

Proceedings

**2024 IEEE/ACM 19th Symposium on Software
Engineering for Adaptive and Self-Managing Systems**

SEAMS 2024

**15 – 16 April 2024
Lisbon, Portugal**

Proceedings

**2024 IEEE/ACM 19th Symposium on Software
Engineering for Adaptive and Self-Managing Systems**

SEAMS 2024

**15 – 16 April 2024
Lisbon, Portugal**



Association for
Computing Machinery



IEEE
COMPUTER
SOCIETY

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

ACM COPYRIGHT NOTICE. Copyright © 2024 by the Association for Computing Machinery, Inc. 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, +1-978-750-8400, +1-978-750-4470 (fax).

Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that was previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG Newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.

ACM ISBN: 979-8-4007-0585-4

Editorial production by Lisa O'Conner
Cover art production by Annie Jiu



IEEE Computer Society
Conference Publishing Services (CPS)

<http://www.computer.org/cps>

2024 IEEE/ACM 19th Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS) SEAMS 2024

Table of Contents

Message from the SEAMS 2024 Chairs	ix
Organizing Committee	x
Program Committee	xi

Session 1: Intro + Keynote 1

Keynote 1: Advances on Symbolic Machine Learning and Recent Applications to Software Engineering	1
<i>Alessandra Russo (Imperial College London)</i>	

Session 2: Uncertainty

Formal Synthesis of Uncertainty Reduction Controllers	2
<i>Marc Carwehl (Humboldt-Universität zu Berlin, Germany), Calum Imrie (University of York, UK), Thomas Vogel (Humboldt-Universität zu Berlin, Germany), Genaina Rodrigues (University of Brasilia, Brasil), Radu Calinescu (University of York, UK), and Lars Grunske (Humboldt-Universität zu Berlin, Germany)</i>	
Automated Planning for Adaptive Cyber-Physical Systems under Uncertainty in Temporal Availability Constraints	14
<i>Raquel Sánchez-Salas (ITIS Software, Universidad de Málaga), Javier Troya (ITIS Software, Universidad de Málaga), and Javier Cámara (ITIS Software, Universidad de Málaga)</i>	
Handling uncertainty in the specification of autonomous multi-robot systems through mission adaptation	25
<i>Gianluca Filippone (University of L'Aquila), Juan Antonio Piñera García (Gran Sasso Science Institute), Marco Autili (University of L'Aquila), and Patrizio Pelliccione (Gran Sasso Science Institute)</i>	

Uncertainty Flow Diagrams: Towards a Systematic Representation of Uncertainty Propagation and Interaction in Adaptive Systems	37
<i>Javier Cámara (ITIS Software, Universidad de Málaga), Sebastian Hahner (Karlsruhe Institute of Technology), Diego Perez-Palacin (Linnaeus University), Antonio Vallecillo (ITIS Software, Universidad de Málaga), Maribel Acosta (Technical University of Munich), Nelly Bencomo (Durham University), Radu Calinescu (University of York), and Simos Gerasimou (University of York)</i>	

Session 3: Unmanned Aerial Vehicles and LLMs

ADAM: Adaptive Monitoring of Runtime Anomalies in Small Uncrewed Aerial Systems	44
<i>Md Nafee Al Islam (University of Notre Dame, USA), Jane Cleland-Huang (University of Notre Dame, USA), and Michael Vierhauser (University of Innsbruck, Austria)</i>	
Towards Proactive Decentralized Adaptation of Unmanned Aerial Vehicles for Wildfire Tracking	56
<i>Enrique Vilchez (ITIS Software, Universidad de Málaga), Javier Troya (ITIS Software, Universidad de Málaga), and Javier Cámara (ITIS software, Universidad de Málaga)</i>	
Wildfire-UAVSim: An Exemplar for Evaluation of Adaptive Cyber-Physical Systems in Partially-Observable Environments	63
<i>Enrique Vilchez (ITIS Software, Universidad de Málaga), Javier Troya (ITIS Software, Universidad de Málaga), and Javier Cámara (ITIS Software, Universidad de Málaga)</i>	
Aloft: Self-Adaptive Drone Controller Testbed	70
<i>Calum Imrie (University of York), Rhys Howard (University of Oxford), Divya Thuremella (University of Oxford), Nawshin Mannan Proma (University of York), Tejas Pandey (University of York), Paulina Lewinska (University of York), Ricardo Cannizzaro (University of Oxford), Richard Hawkins (University of York), Colin Paterson (University of York), Lars Kunze (University of Oxford), and Victoria Hodge (University of York)</i>	
Exploring the Potential of Large Language Models in Self-adaptive Systems	77
<i>Jialong Li (Waseda University), Mingyue Zhang (Southwest University), Nianyu Li (ZGC National Laboratory), Danny Weyns (KU Leuven), Zhi Jin (Peking University), and Kenji Tei (Tokyo Institute of Technology)</i>	

