EuroSys 2025

editorial

Lawall, Julia; Chen, Haibo; Giuffrida, Cristiano; Liang, Kaitai; Smaragdakis, Georgios

Publication date
2025

Document Version
Final published version

Published in
EuroSys '25: Proceedings of the Twentieth European Conference on Computer Systems

Citation (APA)
Lawall, J., Chen, H., Giuffrida, C., Liang, K., & Smaragdakis, G. (Eds.) (2025). EuroSys 2025: editorial. In *EuroSys '25: Proceedings of the Twentieth European Conference on Computer Systems* (pp. ix-xv). (EuroSys 2025 - Proceedings of the 2025 20th European Conference on Computer Systems).

March 30–April 3, 2025
Rotterdam, Netherlands

Association for
Computing Machinery

Advancing Computing as a Science & Profession

EuroSys '25

Proceedings of the 2025

Twentieth European Conference on Computer Systems

Sponsored by:

ACM SIGOPS

Supported by:

Huawei (diamond sponsor); Amazon Web Services and Microsoft (silver sponsors);AMD, Arm, Google, ING, Intel, MangoBoost, Qualcomm, Red Hat, Rotterdam Partners, Samsung, SK Hynix, VMWare, Western Digital (bronze sponsors)

Association for
Computing Machinery

Advancing Computing as a Science & Profession

The Association for Computing Machinery
2 Penn Plaza, Suite 701
New York, New York 10121-0701

ISBN: 979-8-4007-1196-1

Additional copies may be ordered prepaid from:
ACM Order Department
PO Box 30777
New York, NY 10087-0777, USA

Phone: +1 800 342-6626 (USA and Canada)
+1 212 626-0500 (Global)
Fax: +1 212 944-1318
Email: acmhelp@acm.org
Hours of Operation: 8:30 am–4:30 pm ET

Contents

EuroSys 2025

We are delighted to welcome you to EuroSys 2025, the 20th edition of the European Conference on Computer Systems! We are excited to host EuroSys 2025 in the modern and dynamic city of Rotterdam, Netherlands. This year's EuroSys is very special as it is co-located (for the first time) with the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2025). We hope you will enjoy an excellent technical program, engaging discussions, and networking opportunities in this vibrant city known for its innovative architecture, bustling port, and rich cultural scene.

Main Conference Program: This year's program features papers that advance the state of the art across a range of systems research topics -- operating systems, virtualization, ML systems, cloud and serverless computing, LLM serving, GPU systems, bug finding, networking, distributed systems, graph and stream processing, fault tolerance, and storage.

EuroSys 2025 followed the example of EuroSys 2024 and implemented a dual-deadline format with separate submission cycles in the spring and in the fall. Across the two deadlines, we received 696 submissions, a 42% increase in the number of submissions over last year. The submissions were reviewed by a program committee of 122 members (chairs excluded). Each valid submission received 3-6 reviews across two rounds of reviewing for a total of 2,571 reviews. Overall, we accepted 85 papers for an acceptance rate of 12.4%. While the acceptance rate is somewhat lower than in previous years, we found that the difference primarily stems from fewer papers passing to the second round of reviews. The acceptance rate in the second round is comparable to prior installments of EuroSys. All accepted papers were assigned a shepherd from the program committee, with acceptance of the final version for publication being conditional on shepherd's approval. Like prior editions, EuroSys 2025 also provided one-shot revisions. Twenty-five papers across the two deadlines (corresponding to 3.6% of all submissions) were offered the opportunity to submit a revised version of the paper to the following EuroSys deadline. Nine of the papers submitted to the spring deadline were revised papers from the EuroSys 2024 fall deadline, and 14 papers submitted to the fall deadline were revised papers

from the EuroSys 2025 spring deadline. All such revised papers embodied significant work that successfully addressed the reviewers' requirements, and thus were finally accepted. EuroSys 2025 offers two parallel sessions (going multi-track for the first time) and ASPLOS 2025 offers four parallel sessions. Attendees can freely attend any of the sessions.

EuroSys 2025 also features a poster session, hosted during the welcome reception. ASPLOS 2025 offers poster sessions across the three main conference days, where attendees can have in-depth discussions with the leading authors of the ASPLOS 2025 papers. To foster social interactions between the ASPLOS and EuroSys communities, the joint program also features a banquet at the SS Rotterdam, a former ocean liner and cruise ship.

