DAGSTUHL
ARTIFACTS SERIES

Volume 12 | Issue 2 | July 2026

Special Issue of the 38th European Conference on Real-Time Systems (ECRTS 2026)

Edited by

Bryan C. Ward
Federico Aromolo

ISSN 2509-8195

DARTS Special Issue Editors

Bryan C. Ward
Vanderbilt University, Nashville, TN, USA
bryan.ward@vanderbilt.edu
https://orcid.org/0000-0001-7168-6693

Federico Aromolo
Scuola Superiore Sant'Anna, Pisa, Italy
federico.aromolo@santannapisa.it
https://orcid.org/0009-0007-3537-8782

ACM Classification 2012
Computer systems organization → Real-time systems

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.
Online available at
`https://drops.dagstuhl.de/darts`.

Publication date
July 2026

Digital Object Identifier
10.4230/DARTS.12.2.0

Aims and Scope
The Dagstuhl Artifacts Series (DARTS) publishes evaluated research data and artifacts in all areas of computer science. An artifact can be any kind of content related to computer science research, e.g., experimental data, source code, virtual machines containing a complete setup, test suites, or tools.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
DARTS, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
publishing@dagstuhl.de

`https://www.dagstuhl.de/darts`

Contents

Preface
Bryan C. Ward and Federico Aromolo 0:vii

Artifact Evaluation Process
..... 0:ix

Artifact Evaluation Committee
..... 0:xi

Artifacts

PREEMPT-FaaS: Taming Orchestration Times in Latency-Sensitive Serverless Environments (Artifact)
Marcello Cinque, Luigi De Simone, Raffaele Della Corte, and Stefano Toscano ... 1:1–1:21

HyperSSE: Cross-Domain Static Analysis of Partitioned Real-Time Hypervisor Systems (Artifact)
Andreas Kässens, Mareike Burg, and Daniel Lohmann 2:1–2:3

BRUMM: A Case for Predictable Memory Reclamation (Artifact)
Viktor Reusch, Michael Roitzsch, and Horst Schirmeier 3:1–3:3

DySM: Dynamic Scaling of GPU Streaming Multiprocessor in Spatially Shared Real-Time Embedded GPU Systems (Artifact)
Srinivasan Subramaniyan and Xiaorui Wang 4:1–4:3

Controlling Adaptive HARQ Erasure Coding for Real-Time Transport under Channel Model Mismatch (Artifact)
Moritz Miodek, Marlene Böhmer, and Thorsten Herfet 5:1–5:4

Nancy-Playground: A Console Calculator for Deterministic Network Calculus (Artifact)
Raffaele Zippo and Giovanni Stea 6:1–6:5

CacheFlow: Using Maximum Flow to Bound Cache-Based Preemption Delays (Artifact)
Tiancheng He and Bryan C. Ward 7:1–7:5

DART: A Real-Time Address-Randomization Defense with Predictable Timing (Artifact)
Patrick Dobranowski, Owen Rice, Ryan Burrow, Nathan Burow, and Bryan C. Ward 8:1–8:4

Uncertainty-Aware Resource Allocation for Multi-Path Programs with In-Kernel Predictions (Artifact)
Abigail Eisenklam[1], Carlos A. Montenegro G.[1], Xian Wang, Yifan Cai, Robert Gifford, Linh Thi Xuan Phan, and Ricardo G. Sanfelice 9:1–9:4

Boost-At-The-Tail: Work-Triggered Frequency Boosting for Fixed-Priority Scheduling (Artifact)
Behnam Khodabandeloo, Chengzi Huang, and Pontus Ekberg 10:1–10:2

38th European Conference on Real-Time Systems (ECRTS 2026).
Editors: Bryan C. Ward and Federico Aromolo
Dagstuhl Artifacts Series
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Probabilistic Schedulability Analysis for Mixed-Criticality DAG Tasks on Multiprocessors (Artifact)
Hiroto Takahashi, Atsushi Yano, and Takuya Azumi 11:1–11:3

Preface

This volume of the Dagstuhl Artifacts Series (DARTS) collects the artifacts associated with papers accepted to the 38th European Conference on Real-Time Systems (ECRTS 2026), held in Lund, Sweden, July 7–10, 2026. ECRTS was among the first venues in the real-time systems community to establish an artifact evaluation (AE) process, and it has continued to refine that process each year with the goal of promoting the reproducibility, reusability, and long-term availability of the research published at the conference. This year, ECRTS adopted a tiered seal system that distinguishes among artifacts that are made available, that are demonstrated to be functional, that are documented and structured for reuse, and that reproduce the results reported in the paper.

Artifact evaluation gives authors of accepted papers the opportunity to have the software, data, and experimental infrastructure underlying their results independently examined by an Artifact Evaluation Committee, and successful artifacts are recognized with one or more of these seals on the published paper. The accompanying artifacts are archived and published as part of this volume of DARTS so that they remain openly accessible to the community. Participation is entirely voluntary and has no bearing on whether a paper is accepted at ECRTS.