Session 4: Testing and Community Debate

Automating Pipelines of A/B Tests with Population Split Using Self-Adaptation and Machine Learning	84
<i>Federico Quin (KU Leuven, Belgium) and Danny Weyns (Linnaeus University, Sweden and KU Leuven, Belgium)</i>	

Generating Executable Test Scenarios from Autonomous Vehicle Disengagements using Natural Language Processing	98
<i>Qunying Song (Lund University, Sweden), Rune Anderberg (Lund University, Sweden), Henrik Olsson (Lund University, Sweden), and Per Runeson (Lund University, Sweden)</i>	
Swarm intelligence-based bio-inspired algorithms	105
<i>Darko Bozhinoski (IRIDIA, Université Libre de Bruxelles)</i>	
Bio-inspired computing systems: handle with care, discard if need it	107
<i>Rogério de Lemos (University of Kent)</i>	

Session 5: Awards + Keynote 2

Keynote 2: Towards Always Law-Abiding Self-Driving	109
<i>Sun Jun (Singapore Management University (SMU))</i>	

Session 6: Self-Recovery & Evaluation

Raft Protocol for Fault Tolerance and Self-Recovery in Federated Learning	110
<i>Rustem Dautov (SINTEF Digital) and Erik Johannes Husom (SINTEF Digital)</i>	
Integrating Graceful Degradation and Recovery through Requirement-driven Adaptation	122
<i>Simon Chu (Carnegie Mellon University), Justin Koe (The Cooper Union), David Garlan (Carnegie Mellon University), and Eunsuk Kang (Carnegie Mellon University)</i>	
Learning Recovery Strategies for Dynamic Self-healing in Reactive Systems	133
<i>Mateo Sanabria (Universidad de los Andes), Ivana Dusparic (Trinity College Dublin), and Nicolás Cardozo (Universidad de los Andes)</i>	
SWITCH: An Exemplar for Evaluating Self-Adaptive ML-Enabled Systems	143
<i>Arya Marda (IIIT Hyderabad, India), Shubham Kulkarni (IIIT Hyderabad, India), and Karthik Vaidhyanathan (IIIT Hyderabad, India)</i>	

Session 7: SAS Applications

Patterns of Applied Control for Public Health Measures on Transportation Services under Epidemic	150
<i>Kenneth Johnson (Auckland University of Technology), Samaneh Madanian (Auckland University of Technology), and Catia Trubiani (Gran Sasso Science Institute)</i>	
RAMSES: An Artifact Exemplar for Engineering Self-Adaptive Microservice Applications	161
<i>Vincenzo Riccio (Politecnico di Milano, Italy), Giancarlo Sorrentino (Politecnico di Milano, Italy), Ettore Zamponi (Politecnico di Milano, Italy), Matteo Camilli (Politecnico di Milano, Italy), Raffaella Mirandola (Karlsruhe Institute of Technology, Germany), and Patrizia Scandurra (University of Bergamo, Italy)</i>	

Self-adaptive, Requirements-driven Autoscaling of Microservices	168
<i>João Paulo Karol Santos Nunes (IBM Brazil and University of São Paulo), Shiva Nejati (University of Ottawa), Mehrdad Sabetzadeh (University of Ottawa), and Elisa Yumi Nakagawa (University of São Paulo)</i>	
GreenhouseDT: An Exemplar for Digital Twins	175
<i>Eduard Kamburjan (University of Oslo, Norway), Riccardo Sieve (University of Oslo, Norway), Chinmayi Prabhu Baramashetru (University of Oslo, Norway), Marco Amato (University of Turin, Italy), Gianluca Barmina (University of Turin, Italy), Eduard Occhipinti (University of Turin, Italy), and Einar Broch Johnsen (University of Oslo)</i>	
Latency-aware RDMSim: Enabling the Investigation of Latency in Self-Adaptation for the Case of Remote Data Mirroring	182
<i>Sebastian Götz (Technische Universität Dresden), Nelly Bencomo (Durham University), and Huma Samin (Durham University)</i>	

