

1st ISC HPC International Workshop on “Monitoring and Operational Data Analytics” (MODA)

June 25, 2020

Abstract

This workshop aims to provide insight into the current state and trends in Monitoring and Operational Data Analytics (MODA), to identify potential gaps, and to offer an outlook into the future of MODA at large scale together with possible solutions for the upcoming Exascale systems. The focus of the 1st MODA workshop will be on currently envisioned solutions and best practices for monitoring systems at data centers and HPC sites, as well as on effective strategies for analyzing and interpreting the collected operational data. We envision a full day workshop with a balanced mix between technical paper presentations, keynote and invited talks. The workshop is unique to the European HPC arena being the first to address this topic. The load of reviewing the submitted papers will be balanced among the members of the large and diverse program committee. This will also ensure a high quality of the reviews. We have the positive confirmation that Prof. Martin Schulz (TUM and LRZ, Germany) will be the plenary keynote speaker at MODA'20 and will share insights from his work on this topic conducted at LLNL (USA) as well as from his current work at LRZ (Germany).

1. Workshop scope

The race to Exascale poses significant challenges for the collection and analysis of the vast amount of data that future HPC systems will produce, in terms of the increasing complexity of the machines, the scalability and intrusiveness of the adopted monitoring solution, and the interpretability and effective inference driven by the acquired data. The main scope of the **1st ISC-HPC International Workshop on Monitoring and Operational Data Analytics (MODA)** is to provide insight into current trends in MODA, to identify potential gaps, and to offer an outlook into the future of the involved fields high performance-computing, databases, machine learning, and possible solutions for upcoming Exascale systems. Contributions matching the scope of the workshop will be related to:

- a) Currently envisioned solutions and practices for monitoring systems at data centers and HPC sites. Significant focus will be placed on *operational data collection mechanisms* respectively i) covering different system levels, from building infrastructure sensor data to CPU-core performance metrics, and ii) targeting different end-users, from system administrators to application developers and computational scientists.
- b) Effective strategies for *analyzing and interpreting the collected operational data*. Such strategies should particularly include (but are not limited to) different visualization approaches and machine learning-based techniques, potentially inferring knowledge of the system behavior and allowing for the realization a proactive control loop.

This workshop is not targeting new solutions proposed in the context of application performance modeling and/or application performance analysis tools. Novel contributions in the area of compiler analysis, debugging, programming models and/or sustainability of scientific software are also considered out of the scope of the workshop.

MODA is becoming common practice at various international HPC sites. However, each site adopts a different, insular approach, rarely adopted in production environments and mostly limited to the visualisation of system and building infrastructure metrics for health check purposes. In this regard, *we observe a gap between the collection of operational data and its meaningful and effective analysis and exploitation*, which prevents the closing of the feedback loop between the monitored HPC system, its operation, and its end-users. Under these premises, the goals of the workshop can then be summarized in the following way:

- 1) *Gather and share knowledge and establish a common ground* within the international community with respect to best practices in monitoring and operational data analytics.
- 2) *Discuss future strategies and alternatives for MODA*, potentially improving existing solutions and envisioning a common baseline approach in HPC sites and data centers.
- 3) *Establish a debate on the usefulness and applicability of AI techniques* on collected operational data for optimizing the operation of production systems (e.g. for practices such as predictive maintenance, runtime optimization, optimal resource allocation and scheduling).

1.1 Topics of interest

The contributions submitted to MODA will ideally address:

- State-of-the-practice method, tools, techniques in monitoring at various HPC sites
- Solutions for monitoring and analysis of operational data that work very well on large- to extreme-scale systems with a large number of users
- Solutions that have proven limitations in terms of efficiency of operational data collection in real-time or in terms of the quality of the collected data
- Opportunities and challenges of using machine learning methods for efficient monitoring and analysis of operational data
- Integration of monitoring and analysis practices into production system software (energy and resource management) and runtime systems (scheduling and resource allocation)
- Discuss explicit gaps between operational data collection, processing, effective analysis, highly useful exploitation, and propose new approaches to closing these gaps for the benefit of improving HPC centre planning, operations, and research
- Other monitoring and operational data analysis challenges and approaches (data storage, visualization, integration into system software, adoption)

1.2 Submission and publication

We will solicit original contributions in the form of original papers (6-12 pages) which will be peer-reviewed by the program committee members. All accepted papers will be presented during the workshop. We aim at a minimum of 4 and a maximum of 8 accepted papers, for the 8 x 30 minutes slots in the tentative workshop program (see Section 4).

We will publish the workshop papers together with the ISC 2020 proceedings, including an abstract of the keynote and invited talks, and a short white paper of the panel session.

High quality contributions may be considered for a full-length submission to a special journal issue in collaboration with ParCo, CPE, or other journals.

2. Relevance and impact for ISC 2020 and beyond

Data analytics techniques and artificial intelligence are important topics in the context of scientific computing ([Scalable Data Analytics in Scientific Computing workshop](#)¹) and large scale simulations ([Second Workshop on the Convergence of Large Scale Simulation/HPC and Artificial Intelligence](#)²). The proposed workshop addresses these challenges in the context of Operational Data Analytics, with a focus on analyzing and interpreting Monitoring data.

This workshop aligns with the goals of the [HPCMASPA Workshop](#)³, which was held as part of the IEEE 2019 Cluster Conference, in the United States. Similarly, we expect the MODA workshop to attract an important community at ISC in 2020 and beyond, which up to the present did not have a dedicated venue in Europe to discuss the latest results, challenges, and exchange ideas on the topics of HPC monitoring, logging, analysis and use of ML methods for these purposes.