Artifact Evaluation (AE): EuroSys 2025 includes an Artifact Evaluation (AE) process, organized by Thaleia Dimitra Doudali (IMDEA Software Institute, Spain), Christian Pinto (IBM Research Europe, Ireland), and Georgios Portokalidis (IMDEA Software Institute, Spain). The AE process was supported by a dedicated committee of 98 volunteers. Authors of accepted papers were invited to submit artifacts related to their research, including source code, systems, datasets, models, test suites, benchmarks, and other materials supporting their contributions. Each submission was rigorously evaluated by two to four AE committee members to assess its availability, functionality, and reproducibility. This year, the AE committee received submissions from 45 accepted EUROSYS papers (15 in the spring round and 30 in the fall round), with a total of 44 submissions to the Artifacts Available badge, 43 submissions to the Artifacts Functional badge and 28 to the Results Reproduced Badges. The evaluation process resulted in an acceptance rate of 100% for Artifacts Available badges (44 granted out of 44 submissions), 98% for Artifacts Functional badges (42 granted out of 43 submissions) and 75% for Results Reproduced badges (21 granted out of 28 submissions).

Shadow PC: In parallel to the main PC, we also organized a Shadow PC. The Shadow PC was chaired by Antonio Barbalace (The University of Edinburgh, Scotland) and Baptiste Lepers (Inria, France). The Shadow PC is an integral part of the EuroSys community-building exercise: serving on a Shadow PC is an excellent opportunity for young systems researchers (PhD students, postdocs, and new faculty members) to gain experience in program committee practices. The 2025 Shadow PC included 71 members evenly split across three continents: a third from Europe, a third from the US and Canada, and a third from Asia. The Shadow PC received 63 papers from the main track. All received papers were reviewed extensively by the Shadow PC members and got between 4 and 5 reviews each. The Shadow PC reached decisions with online discussions on HotCRP (no virtual PC meeting was needed). The Shadow PC finally decided to accept 21 papers, and to revise 7 papers (~44% acceptance ratio including the revisions).

Workshops: In addition to the main program, EuroSys 2025 and ASPLOS 2025 jointly feature 34 workshops and tutorials, coordinated by ASPLOS 2025 workshop chairs Magnus Jahre

(Norwegian University of Science and Technology, Norway) and Alexandra Jimborean (University of Murcia, Italy), and EuroSys 2025 workshop chairs Asterios Katsifodimos (TU Delft, Netherlands) and Soham Chakraborty (TU Delft, Netherlands). Among these, 9 workshops are specific to EuroSys 2025, and 5 workshops are relevant for both EuroSys and ASPLOS communities. These workshops cover a large range of topics reflecting the breadth of the EuroSys community and include well-established workshops as well as one workshop that is being offered for the first time. Continuing from previous years, we have the 5th Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS), the 8th International Workshop on Edge Systems, Analytics and Networking (EdgeSys), the 20th EuroSys Doctoral Workshop (EuroDW), the 5th Workshop on Machine Learning and Systems (EuroMLSys), the 18th European Workshop on System Security (EuroSec), the 12th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC), and the 3rd Workshop on SErverless Systems, Applications and MEthodologies (SESAME). New this year is the 1st Workshop on Systems and Architectures for Encrypted AI (SAFE-AI). We look forward to the engaging discussions and innovative ideas these workshops will foster.

The EuroDW Doctoral Workshop holds a special place at EuroSys, serving as an essential platform to foster the growth of young systems researchers. This year's EuroSys Doctoral Workshop is chaired by Haitham Al Hassanieh (EPFL, Switzerland) and Sanidhya Kashyap (EPFL, Switzerland). The 20th EuroDW aims to provide a supportive environment for PhD students to present their work and receive valuable feedback from seasoned professionals and fellow students alike. In addition to technical presentations, the workshop will offer insights and discussions about pursuing a PhD, conducting research, and exploring various career paths. PhD students at any stage of their doctoral journey are welcome to participate and benefit from this enriching experience.

Acknowledgements: We express our deepest gratitude to the numerous individuals who contributed to the successful organization of EuroSys 2025. The core organization team comprises: the local organization chair Lucienne Dado (Support Delft Events, The Netherlands); the workshop chairs Soham Chakraborty (Delft University of Technology, The Netherlands) and Asterios Katsifodimos (Delft University of Technology, The Netherlands); the poster chairs Nitinder Mohan (Delft University of Technology, The Netherlands), Harm Griffioen (Delft University of Technology, The Netherlands), and Klaus von Gleissenthall (Vrije Universiteit Amsterdam, The Netherlands); the proceedings chairs Alexios Voulimeneas (Delft University of Technology, The Netherlands) and Habib Mostafaei (Eindhoven University of Technology, The Netherlands); the web chair Roland Kromes, (Delft University of Technology, The Netherlands); the artifact evaluation chairs Thaleia Doudali (IMDEA Software Institute, Spain) Georgios Portokalidis (IMDEA Software Institute, Spain) Christian Pinto (IBM Research Europe, Ireland); the shadow PC chairs Antonio Barbalace (University of Edinburgh, Scotland) and Baptiste Lepers (Université de Neuchâtel, Switzerland); the doctoral workshop chairs Haitham Al

Hassanieh (EPFL, Switzerland) and Sanidhya Kashyap (EPFL, Switzerland); the travel grant chairs Asia Slowinska (Vrije Universiteit Amsterdam, The Netherlands) and Daniele Cono D'Elia (Sapienza University of Rome, Italy); the publicity chairs Manolis Stamatogiannakis (Vrije Universiteit Amsterdam, The Netherlands) and Vijay Chidambaram (UT Austin, USA); the sponsorship chairs Minsoo Rhu (KAIST, South Korea), Rakesh Kumar (Norwegian University of Science and Technology, Norway), and Benjamin Lee (University of Pennsylvania, USA); the registration chair Enrico Bassetti (Delft University of Technology, The Netherlands); and the AV chair Aleksandar Markovic (Vrije Universiteit Amsterdam, The Netherlands). Additionally, we would like to extend our thanks to the EuroSys Officers and the Steering Committee for their continued support throughout this process.

We would also like to thank the ASPLOS steering committee and organizers for the excellent collaboration that made it possible to co-locate the two conferences for the first time and bring together researchers who work on similar topics.

We also want to acknowledge the generous support by the EuroSys 2025 sponsors: Huawei (diamond sponsor); Amazon Web Services and Microsoft (silver sponsors); AMD, Arm, Google, ING, Intel, MangoBoost, Qualcomm, Red Hat, Rotterdam Partners, Samsung, SK Hynix, VMWare, Western Digital (bronze sponsors).

Finally, we would like to thank the Vrije Universiteit Amsterdam and the Delft University of Technology for their generous support and the Postillion Hotel & Convention Centre WTC Rotterdam for hosting the conference.

Julia Lawall (Inria), Program Co-Chair
Haibo Chen (Shanghai Jiao Tong University), Program Co-Chair
Cristiano Giuffrida, Vrije Universiteit Amsterdam, The Netherlands, General Co-Chair
Kaitai Liang, Delft University of Technology, The Netherlands, General Co-Chair
Georgios Smaragdakis, Delft University of Technology, The Netherlands, General Co-Chair

Program Committee

Program Committee Chairs

Julia Lawall (Inria)
Haibo Chen (Shanghai Jiao Tong University)

Program Committee Members

Ada Gavrilovska, Georgia Tech
Adil Ahmad, Arizona State University
Aishwarya Ganesan, UIUC