Session 8: Human Aspects

Explanation-driven Self-adaptation using Model-agnostic Interpretable Machine Learning	189
<i>Francesco Renato Negri (Politecnico di Milano), Niccolò Nicolosi (Politecnico di Milano), Matteo Camilli (Politecnico di Milano), and Raffaella Mirandola (Karlsruhe Institute of Technology)</i>	
Human empowerment in self-adaptive socio-technical systems	200
<i>Nicolas Boltz (Karlsruhe Institute of Technology, Germany), Sinem Getir Yaman (University of York, United Kingdom), Paola Inverardi (Gran Sasso Science Institute, Italy), Rogério de Lemos (University of Kent, United Kingdom), Dimitri Van Landuyt (KU Leuven, Belgium), and Andrea Zisman (The Open University, United Kingdom)</i>	
Towards Understanding Trust in Self-adaptive Systems	207
<i>Dimitri Van Landuyt (KU Leuven, Belgium), Dávid Halász (Masaryk University, Czech Republic), Stef Verreydt (KU Leuven, Belgium), and Danny Weyns (KU Leuven, Belgium/Linnaeus University, Sweden)</i>	
SafeDriveRL: Combining Non-cooperative Game Theory with Reinforcement Learning to Explore and Mitigate Human-based Uncertainty for Autonomous Vehicles	214
<i>Kenneth H. Chan (Michigan State University), Sol Zilberman (Michigan State University), Nick Polanco (Michigan State University), Joshua E. Siegel (Michigan State University), and Betty H.C. Cheng (Michigan State University)</i>	

Author Index	221
---------------------------	------------

Message from the Chairs

SEAMS 2024

SEAMS aims to showcase and discuss cutting-edge research and advancements in the area of self-adaptive systems. This includes software engineering methods, techniques, processes, and tools to support the construction of safe, performant, and cost-effective self-adaptive and autonomous systems that provide self-* properties like self-configuration, self-healing, self-optimization, and self-protection. The objective of SEAMS is to bring together researchers and practitioners from academia, industry, and government to investigate, discuss, examine, and advance the fundamental principles, state-of-the-art, and solutions addressing critical challenges of engineering self-adaptive and self-managing systems.

2024 is the 19th year of the SEAMS series. New to this year, we used the light double-blind submission system. We maintained two submission rounds for the Research Track, with deadlines in October 2023 and December 2023, with the possibility of submitting a revised version from the first round to the second one. Besides research and short papers, we also solicited community debate papers to discuss the role of bio-inspired algorithms in the engineering of self-adaptive systems.

This year the track received 69 submissions to the research track, of which 11 full papers, 7 short papers and 2 community debates were selected for presentation. We also received 8 submissions to the artifact track, of which 6 artifacts were selected for presentation. We will be able to hear and exchange ideas and views about techniques and artifacts to address uncertainties, test self-adaptive systems, engineer recovery of self-adaptive systems, and tackle human aspects of self-adaptation. Some papers and artifacts considered specific application domains, such as unmanned aerial vehicles, transportation during an epidemic, micro-services, and digital twins. We will give an award to the best paper and artifact and the most influential paper from those presented ten years ago. In addition, authors of selected papers will be invited to submit a revised version of their paper to the special issue of the *ACM Transactions on Autonomous and Adaptive Systems*.

We would like to warmly thank our programme committee members for their thorough and timely reviews of the submissions. We also thank the members of the organizing committee for their tireless support. It is because of all of them that the conference attendees and the readers of the proceedings are able to enjoy a strong program!

We hope that these proceedings will serve as a valuable reference for software engineering researchers and software users.