The organizers and the program committee members bring together world-renowned scientists in the field as well as experts from important HPC centers in Europe and worldwide. This is a testament that MODA not only bridges between HPC and analytics, but also succeeds in bringing together academia and data and compute research centers. This constellation will ensure the fulfillment of the workshop goals.

3. Committee

3.1 Workshop organizing committee

- 1 Florina Ciorba (florina.ciorba@unibas.ch) – University of Basel, Switzerland
- 2 Nicolas Lachiche (nicolas.lachiche@unistra.fr) - University of Strasbourg, France
- 3 Aurélien Cavelan (aurelien.cavelan@unibas.ch) - University of Basel, Switzerland
- 4 Daniele Tafani (Daniele.Tafani@lrz.de) - Leibniz Supercomputing Centre, Germany
- 5 Utz-Uwe Haus (uhaus@cray.com) - Cray EMEA Research Lab, Switzerland

3.2 Tentative program committee

- 1 Andrea Bartolini (a.bartolini@unibo.it) - University of Bologna, Italy
- 2 Valeria Bartsch (valeria.bartsch@itwm.fraunhofer.de) - Fraunhofer ITWM Kaiserslautern, Germany
- 3 Norm Bourassa (njbourassa@lbl.gov) - NERSC LBNL, USA
- 4 Jim Brandt (brandt@sandia.gov) - Sandia National Labs, USA

¹ <https://sdasconf.github.io>

² <https://sites.google.com/nvidia.com/hpc-ai-2019>

³ <https://sites.google.com/site/hpcmaspa2019/home>

- 5 Rubén Cabezón (ruben.cabazon@unibas.ch) - sciCORE, University of Basel, Switzerland
- 6 Carlo Cavazzoni (c.cavazzoni@cineca.it) - CINECA, Italy
- 7 Thomas Ilsche (thomas.ilsche@tu-dresden.de) - Technische Universität Dresden, Germany
- 8 Jacques-Charles Lafoucriere (Jacques-Charles.LAFOUCRIERE@CEA.FR) -- CEA, France
- 9 Erwin Laure (erwinl@pdc.kth.se) - KTH, Sweden
- 10 Fiilippo Mantovani (filippo.mantovani@bsc.es) - BSC, Spain
- 11 Ariel Oleksia (ariel@man.poznan.pl) - Poznan Supercomputing Center, Poland
- 12 Melissa Romanus (MRomanus@lbl.gov) - NERSC LBNL, USA
- 13 Karthee Sivalingam (ksivalinga@cray.com) - CRAY EMEA Research Lab, UK
- 14 Heiko Schuldt (heiko.schuldt@unibas.ch) - University of Basel, Switzerland
- 15 Martin Schulz (schulzm@in.tum.de) - TU Munich / Leibniz Supercomputing Centre, Garching, Germany
- 16 Keiji Yamamoto (keiji.yamamoto@riken.jp) - RIKEN, Japan

4. Format of the workshop

The MODA Workshop is meant initially as a full-day workshop offering a unique platform for a combination of high quality contributed manuscripts, invited talks, and a panel discussion. Also, we will ensure a balanced program and even load of the reviewers to provide the required quality of the paper reviews by having a sufficiently large and diverse program committee.

09:00 - 09:05	Introduction (Workshop Organizers)
09:05 - 09:50	Keynote presentation - confirmed Martin Schulz, TU Munich / Leibniz Supercomputing Centre, Germany
09:50 - 10:00	Keynote Q&A Session
10:00 - 10:30	Paper presentation 1
10:30 - 11:00	Paper presentation 2
11:00 - 11:30	<i>Coffee Break</i>
11:30 - 12:00	Paper presentation 3
12:00 - 12:30	Paper presentation 4
12:30 - 14:00	<i>Lunch Break</i>
14:00 - 14:30	Invited talk 1 (LASSi -- Analysis of parallel IO at Archer; Cray CERL)
14:30 - 15:00	Invited talk 2 (TBD)
15:00 - 15:30	Invited talk 3 (TBD)
15:30 - 16:00	Invited talk 4 (TBD)
16:00 - 16:30	<i>Coffee Break</i>
16:30 - 17:30	Panel Discussion
17:30 - 17:45	Closing (Workshop Organizers)

5. Expected outcomes

The workshop aims to establish a forum of discourse and exchange ideas of current practices in HPC monitoring, logging, analysis and the integration of data analytics methods using machine-learning. The outcomes will include post-workshop proceedings and an up to date overview of existing efforts in Europe and beyond.

6. Outreach strategies

We will advertise the MODA workshop and attract both high quality submissions and attendees via the following channels:

- Social Media: LinkedIn (user re/posts, special interest groups, hashtags), Twitter (user re/posts, hashtags)
- Professional societies: ACM, IEEE, SIAM, and their special interest groups
- Community multipliers: Ann Gentile and Jim Brandt's Google Group on Monitoring Large-Scale HPC Systems⁴
- Mailing lists: HPC-Announce, TCPP-Announce, SIAM-SC, etc.
- Direct contact: Colleagues emails, word of mouth, Slack/Mattermost workspaces, Conference presentations

7. Website

The website for the workshop including the program and speakers will be created under: <https://moda20.sciencesconf.org/>

⁴ <https://groups.google.com/forum/#!forum/monitoring-large-scale-hpc-systems>