



18th Workshop on Workflows in Support of Large-Scale Science (WORKS)

Welcome from the Chairs

Silvina Caino-Lores, French Institute for Research in Computer Science and Automation (Inria), silvina.caino-lores@inria.fr

Anirban Mandal, Renaissance Computing Institute (RENCI), UNC-Chapel Hill, anirban@renci.org



November 12-13, Denver, CO, USA

Co-located with The International Conference for High Performance Computing, Networking, Storage, and Analysis

Workshop Overview

- Submissions: 29 → All-time record!
 - Full papers: 21
 - Abstracts: 8
- Accepted: 17
 - Full papers: 12 (57%)
 - Abstracts: 5 (63%)
- Reviews per submission: 3.5

- Full papers and combined paper with materials from abstract papers published in the Workshop Proceedings of SC
 - <https://dl.acm.org/doi/proceedings/10.1145/3624062>
- Two invited talks
- **Upgraded to full-day workshop consisting of two parts**



Part I (Sunday 12th, 2pm to 5:30pm, Rooms 501-502)

14:00	14:10	Welcome - Part I
14:10	14:42	Workflow Building Blocks: The Success Story of Environmental Modeling, HPC, and AI for Predicting Farmed Seafood Bacteria Contamination <i>Montella</i>
14:42	15:00	End-to-end Workflows for Climate Science: Integrating HPC Simulations, Big Data Processing and Machine Learning <i>Elia, Scardigno, Ejarque, D'Anca, Accarino, Scoccimarro, Donno, Peano, Immorlano, Aloisio</i>
15:00	15:30	Break
15:30	15:48	Scale Composite BaaS Services With AFCL Workflows <i>Larcher, Ristov</i>
15:48	16:06	A Systematic Mapping Study of Italian Research on Workflows <i>Aldinucci, Baralis, Cardellini, Colonnelli, Danelutto, Decherchi, Di Modica, Ferrucci, Gribaudo, Iannone, Lapegna, Medic, Muscianisi, Righetti, Sciacca, Tonello, Tortonesi, Trunfio, Vardanega</i>
16:06	16:16	Transcriptomics Atlas Pipeline: Cloud vs HPC <i>Kica, Lichołaj, Malawski</i>
16:16	16:26	Patterns and Anti-Patterns in Migrating from Legacy Workflows to Workflow Management Systems <i>Cassol, Froula, Kirton, Sul, Melara, Kothadia, Player, Sarrafan, Chan, Fagnan</i>
16:26	16:44	Accelerating Data-Intensive Seismic Research Through Parallel Workflow Optimization and Federated Cyberinfrastructure <i>Adair, Rodero, Parashar, Melgar</i>
16:44	17:02	Laminar: A New Serverless Stream-based Framework with Semantic Code Search and Code Completion <i>Zahra, Li, Filgueira</i>
17:02	17:20	Optimization towards Efficiency and Stateful of dispel4py <i>Liang, Zhang, Yang, Heinis, Filgueira</i>
17:20	17:30	Wrap Up - Part I

Part II (Monday 13th, 9am to 12:30pm, Rooms 704-706)

09:00	09:05	Welcome - Part II
09:05	09:37	FAIRIST of Them All: Meeting Researchers Where They Are With Just-in-Time, FAIR Implementation Advice <i>Kirkpatrick</i>
09:37	09:55	A data science pipeline synchronisation method for edge-fog-cloud continuum <i>Sanchez-Gallegos, Gonzalez-Compean, Carretero, Marin-Castro</i>
09:55	10:25	Break
10:43	10:53	Leveraging Large Language Models to Build and Execute Computational Workflows <i>Duque, Syed, Day, Berry, Katz, Kindratenko</i>
10:53	11:11	Delivering Rules-Based Workflows for Science <i>Marchant, Blomqvist, Jensen, Lilholm, Nørgaard</i>
11:11	11:29	Julia as a Unifying End-to-End Workflow Language on the Frontier Exascale System <i>Godoy, Valero-Lara, Anderson, Lee, Gainaru, Ferreira da Silva, Vetter</i>
11:29	11:39	Scaling on Frontier: Uncertainty Quantification Workflow Applications using ExaWorks to Enable Full System Utilization <i>Titov, Carson, Rolchigo, Coleman, Belak, Bement, Laney, Turilli, Jha</i>
11:39	11:57	Distributed Data Locality-Aware Job Allocation <i>Markovic, Kolovos, Soares Indrusiak</i>
11:57	12:15	Fluxion: A Scalable Graph-Based Resource Model for HPC Scheduling Challenges <i>Patki, Ahn, Milroy, Yeom, Garlick, Grondona, Herbein, Scogland</i>
12:15	12:25	The Common Workflow Scheduler Interface: Status Quo and Future Plans <i>Lehmann, Bader, Thamsen, Leser</i>
12:25	12:30	Wrap Up - Part II