Alain Tchana, Grenoble INP
Alexandre da Silva Veith, Nokia
Ali Anwar, University of Minnesota
Andreas Schmidt, Saarland University
Anjo Vahldiek-Oberwagner, Intel Labs
Antoine Kaufmann, Max Plank Institute for Software Systems (MPI-SWS)
Anton Burtsev, University of Utah
Antonio Barbalace, The University of Edinburgh
Anurag Khandelwal, Yale University
Ashvin Goel, University of Toronto
Baptiste Lepers, The University of Neuchâtel
Beomseok Nam, Sungkyunkwan University
Boris Grot, University of Edinburgh, Huawei Research
Călin Iorgulescu, Oracle Labs
Chang Lou, University of Virginia
Cheng Li, University of Science and Technology of China
Cheng Tan, Northeastern University
Chia-Che Tsai, Texas A&M University
Christian Dietrich, Hamburg University of Technology
Cristina Nita-Rotaru, Northeastern University
Dan Williams, Virginia Tech
Daniel Lohmann, Leibniz Universität Hannover
Dezun Dong, National University of Defense Technology
Diogo Behrens, Huawei
Donald Porter, The University of North Carolina at Chapel Hill
Emmanuel Baccelli, Inria
Emmanuel Cecchet, University of Massachusetts Amherst
Erci Xu, Alibaba Group
Etienne Rivière, UCLouvain
Fan Lai, UIUC
Fan Yang, Microsoft Research Asia
Gaël Thomas, Inria Saclay
Haifeng Yu, National University of Singapore
Hong Zhang, University of Waterloo
Horst Schirmeier, TU Dresden
Hyungon Moon, UNIST
Indrajit Roy, Google Inc
Jacob Gorm Hansen, Jamscape
Jayashree Mohan, Microsoft Research
Jean-Pierre Lozi, Inria

Jeff Mogul, Google
Jia-Ju Bai, Beihang University
Jialin Li, National University of Singapore
Jian Huang, UIUC
Jia Rao, The University of Texas at Arlington
Jinyu Gu, Shanghai Jiao Tong University
John Wilkes, Google
Jon Crowcroft, University of Cambridge
Kai Chen, Hong Kong University of Science and Technology
Kang Chen, Tsinghua University
Keval Vora, Simon Fraser University
Laiping Zhao, TianJin University
Larry Rudolph, Two Sigma Investments, LP
Laurent Bindschaedler, Max Planck Institute for Software Systems
Lingxiao Ma, Microsoft Research Asia
Liuba Shrira, Brandeis University
Luo Mai, University of Edinburgh
Marcus Brandenburger, IBM Research Europe
Malte Schwarzkopf, Brown University
Mangpo Phothilimthana, Google DeepMind
Marc Shapiro, Sorbonne-Université–LIP6 & Inria
Marco Serafini, University of Massachusetts Amherst
Maria Carpen-Amarie, Huawei Research
Martin Kleppmann, University of Cambridge
Matteo Interlandi, Microsoft
Michael Swift, University of Wisconsin-Madison
Michio Honda, University of Edinburgh
Mingxing Zhang, Tsinghua University
Djob Mvondo, Univ Rennes, CNRS, INRIA, IRISA
Nuno Preguica, Universidade Nova de Lisboa
Olaf Spinczyk, Osnabrueck University
Patrick P. C. Lee, The Chinese University of Hong Kong
Philippe Bonnet, DIKU
Pierre Olivier, The University of Manchester
Pramod Bhatotia, Technical University of Munich
Qian Ge, Google DeepMind
Qingyang Wang, Louisiana State University
Redha Gouicem, RWTH Aachen University
Renaud Lachaize, Université Grenoble Alpes
Reto Achermann, University of British Columbia

Rishabh Iyer, UC Berkeley
Roberto Natella, University of Naples Federico II
Rong Chen, Shanghai Jiao Tong University
Sam H. Noh, Virginia Tech
Sanidhya Kashyap, EPFL
Sara Hamouda, Google
Serdar Tasiran, Amazon Web Services
Shimin Chen, ICT, CAS & University of Chinese Academy of Sciences
Shuai Mu, Stony Brook University
Soujanya Ponnapalli, UC Berkeley
Steve Blackburn, Google
Thaleia Dimitra Doudali, IMDEA Software Institute
Tianyin Xu, UIUC
Tim Harris, Microsoft
Timo Hönig, Ruhr-Universität Bochum (RUB)
Tobias Distler, FAU Erlangen-Nürnberg
Thomas Pasquier, UBC
Trinabh Gupta, UCSB
Valerio Schiavoni, University of Neuchatel
Vero Estrada-Galiñanes, EPFL
Wenguang Chen, Tsinghua University
Wenjun Hu, Yale University
Xiaosong Ma, Qatar Computing Research Institute
Y. Charlie Hu, Purdue University
Yang Wang, Meta/The Ohio State University
Yanyan Jiang, Nanjing University
Yérom-David Bromberg, University of Rennes, INRIA
Yiming Zhang, Xiamen University
Yizhou Shan, Huawei Cloud
Yongle Zhang, Purdue
Youngjin Kwon, KAIST
Youyou Lu, Tsinghua University
Yu Hua, Huazhong University of Science and Technology
Yu Liang, ETH Zurich
Yubin Xia, Shanghai Jiao Tong University
Yuvraj Patel, University of Edinburgh
Zhaoguo Wang, Shanghai Jiao Tong University
Zsolt István, TU Darmstadt

External Reviewers

Alexander Merritt, Unaffiliated
Ashraf Mahgoub, Purdue University
Bernard Wong, University of Waterloo
Deepak Narayanan, NVIDIA
Divya Mahajan, Georgia Institute of Technology
Evangelia Kalyvianaki, University of Cambridge
Guilherme Cox, NVIDIA
Indranil Gupta, University of Illinois Urbana-Champaign, USA
Junchen Jiang, University of Chicago
Luis Rodrigues, INESC-ID, IST, ULisboa
Lydia Y. Chen, TU Delft
Marcos K. Aguilera, VMware Research Group
Matthew Burke, Databricks
Miguel Matos, IST Lisbon
Mike Mesnier, Intel Labs
Ming Wu, Shanghai Tree-Graph Blockchain Research Institute
Minsoo Rhu, KAIST
Naama Ben-David, Technion
Panos Kalnis, King Abdullah University of Science and Technology
Patrick Stuedi, Meta
Poulami Das, UT Austin
Prashant J. Nair, The University of British Columbia
Qinbin Li, UC Berkeley
Ric Wheeler, IonQ
Robert Ladig, University of Southern Denmark
Roberto Palmieri, Lehigh University
Ronghui Gu, Columbia University
Samuel King, UC Davis
Shadi Noghabi, Microsoft Research
Swamit Tannu, U. of Wisconsin
Tanvir Ahmed Khan, Columbia University
Tianzheng Wang, Simon Fraser University

Organization Committee

General Chairs

Cristiano Giuffrida, Vrije Universiteit Amsterdam, The Netherlands

Kaitai Liang, Delft University of Technology, The Netherlands
Georgios Smaragdakis, Delft University of Technology, The Netherlands

Program Committee Chairs
Julia Lawall (Inria)
Haibo Chen (Shanghai Jiao Tong University)

Local Organization Chair
Lucienne Dado, Support Delft Events, The Netherlands

Workshop Chairs
Soham Chakraborty, Delft University of Technology, The Netherlands
Asterios Katsifodimos, Delft University of Technology, The Netherlands

Posters Chairs
Nitinder Mohan, Delft University of Technology, The Netherlands
Harm Griffioen, Delft University of Technology, The Netherlands
Klaus von Gleissenthall, Vrije Universiteit Amsterdam, The Netherlands

Proceedings Chairs
Alexios Voulimeneas, Delft University of Technology, The Netherlands
Habib Mostafaei, Eindhoven University of Technology, The Netherlands

Web Chair
Roland Kromes, Delft University of Technology, The Netherlands

Artifact Evaluation Chairs
Thaleia Doudali, IMDEA Software Institute, Spain
Georgios Portokalidis, IMDEA Software Institute, Spain
Christian Pinto, IBM Research Europe, Ireland

Shadow PC Chairs
Antonio Barbalace, University of Edinburgh, Scotland
Baptiste Lepers, Université de Neuchâtel, Switzerland

Doctoral Workshop Chairs
Haitham Al Hassanieh, EPFL, Switzerland
Sanidhya Kashyap, EPFL, Switzerland

Travel Grants Chair
Asia Slowinska, Vrije Universiteit Amsterdam, The Netherlands

Daniele Cono D'Elia, Sapienza University of Rome, Italy

Publicity Chairs
Manolis Stamatogiannakis, Vrije Universiteit Amsterdam, The Netherlands
Vijay Chidambaram, UT Austin, USA

Sponsorship Chairs
Minsoo Rhu, KAIST, South Korea
Rakesh Kumar, Norwegian University of Science and Technology, Norway
Benjamin Lee, University of Pennsylvania, USA

Registration Chair
Enrico Bassetti, Delft University of Technology, The Netherlands

AV Chair
Aleksandar Markovic, Vrije Universiteit Amsterdam, The Netherlands