General Chair

Luciano Baresi, Politecnico di Milano, Italy

Program Co-Chairs

Xiaoxing Ma, Nanjing University, China

Liliana Pasquale, University College Dublin and Lero, Ireland

Organizing Committee

SEAMS 2024

General Chair

Luciano Baresi, *Politecnico di Milano, Italy*

Program Co-Chairs

Liliana Pasquale, *University College Dublin & Lero, Ireland*

Xiaoxing Ma, *Nanjing University, China*

MIP Award Chair

Amel Bennaceur, *The Open University, United Kingdom*

Proceedings Chair

Vivek Nallur, *University College Dublin, Ireland*

Community Debate Chair

Siobhán Clarke, *Trinity College Dublin, Ireland*

Web Chair

Giovanni Quattrocchi, *Politecnico di Milano, Italy*

Publicity Co-Chairs

Erik Fredericks, *Grand Valley State University, United States*

Sona Ghahremani, *Hasso Plattner Institute, University of Potsdam, Germany*

Wenhua Yang, *Nanjing University of Aeronautics and Astronautics, China*

Artifact Evaluation Co-Chairs

Christos Tsigkanos, *University of Athens, Greece*

Shanshan Li, *National University of Defense Technology, China*

Social Media Co-Chairs

Livia Lestingi, *DEIB, Politecnico di Milano, Italy*

Yuan Yao, *Nanjing University, China*

Program Committee

Alessandro V. Papadopoulos, *Mälardalen University, Sweden*
Andrea Zisman, *The Open University, UK*
Antonio, Filieri, *AWS and Imperial College London, UK*
Bradley Schmerl, *Carnegie Mellon University, USA*
Catia Trubiani, *Gran Sasso Science Institute, Italy*
Claudia Raibulet, *University of Milano-Bicocca, Italy*
Dalal Alrajeh, *Imperial College London, UK*
Danny Weyns, *KU Leuven, Belgium and Linnaeus University, Sweden*
Dimitri Van Landuyt, *KU Leuven, Belgium*
Elisa Yumi Nakagawa, *University of São Paulo, Brazil*
Elisabetta Di Nitto, *Politecnico di Milano, Italy*
Enes Yigitbas, *Paderborn University, Germany*
Erick Fredericks, *Grand Valley State University, USA*
Eunsuk Kang, *Carnegie Mellon University, USA*
Fatemeh Golpayegani, *University College Dublin, Ireland*
Gabriel A. Moreno, *SEI and Carnegie Mellon University, USA*
Genaína Nunes Rodrigues, *University of Brasília, Brazil*
Giacomo Cabri, *Università di Modena e Reggio Emilia, Italy*
Giovanni Quattrocchi, *Politecnico di Milano, Italy*
Hausi Muller, *University of Victoria, Canada*
Henry Muccini, *University of L'Aquila, Italy*
Hiroyuki Nakagawa, *Osaka University, Japan*
Ilias Gerostathopoulos, *Vrije Universiteit Amsterdam, Netherlands*
Ivana Dusparic, *Trinity College Dublin, Ireland*
Javier Camara, *University of Malaga, Spain*
Jesper Andersson, *Linnaeus University, Sweden*
Karthik Vaidhyanathan, *IIT Hyderabad, India*
Kenji Tei, *Tokyo Institute of Technology, Japan*
Marco Autili, *University of L'Aquila, Italy*
Marin Litoiu, *York University, Canada*
Michael Vierhauser, *University of Innsbruck, Austria*
Mingyue Zhang, *Southwest University, China*
Nelly Bencomo, *Durham University, UK*
Nicolas Cardozo, *Universidad de los Andes, Colombia*
Patrizia Scandurra, *University of Bergamo, Italy*
Radu Calinescu, *University of York, UK*
Raffaella Mirandola, *Politecnico di Milano, Italy*
Rami Bahsoon, *University of Birmingham, UK*
Ridwan Shariffdeen, *National University of Singapore, Singapore*
Rogerio de Lemos, *University of Kent, UK*
Sebastian Götz, *Technische Universität Dresden, Germany*
Shihong Huang, *Carnegie Mellon University, USA*
Shiva Nejati, *University of Ottawa, Canada*
Simos Gerasimou, *University of York, UK*

Sinem Getir Yaman, *University of York, UK*
Siobhan Clarke, *Trinity College Dublin, Ireland*
Sona Ghahremani, *Hasso Plattner Institute, Germany*
Sophie Cerf, *INRIA, France*
Thomas Vogel, *Humboldt-Universität zu Berlin, Germany*
Thomas Welsh, *University of Iceland, Iceland*
Tianxiao Gu, *Tiktok, China*
Timo Kehrer, *University of Bern, Switzerland*
Wenhua Yang, *Nanjing University of Aeronautics and Astronautics, China*
Xiaoyu Sun, *Australian National University, Australia*
Xin Pen, *Fudan University, China*
Yi Qin, *Nanjing University, China*
Zhaoguo Wang, *Shanghai Jiao Tong University, China*
Zhi Jin, *Peking University, China*

External Reviewers

Huiyan Wang, *Nanjing University, China*
Jingwei Xu, *Nanjing University, China*
Kushal Ramkumar, *University College Dublin and Lero, Ireland*
Liang Wang, *Nanjing University, China*
Yuan Zhang, *Nanjing University, China*

SEAMS 2024 Keynote 1

Alessandra Russo
Imperial College London

Alessandra Russo is a Professor on Applied Computational Logic, at the Department of Computing, Imperial College London, Deputy director of the UKRI Centre for Doctoral Training in “Safe and Trusted AI”, and promoter of the Imperial-X inter-disciplinary research initiative “Intelligible AI” on explainable, safe and trustworthy AI. She leads the “Structured and Probabilistic Intelligent Knowledge Engineering (SPIKE)” research group at the Department of Computing. She has pioneered several state-of-the-art symbolic machine learning systems, including the recent state-of-the-art LAS (Learning from Answer Sets) system for learning interpretable knowledge from labelled data. More recently she has explored novel methodologies for neuro-symbolic learning that integrate machine learning and probabilistic inference with symbolic learning to support generalisation and transfer learning from multimodal unstructured data. She has published over 200 articles in flagship conferences and high impact journals in Artificial Intelligence and Software Engineering, and led various projects funded by the EPSRC, the EU and Industry.