Invited talk

Paper talk

Lightning talk



Part I (Sunday 12, 2pm to 5:30pm, Rooms 501-502)

Part II (Monday 13, 9am to 12:30pm, Rooms 704-706)

14:00	14:10	Welcome - Part I
14:10	14:42	Workflow Building Blocks: The Success Story of Environmental Modeling, HPC, and AI for Predicting Farmed Seafood Bacteria Contamination <i>Montella</i>
14:42	15:00	End-to-end Workflows for Climate Science: Processing and Machine Learning <i>Elia, Scardigno, Ejarque, D'Anca, Accarino, Sc</i>
15:00	15:30	Break
15:30	15:48	Scale Composite BaaS Services With AFCL W <i>Larcher, Ristov</i>
15:48	16:06	A Systematic Mapping Study of Italian Resea <i>Aldinucci, Baralis, Cardellini, Colonnelli, Dane</i> <i>Iannone, Lapegna, Medic, Muscianisi, Righet</i> <i>Vardanega</i>
16:06	16:16	Transcriptomics Atlas Pipeline: Cloud vs HPC <i>Kica, Lichołai, Malawski</i>
16:16	16:26	Patterns and Anti-Patterns in Migrating from Systems <i>Cassol, Froula, Kirton, Sul, Melara, Kothadia,</i>
16:26	16:44	Accelerating Data-Intensive Seismic Research Federated Cyberinfrastructure <i>Adair, Rodero, Parashar, Melgar</i>
16:44	17:02	Laminar: A New Serverless Stream-based Framework with Semantic Code Search and Code Completion <i>Zahra, Li, Filgueira</i>
17:02	17:20	Optimization towards Efficiency and Stateful of dispel4py <i>Liang, Zhang, Yang, Heinis, Filgueira</i>
17:20	17:30	Wrap Up - Part I

09:00	09:05	Welcome - Part II
09:05	09:37	FAIRIST of Them All: Meeting Researchers Where They Are With Just-in-Time, FAIR Implementation Advice <i>Kirkpatrick</i>
		synchronisation method for edge-fog-cloud continuum <i>alez-Compean, Carretero, Marin-Castro</i>
		age Models to Build and Execute Computational Workflows <i>, Katz, Kindratenko</i>
		Workflows for Science <i>nsen, Lilholm, Nørgaard</i>
		to-End Workflow Language on the Frontier Exascale System <i>erson, Lee, Gainaru, Ferreira da Silva, Vetter</i>
		certainty Quantification Workflow Applications using l System Utilization <i>Coleman, Belak, Bement, Laney, Turilli, Jha</i>
		ty-Aware Job Allocation <i>es Indrusiak</i>
		oph-Based Resource Model for HPC Scheduling Challenges <i>, Garlick, Grondona, Herbein, Scogland</i>
12:15	12:25	The Common Workflow Scheduler Interface: Status Quo and Future Plans <i>Lehmann, Bader, Thamsen, Leser</i>
12:25	12:30	Wrap Up - Part II

More details on the website
<https://works-workshop.org/>



- Invited talk
- Paper talk
- Lightning talk



Join us for Part II tomorrow at 9am on rooms 704-706

More details on the website
<https://works-workshop.org/>



09:00	09:05	Welcome - Part II
09:05	09:37	FAIRIST of Them All: Meeting Researchers Where They Are With Just-in-Time, FAIR Implementation Advice <i>Kirkpatrick</i>
09:37	09:55	A data science pipeline synchronisation method for edge-fog-cloud continuum <i>Sanchez-Gallegos, Gonzalez-Compean, Carretero, Marin-Castro</i>
09:55	10:25	Break
10:43	10:53	Leveraging Large Language Models to Build and Execute Computational Workflows <i>Duque, Syed, Day, Berry, Katz, Kindratenko</i>
10:53	11:11	Delivering Rules-Based Workflows for Science <i>Marchant, Blomqvist, Jensen, Lilholm, Nørgaard</i>
11:11	11:29	Julia as a Unifying End-to-End Workflow Language on the Frontier Exascale System <i>Godoy, Valero-Lara, Anderson, Lee, Gainaru, Ferreira da Silva, Vetter</i>
11:29	11:39	Scaling on Frontier: Uncertainty Quantification Workflow Applications using ExaWorks to Enable Full System Utilization <i>Titov, Carson, Rolchigo, Coleman, Belak, Bement, Laney, Turilli, Jha</i>
11:39	11:57	Distributed Data Locality-Aware Job Allocation <i>Markovic, Kolovos, Soares Indrusiak</i>
11:57	12:15	Fluxion: A Scalable Graph-Based Resource Model for HPC Scheduling Challenges <i>Patki, Ahn, Milroy, Yeom, Garlick, Grondona, Herbein, Scogland</i>
12:15	12:25	The Common Workflow Scheduler Interface: Status Quo and Future Plans <i>Lehmann, Bader, Thamsen, Leser</i>
12:25	12:30	Wrap Up - Part II

Invited talk
Paper talk
Lightning talk



Best Paper Distinction



November 12-13, Denver, CO, USA

Co-located with The International Conference for High Performance Computing, Networking, Storage, and Analysis

Best Paper Distinction

Julia as a Unifying End-to-End Workflow Language on the Frontier Exascale System

*William Godoy, Pedro Valero-Lara, Caira Anderson, Katrina Lee, Ana Gainaru,
Rafael Ferreira da Silva, Jeffrey Vetter*

Presentation during **Part II on Monday, November 13th at 11:11am in rooms 704-706**



Logistics

- We have a tight schedule, please adhere to the following time restrictions
 - Invited talks: 32-minute time slot (28+4)
 - Paper talks (i.e., full papers): 18-minute time slot (15+3)
 - Lightning talks (i.e., abstracts): 10-minute time slot (8+2)
- Online participants must submit their questions to the Digital Experience platform. Chair will monitor and select the questions to the room
- Make sure you test your laptop during the break preceding your presentation
- In case of technical difficulties, we will show the copy of the presentation you had submitted
- Presenters will come to the front to set up their presentation during the last minute of questions of the previous presentation



Thank you!

Program Committee

Ilkay Altintas (San Diego Supercomputing Center)
Ivona Brandic (Technical University of Vienna)
Kyle Chard (University of Chicago)
Jesus Carretero (University Carlos III of Madrid)
Tainã Coleman (University of Southern California)
Ewa Deelman (University of Southern California)
Vincenzo de Maio (Technical University of Vienna)
Frank Di Natale (Nvidia)
Rosa Filgueira (University of St. Andrews)
Daniel Garijo (Universidad Politécnica de Madrid)
Sandra Gesing (University of Illinois Chicago)
Tristan Glatard (Concordia University)
William Godoy (Oak Ridge National Laboratory)
Shantenu Jha (Rutgers University)
Daniel S. Katz (University of Illinois at Urbana-Champaign)
Tamas Kiss (University of Westminster)
Jakob Luetzgau (University of Tennessee)
Ketan C. Maheshwari (Oak Ridge National Laboratory)
Raffaele Montella (University of Naples Parthenope)

Paolo Missier (Newcastle University)
Bogdan Nicolae (Argonne National Laboratory)
Paola Olaya (University of Tennessee)
Tom Peterka (Argonne National Laboratory)
Loïc Pottier (Lawrence Livermore National Laboratory)
Radu Prodan (University of Klagenfurt)
Omer Rana (Cardiff University)
Ivan Rodero (University of Utah)
Raul Sirvent (Barcelona Supercomputing Center)
Tyler Skluzacek (Oak Ridge National Laboratory)
Renan Souza (Oak Ridge National Laboratory)
Frédéric Suter (Oak Ridge National Laboratory)
Domenico Talia (University of Calabria)
Francois Tessier (INRIA Rennes)
Douglas Thain (University of Notre Dame)
Rafael Tolosana-Calasanz (Universidad de Zaragoza)
Sean R. Wilkinson (Oak Ridge National Laboratory)
Justin Wozniak (Argonne National Laboratory)
Orcun Yildiz (Argonne National Laboratory)

Steering Committee

David Abramson
Malcom Atkinson
Ewa Deelman
Michela Taufer

Session Chairs

Justin Wozniak
Frédéric Suter
Ketan C. Maheshwari
Jakob Luetzgau

SC organizers and volunteers



Thank you!