In addition to the tiered seals, the process this year introduced an optional, author-controlled provision allowing reviewers to use generative AI assistance for the limited purpose of resolving technical issues encountered while running an artifact. We observed that the overwhelming majority of authors were amenable to reviewers using AI in their reviewing. The artifact evaluation process is described in detail in the following chapter.

We thank the members of the Artifact Evaluation Committee for their careful and generous work, and we thank the authors who chose to submit their artifacts and to share them with the community.

Bryan C. Ward and Federico Aromolo

38th European Conference on Real-Time Systems (ECRTS 2026).
Editors: Bryan C. Ward and Federico Aromolo
Dagstuhl Artifacts Series
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Artifact Evaluation Process

The ECRTS Artifact Evaluation (AE) process took place separately from paper acceptance. We sought to achieve the benefits of the AE process without disturbing the current process through which ECRTS has generated high-quality programs in the past. Therefore, the submission, review, and acceptance procedure for papers was completely unaltered by the AE process. The decision to submit an artifact and the outcomes of the AE process had no impact on whether a paper was accepted at ECRTS, and the process was entirely optional and up to the authors. Moreover, there is no disclosure of the title and authors of papers whose artifacts did not pass evaluation. This is to avoid negative bias towards submitting artifacts.

The authors of accepted or shepherded papers were invited to submit an artifact evaluation (or replication) package; artifacts associated with shepherded papers were allowed to be submitted as there was a slight overlap in reviewing timelines. Once authors who desired to do so submitted their artifacts, an Artifact Evaluation Committee evaluated the artifacts. The evaluation was single-blind, success-oriented, and non-competitive: the goal of the committee was to help authors demonstrate the value of their artifacts rather than to find reasons to reject them.

Each submitted artifact included two components:

- a document explaining how to use the artifact and which of the experiments presented in the paper are repeatable (with reference to specific results, figures, and tables in the paper), the system requirements, and instructions for installing and using the artifact; and
- the software and any accompanying data.

To minimize dependency and portability problems, authors were encouraged to package their artifacts as Docker containers or as self-contained virtual machine images. The committee also accommodated artifacts with special requirements, such as those depending on particular hardware, on a case-by-case basis.

Based on the evaluation, artifacts could be awarded one or more of the following seals, which authors may display on their published papers:

- *Artifact Available*: the artifact is archived in a publicly accessible repository with appropriate documentation.
- *Artifact Evaluated – Functional*: the evaluators are able to run the artifact and verify that it is functional.
- *Artifact Evaluated – Reusable*: the artifact is functional and, in addition, is carefully documented and well structured so as to facilitate reuse and repurposing by others.
- *Artifact Evaluated – Results Reproduced*: the evaluators are able to reproduce the results reported in the paper using the artifact.

Evaluators bid on artifacts based on their ability to meet the system requirements listed by the authors. Evaluators were assigned to artifacts on this basis to the extent possible, with the goal of three reviewers per artifact. For any artifacts for which one of the co-chairs had a conflict of interest, the other co-chair managed reviewer assignment, communication, and determination of the final decision.

During the first week, all the evaluators checked that they could run the code of the artifacts assigned to them. In case of problems, these were promptly (and anonymously) reported to the authors of the artifact in order to help fix them. Authors could optionally elect to allow reviewers to use generative AI while evaluating their artifacts. When permitted

38th European Conference on Real-Time Systems (ECRTS 2026).
Editors: Bryan C. Ward and Federico Aromolo
Dagstuhl Artifacts Series
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

by the authors, generative AI was allowed only for the purpose of overcoming technical issues encountered in running the artifact, and the human reviewers retained ultimate responsibility for all evaluation decisions. Authors largely supported reviewer AI use, and reviewers largely found that AI helped to resolve simple debugging and platform issues. The total evaluation period was two weeks. After the review period, a brief online discussion took place if and when necessary, and notifications were sent to the authors.

Artifact Evaluation Committee

Artifact Evaluation Co-Chairs

Bryan C. Ward
Vanderbilt University
Nashville, TN, United States of America
bryan.ward@vanderbilt.edu

Federico Aromolo
Scuola Superiore Sant'Anna
Pisa, Italy
federico.aromolo@santannapisa.it

Evaluation Committee

Shareef Ahmed
University of South Florida
Tampa, FL, United States of America

Tanya Amert
Carleton College
Northfield, MN, United States of America

Vijay Banerjee
Washington State University
Pullman, WA, United States of America

Hyunjong Choi
San Diego State University
San Diego, CA, United States of America

Xiaotian Dai
University of York
York, United Kingdom

Tomasz Kłoda
LAAS-CNRS
Toulouse, France

Daniele Ottaviano
Technical University of Munich
Munich, Germany

Binqi Sun
Technical University of Munich
Munich, Germany

Harun Teper
TU Dortmund University
Dortmund, Germany

Raffaele Zippo
University of Pisa
Pisa, Italy

38th European Conference on Real-Time Systems (ECRTS 2026).
Editors: Bryan C. Ward and Federico Aromolo
Dagstuhl Artifacts Series
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany