



GraphSys-2024: 2nd Workshop on Serverless, Extreme-Scale, and Sustainable Graph Processing Systems

Chairs' Welcome

Alexandru Iosup
a.iosup@vu.nl
VU University of Amsterdam
Amsterdam, The Netherlands

Radu Prodan
radu.prodan@aau.at
Alpen-Adria-Universität Klagenfurt
Klagenfurt, Austria

Ana-Lucia Varbanescu
a.l.varbanescu@utwente.nl
University of Twente
Enschede, The Netherlands

Welcome!

It is our great pleasure to welcome you to *GraphSys'24*, the 2nd edition of the *ACM/SPEC Workshop on Serverless, Extreme-Scale, and Sustainable Graph Processing Systems*. This is a returning workshop, where we continue to facilitate the exchange of ideas and expertise in the broad field of high-performance large-scale graph processing.

Graphs and GraphSys

The use, interoperability, and analytical exploitation of graph data are essential for modern digital economies. Today, thousands of computational methods (algorithms) and findable, accessible, interoperable, and reusable (FAIR) graph datasets exist. However, current computational capabilities lag when faced with the complex workflows involved in graph processing, the extreme scale of existing graph datasets, and the need to consider sustainability metrics in graph-processing operations. Needs are emerging for graph-processing platforms to provide multilingual information processing and reasoning based on the massive graph representation of extreme data in the form of general graphs, knowledge graphs, and property graphs. Because graph workloads and graph datasets are strongly irregular, and involve one or several big data “Vs” (e.g., volume, velocity, variability, vicissitude), the community needs to reconsider traditional approaches in performance analysis and modeling, system architectures and techniques, serverless and “as a service” operation, real-world and simulation-driven experimentation, etc., and provide new tools and instruments to address emerging challenges in graph processing.

Graphs or linked data are crucial to innovation, competition, and prosperity and establish a strategic investment in technical processing and ecosystem enablers. Graphs are universal abstractions that capture, combine, model, analyze, and process knowledge about real and digital worlds into actionable insights through item representation and interconnectedness. For societally relevant problems, graphs are extreme data that require further technological innovations to meet the needs of the European data economy. Digital graphs help pursue the United Nations Sustainable Development

Goals (UN SDG) by enabling better value chains, products, and services for more profitable or green investments in the financial sector and deriving trustworthy insight for creating sustainable communities. All science, engineering, industry, economy, and society-at-large domains can leverage graph data for unique analysis and insight, but only if graph processing becomes easy to use, fast, scalable, and sustainable.

GraphSys is a cross-disciplinary meeting venue focusing on state-of-the-art and the emerging (future) graph processing systems. We invite experts and trainees in the field, across academia, industry, governance, and society, to share experience and expertise leading to a shared body of knowledge, to formulate together a vision for the field, and to engage with the topics to foster new approaches, techniques, and solutions.

Technical content

In this second edition, the workshop received 10 submissions from which, after peer review, it accepted six full papers for publication. Topics include graph sampling characterization, new models for pipeline/streaming computation on graphs, and using graphs for thermal modeling in HPC centers, telemetry data analysis, and improvements in collision detection for autonomous vehicle driving. We are sure this mix of topics will be reflected in a lively workshop, from presentations to discussions. We are looking forward to it! We further invited the other four submissions to submit a work-in-progress short papers. Three teams have accepted our invitation, and thus the workshop will feature three short, work-in-progress talks on creating massive knowledge graphs, state-of-the-art in serverless workflow management, and graph-sampling algorithms. Last, but not least, we are happy to welcome two invited speakers to our workshop. Dr. Johannes Langguth (SIMULA, Norway) will talk about *Graph Algorithms on Emerging Tile-Centric Accelerators*. Dr. Gabor Szarnyas from the LDBC (Linked Data Benchmark Council, The Netherlands) will talk about *Linked Data Benchmark Council: 12 years of fostering competition in the graph processing space*.

Thanks and Acknowledgments

This year's GraphSys is the result of the collaboration of many people, authors, reviewers, and the organizing committee. Thank you! We further thank the ICPE 2024 Program and Workshop Chairs, for their active support. We could not have done this without you! The GraphSys workshop is technically sponsored by the Graph-Massivizer project funded by the Horizon Europe research and innovation program of the European Union for the period 2024-2026,

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPE Companion '24, May 7–11, 2024, London, United Kingdom

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0445-1/24/05.

<https://doi.org/10.1145/3629527.3651435>

which studies and aims to develop a high-performance, scalable, gender-neutral, secure, and sustainable platform for massive graph processing.

Concluding remarks

We aim to continue to align GraphSys with the Standard Performance Evaluation Corporation (SPEC)'s Research Groups (RG), and,

in particular, the RG Cloud Group that is taking a broad approach, relevant for both academia and industry, to cloud benchmarking, quantitative evaluation, and experimental analysis.

We hope GraphSys will continue its growth towards a focused yearly series, aiming to develop the topic of large scale high performance, sustainable graph processing, and around it a community of knowledge and practice.