. One single cloud provider hosts data for multiple customers and an important financial gain may be achieved by

exploiting their sensitive data. Trusted hardware features have recently hit the mainstream through their inclusion in commodity CPUs. We witness a rising interest in how this new hardware can be used to build trustworthy cloud applications.

This tutorial will provide an overview of the recent support in Intel CPUs for trusted execution of security-sensitive software using Software Guard Extensions (SGX). It will introduce the key concepts behind trusted execution and discuss the required support to create trustworthy applications in otherwise untrusted environments. It is structured in lectures on specific topics related to trusted hardware, a presentation of Intel SGX, and a hands-on practical exercise using the technology. The goal is to convey opportunities and challenges that the widespread availability of trusted execution features will bring, and how it will affect future cloud computing software.

Speaker's short biography: Marcelo Pasin is a researcher in the University of Neuchâtel (Switzerland) and an associate professor in the Arc Engineering School of the University of Applied Sciences and Arts Western Switzerland (HES-SO). He obtained a PhD degree in Computer Science from the National Polytechnic Institute of Grenoble (France, 1999), a master's degree in computer science from the Federal University of Rio Grande do Sul (Porto Alegre, Brazil, 1994) and an electrical engineering degree from the Federal University of Santa Maria (Brazil, 1988). Marcelo's research interests are large-scale distributed systems, with focus and publications on cloud computing security, resource management, fault- and intrusion tolerance, high performance computing and networking. He is a member of two research teams: the complex systems group in the University of Neuchatel, and the data analytics group in the Arc school of HES-SO.

Tutorial 2 : Singularity, an HPC Application Container

Speakers: Alexandre Dehne Garcia (INRA, France) and Martin Souchal (IN2P3, France)

Overview: Docker is mainstream container technology adopted by computer scientists and practitioners. There are, however, many other containers technologies with their advantages and drawbacks. Although Docker is scalable, enabling applications to be structured as micro-services, it is not suited to High Performance Computing (HPC). Singularity on the other hand strikes a balance considering maturity, ease of use and administration, functionality and security. In this tutorial the speakers will discuss container technologies suitable to HPC, present Singularity, and then offer out a hands-on section.

Speakers' short biographies: Alexandre Dehne Garcia, INRA (French National Institute for Agricultural Research) is a DevOps HPC system administrator engineer at CBGP Lab (Biology Center for Populations Husbandry) and INGENUM (INRA E-infrastructures for research).

Martin Souchal, IN2P3 (Institut National de Physique Nucléaire et de Physique des Particules) is a DevOps system administrator at APC Lab (Astroparticule et Cosmologie) and the technical manager of The François Arago Center, an HPC platform for space.

Both are part of ComputeOps (IN2P3 project on HPC research and development).

Tutorial 3: Deploying Distributed Computing Experiments using Grid'5000 / SILECS Testbed

Speaker : Simon Delamare (CNRS, ENS de Lyon, France), Frédéric Desprez (Inria, LIG, France) and David Loup (CNRS, ENS de Lyon, France)

Overview: This tutorial will present the Grid'5000 / SILECS platform and introduce its usage to set up distributed computing experiments. SILECS is a scientific instrument designed to support experiment-driven research in many areas of computer science. Among available SILECS' testbeds, Grid'5000 focuses on Cloud, HPC and other parallel and large-scale distributed computing deployments. Grid'5000 gives to users an easy access to a large number of hardware resources with high capacity of customization at every level of computer stack (from hardware to high level applications). Grid'5000 allows users to set-up large-scale and reproducible experiments. The platform is supported by major French scientific institutions (Inria, CNRS, Renater, Universities) and its official Grid'5000 website is: https://www.grid5000.fr.

SBAC-PAD 2018 Location

Venue : Institut Français de l'Education (19 Allée de Fontenay, 69007 Lyon)

From the subway (Line B, Debourg Station) to the conference site (IFE, Institut français de l'éducation), see the figure on the left below. The IFE site is shown on the figure on the right. The blue line represents the path to the restaurant.

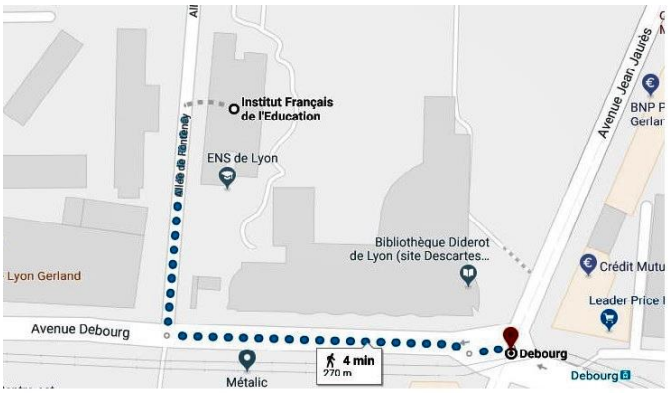

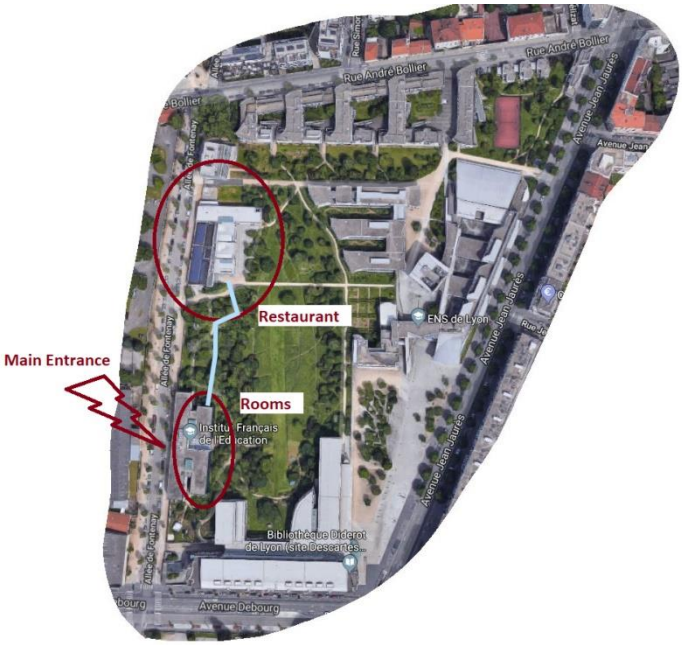

Social Event
Thursday, September 27, 2018
Banquet Reception on the Hermes boat

Best Paper
Award Ceremony

Announcement of
SBAC-PAD 2019

How to reach the Hermes boat (16 Quai Claude Bernard)

Catch the Tram T1 at the ENS Lyon stop towards Perrache (blue bubbles on left picture below),
Leave the Tram at the Quai Claude Bernard stop,
Walk from there to the entrance of the Hermes Boat (blue bubbles on right picture below).

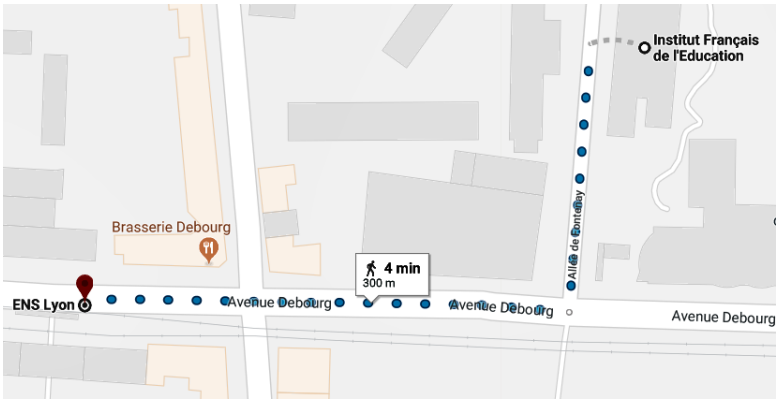

SBAC-PAD 2018 Sponsors

Organized by

Promoted by

Co-Sponsored by