Keynote: Advances on Symbolic Machine Learning and Recent Applications to Software Engineering

Learning interpretable models from data is one of the main challenges of AI. Symbolic Machine Learning, a field of Machine Learning, offers algorithms and systems for learning models that explain data in the context of a given domain knowledge. In contrast to statistical learning, models learned by Symbolic Machine Learning are interpretable: they can be translated into natural language and understood by humans. In this talk, I will overview our state-of-the-art symbolic machine learning system (ILASP) capable of learning different classes of models, (e.g., non-monotonic, non-deterministic and preference-based) for real-world problems, in a manner that is data efficient, scalable, and robust to noise. I will show how such system can be integrated with statistical and deep learning to provide neuro-symbolic AI solutions for learning complex interpretable knowledge from unstructured data. I will then illustrate how these advances can be applied to areas such agent learning, run-time adaptation of security for unmanned aerial vehicles, and online learning of policies for explainable security.

SEAMS 2024

Keynote 2

Sun Jun

Singapore Management University

Sun Jun is currently a professor at Singapore Management University (SMU). He received Bachelor and PhD degrees in computing science from National University of Singapore (NUS) in 2002 and 2006. He has been a faculty member since 2010. He was a visiting scholar at MIT from 2011-2012. Jun's research interests include AI safety, formal methods, program analysis and cyber-security. He is the co-founder of the PAT model checker. He has published many journal articles or peer-reviewed conference papers, many of which are published at top-tier venues. He serves as the technical consultant for multiple companies.

Keynote: Towards Always Law-Abiding Self-Driving

How should an autonomous vehicle behave on the road besides causing no accidents and reaching the destination? Fortunately, rich sets of criteria for how a vehicle should undertake a journey already exist: the various national traffic laws. In addition to avoiding collisions, an autonomous vehicle should satisfy the traffic laws of the country it operates in. Until we design new traffic laws specifically for autonomous vehicles, existing traffic laws remain the gold standard for ensuring road safety. The question is then: how do we systematically make sure that an autonomous vehicle almost always abides the traffic laws? In this work, I will introduce our recent effort on formalizing national traffic laws and use it to adaptively enforce desirable self-driving automatically.

Author Index

Acosta, Maribel	37	Jin, Zhi	77
Amato, Marco	175	Johnsen, Einar Broch	175
Anderberg, Rune	98	Johnson, Kenneth	150
Autili, Marco	25	Jun, Sun	109
Baramashetru, Chinmayi Prabhu	175	Kamburjan, Eduard	175
Barmina, Gianluca	175	Kang, Eunsuk	122
Bencomo, Nelly	37, 182	Karol Santos Nunes, João Paulo	168
Boltz, Nicolas	200	Koe, Justin	122
Bozhinoski, Darko	105	Kulkarni, Shubham	143
Calinescu, Radu	2, 37	Kunze, Lars	70
Cámara, Javier	14, 37, 56, 63	Lewinska, Paulina	70
Camilli, Matteo	161, 189	Li, Jialong	77
Cannizzaro, Ricardo	70	Li, Nianyu	77
Cardozo, Nicolás	133	Madanian, Samaneh	150
Carwehl, Marc	2	Marda, Arya	143
Chan, Kenneth H.	214	Mirandola, Raffaella	161, 189
Cheng, Betty H.C.	214	Negri, Francesco Renato	189
Chu, Simon	122	Nejati, Shiva	168
Cleland-Huang, Jane	44	Nicolosi, Niccolò	189
Dautov, Rustem	110	Occhipinti, Eduard	175
de Lemos, Rogério	107, 200	Olsson, Henrik	98
Dusparic, Ivana	133	Pandey, Tejas	70
Filippone, Gianluca	25	Paterson, Colin	70
Garlan, David	122	Pelliccione, Patrizio	25
Gerasimou, Simos	37	Perez-Palacin, Diego	37
Getir Yaman, Sinem	200	Piñera García, Juan Antonio	25
Götz, Sebastian	182	Polanco, Nick	214
Grunske, Lars	2	Proma, Nawshin Mannan	70
Hahner, Sebastian	37	Quin, Federico	84
Halász, Dávid	207	Riccio, Vincenzo	161
Hawkins, Richard	70	Rodrigues, Genaina	2
Hodge, Victoria	70	Runeson, Per	98
Howard, Rhys	70	Russo, Alessandra	1
Husom, Erik Johannes	110	Sabetzadeh, Mehrdad	168
Imrie, Calum	2, 70	Samin, Huma	182
Inverardi, Paola	200	Sanabria, Mateo	133
Islam, Md Nafee Al	44	Sánchez-Salas, Raquel	14

Scandurra, Patrizia	161
Siegel, Joshua E.	214
Sieve, Riccardo	175
Song, Qunying	98
Sorrentino, Giancarlo	161
Tei, Kenji	77
Thuremella, Divya	70
Troya, Javier	14, 56, 63
Trubiani, Catia	150
Vaidhyanathan, Karthik	143
Vallecillo, Antonio	37
Van Landuyt, Dimitri	200, 207
Verreydt, Stef	207
Vierhauser, Michael	44
Vilchez, Enrique	56, 63
Vogel, Thomas	2
Weyns, Danny	77, 84, 207
Yumi Nakagawa, Elisa	168
Zamponi, Ettore	161
Zhang, Mingyue	77
Zilberman, Sol	214
Zisman, Andrea	200



RAMSES: an Artifact Exemplar for Engineering Self-Adaptive Microservice Applications

Vincenzo Riccio
Giancarlo Sorrentino
Ettore Zamponi
Matteo Camilli
{first.last}@mail.polimi.it
matteo.camilli@polimi.it
Politecnico di Milano
Milano, Italy

Raffaella Mirandola
raffaella.mirandola@kit.edu
Karlsruhe Institute of Technology
Karlsruhe, Germany

Patrizia Scandurra
patrizia.scandurra@unibg.it
University of Bergamo
Bergamo, Italy

ABSTRACT

This paper introduces RAMSES, an exemplar tailored for both practitioners and researchers working on self-adaptive microservice applications. By emphasizing a clear separation of concerns between the application and its adaptation logic, RAMSES realizes a reusable autonomic manager that implements a MAPE-K feedback loop whose components are microservices themselves. Its primary focus lies in addressing user-defined QoS attributes at runtime, like availability and performance. To illustrate its usage, we provide a practical example showing its mechanics in an e-food microservice application. Initial experiments indicate the advantages of utilizing RAMSES, as shown by a comparative analysis of the quality properties of a microservice application with and without self-adaptation.

CCS CONCEPTS

• **Computer systems organization** → *Self-organizing autonomic computing; Distributed architectures*; • **Software and its engineering** → *Software verification and validation; Extra-functional properties*.

KEYWORDS

Microservice applications, self-adaptation, MAPE-K, exemplar

ACM Reference Format:

Vincenzo Riccio, Giancarlo Sorrentino, Ettore Zamponi, Matteo Camilli, Raffaella Mirandola, and Patrizia Scandurra. 2024. RAMSES: an Artifact Exemplar for Engineering Self-Adaptive Microservice Applications. In *19th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '24)*, April 15–16, 2024, Lisbon, AA, Portugal. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3643915.3644110>

1 INTRODUCTION

Microservice architectures have gained widespread popularity due to their scalability, flexibility, and ability to facilitate continuous deployment. Existing development frameworks (like Spring, Flask and GoKit, to name a few) and management infrastructures simplify

the whole development and management process also offering some forms of self-adaptation built-in mechanisms (such autoscaling and circuit breaking for resilience)[6]. However, the dynamic and heterogeneous nature of modern computing environments calls for the capacity to employ self-adaptation strategies that extend beyond these built-in capabilities and can achieve arbitrary adaptation.

As a step in addressing this challenge, this paper proposes the exemplar RAMSES to facilitate practitioners and researchers in engineering a self-adaptive microservice application by clearly separating the microservice application and self-adaptation concerns. RAMSES provides a reusable autonomic manager (a managing subsystem) conforming to the well-known feedback control loop model MAPE-K [4] (Monitor-Analyse-Plan-Execute over a Knowledge base) to make a microservice application self-adaptive. Adaptation concerns are mainly aimed at satisfying user-defined QoS attributes (e.g., availability and response time) and self-* autonomic properties (self-configuring, self-healing, and self-optimizing) of a microservice application at runtime. RAMSES's control loop components themselves are microservices. RAMSES is designed to ease its reuse across microservice applications through an *API-led integration* that exploits a *Contract-First* approach to specify probing/actuating RESTful APIs for connecting the RAMSES autonomic manager to any managed microservice application. RAMSES was implemented using the Java-based SPRING BOOT and SPRING CLOUD frameworks¹.

To illustrate RAMSES, as part of this exemplar, an e-food microservice application has been developed and released to be used as target managed system. This paper concretely reports some preliminary evaluation results about the capability of RAMSES to make such a microservice application self-adaptive by enforcing the adaptation goals at runtime.

Concerning the timeliness of the problem addressed by the proposed artifact and the overall discussion and considerations in [7] regarding generality ("*ability to support a variety of architectural models and adaptation mechanisms*") and reusability ("*ability to be reused without requiring substantial effort from software developer*") in self-adaptive microservices, RAMSES has been conceived with a focus on decoupling managed and managing systems. This is achieved by defining RESTful probing and actuating APIs, providing a means for managing systems to interact with managed systems that expose the same interfaces, thereby facilitating reusability and reuse by-design. At the same time, to better handle and speed up the



This work licensed under Creative Commons Attribution International 4.0 License.

SEAMS '24, April 15–16, 2024, Lisbon, AA, Portugal
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0585-4/24/04.
<https://doi.org/10.1145/3643915.3644110>

¹<https://spring.io/microservices>

implementation of various self-adaptation strategies, the managing system in RAMSES has been developed leveraging microservices principles and mainstream microservice frameworks and infrastructure platforms. In-house mechanisms have also been implemented to enhance its autonomy from the management infrastructure itself. This situates RAMSES at the *Cross-Layer level* in the MAPE integration patterns as defined in [7].

The main contributions of this artifact paper for the community of *Software Engineering for Adaptive and Self-Managing Systems* are: (i) a non-simulated exemplar of self-adaptive microservice application; (ii) a standalone autonomic manager for microservices characterised by a Contract-First and API-led approaches for reusability across other microservice applications; (iii) an in-house developed managed microservice application that can be also reused in a standalone manner as a use case for other benchmarks. Section 5 concludes the paper.

2 RELATED WORK

Exemplars development is an active line of research in the self-adaptive community². We report here the exemplars most related to our work. Hognia [1] is a platform for deploying self-managing web applications on the cloud and automating a set of operations, such as the booting of the instances and their setup. Moreover, it enables the continuous monitoring of the health status of the applications. It provides a modular managing system, which is however specifically focused on the deployment and configuration of cloud applications on platforms such as *Amazon EC2*. Another well-known exemplar is *TAS* [13]. It provides a useful (simulated) service-based system to be used as a managed subsystem. Moreover, it defines a set of generic adaptation scenarios applicable to service-based systems, to deal with uncertainties of the system itself and of the environment it is set in, which inspired the one proposed in Table 1. Another related exemplar is *SEABYTE* which proposes an experimental framework for testing novel self-adaptation solutions to enhance the automation of continuous A/B testing of a micro-service-based system [10]. *SEABYTE* represents a ready-to-use framework, implemented using a well-known technology stack (microservices, Spring, Docker, REST/JSON). However, its application is limited to the specific domain of A/B testing.

Other exemplars (e.g., *SWIM* [9], *RDMSim* [12] and *EWS* [2]) implement specific managed subsystems that can be reused by researchers to evaluate and compare different adaptation logic). In different contexts, [3] proposes a framework that implements automatic container sizing and self-healing features for a microservice-based application deployed in Docker containers by exploiting MAPE-K loops. In [14], an extension of Kubernetes has been developed to monitor microservices data and manage aspects of scalability. When compared to RAMSES, the key differences are non-simulated managed and managing systems, the adoption of approaches of *Contract-First* and *API-led* for reusability across multiple platforms. This last aspect is also facilitated by our design choice of not depending to much on the self-adaptation capabilities supported by some service management infrastructures natively (such as Kubernetes autoscaling), to avoid conflicts/interferences with the adaptation decisions of our autonomic manager.

²www.hpi.uni-potsdam.de/giese/public/selfadapt/category/exemplar/




Figure 1: RAMSES architecture overview

3 EXEMPLAR DESCRIPTION

RAMSES is about an autonomic manager (or managing system) for making microservice applications self-adaptive. Once connected to a target microservice application (the managed subsystem) through probing/actuating APIs, the overall self-adaptive microservice application (see Figure 1) is a two-layer architecture that, according to the principles of architecture-based self-adaptation [5], decouples the managed microservice application from its control loop layer. RAMSES control loop components are microservices themselves that execute the monitoring and adaptation activities conforming to the well-known MAPE-K feedback loop model [4]. The goal of RAMSES is to enforce the satisfaction of user-defined QoS attributes (e.g., availability and response time) and self-* properties at run-time, without human oversight. RAMSES is designed to ease its reuse across microservice applications thanks to an *API-led integration* with the managed microservices. An additional dashboard for configuring admin preferences and inspecting the managed microservices at run-time is also supported. RAMSES has been implemented using the Java-based frameworks *SPRING BOOT* and *SPRING CLOUD*³ for developing microservice applications. The exemplar is publicly available (See Section 6), including all software artifacts for injecting and simulating fictional service failures or slowdowns for test scenarios.

3.1 Adaptation concerns and strategies

RAMSES is primarily tailored to automate typical service maintenance tasks. The main adaptation concerns are summarized in the Table 1, together with the adaptation goals that the system adaptations should achieve, and related strategies. These last could be the result of a combined application of single adaptation actions. The current prototype supports the following adaptation actions: *addInstance*, *removeInstance*, *changeConfiguration* (e.g., for tuning the service weights of the load balancing, or changing the circuit breaker's parameters and timeout thresholds), and *selectNewImplementation* (to replace all instances of a given service type with the spawned instances of another implementation of the same service).

³<https://spring.io/microservices>

Table 1: Adaptation concerns, goals, and strategies in RAMSES

Concern	Goal	Strategies (examples)
C1: Individual service failure	Recover from a silent or crashed service (self-healing)	Start a new service instance
C2: Search for better service implementations	Select a better service implementation (self-optimizing)	Change the current service implementation
C3: Workload distribution upon a service instance activation/deactivation	Configure load balancing (self-configuring)	Change service configuration for load balancing weights
C4: Violation of QoS requirements	Maintain QoSs (availability and response time)	Add service instance(s) and/or change weights for load balancing (e.g., to lighten or even shutdown an instance with low performance)

3.2 Probing and actuating APIs

The probing/actuating interfaces have been defined using a *Contract-First approach* to promote reusability and then implemented as a RESTful API. Such APIs are consumed by the managing layer and provided by the managed microservice application.

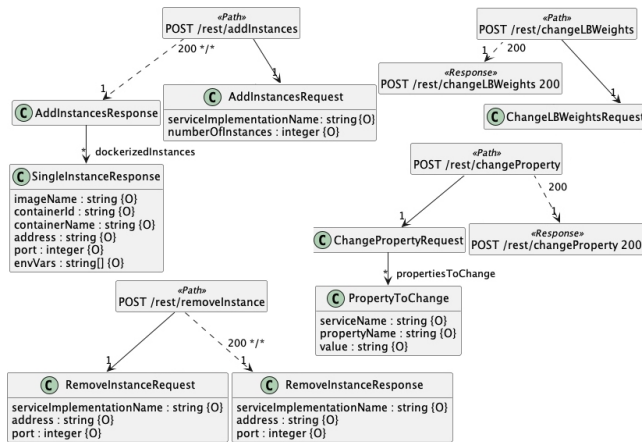
**Figure 2: Actuating REST/JSON API**

Figure 2 shows the actuating API of RAMSES using the PLANTUML class diagram as generated from the API description in the standard format OpenAPI. The actuating API exposes low-level operations to be applied to the managed microservices. Below we provide a description of the main operations exposed by the APIs.

- `POST /rest/addInstances`: allocates and starts a number of new instances of a given service implementation, and returns the details on each new instance (such as its address and port).
- `POST /rest/removeInstance`: stops and deallocates an instance which is currently running.
- `POST /rest/changeLBWeights`: updates the load balancer weights of the instances of a given service.
- `POST /rest/changeProperty`: updates the value of a global property or of a property for a certain service (includes the addition/removal of the property)

The probing API for RAMSES has been defined as well. They are responsible for collecting metrics from all the service instances as a collection of objects *Instance Metrics Snapshot* (or simply *snapshot*). Each service snapshot includes a timestamp, the status of the instance, metrics related to resource usage (e.g., CPU usage), metrics related to HTTP requests (e.g., number of server errors), and circuit breakers. In the current prototype of RAMSES only

availability and average response time are concretely used to extract QoS indicators for adaptations.

3.3 Managing subsystem

This section details the behavior of the MAPE-K microservices of RAMSES. The *Knowledge* microservice maintains an up-to-date data structure with all relevant information about the managed microservices as collected by *Monitor* component through the probing API. These include data about service implementations and their associated operational profile, current values for service configuration parameters (e.g., for load balancing and circuit breaking), running service instances and their snapshots series with the collected metrics (CPU usage, HTTP requests stats, etc.) useful for calculating QoS indicators. Other useful information computed by the MAPE components for coordinating the loop execution itself is also stored in the knowledge.

Through the probing API, the *Monitor* microservice periodically collects snapshots of all running service instances and their associated quality metrics, and stores them in the knowledge. The *Monitor* executes asynchronously w.r.t. the rest of the control loop and its sampling period can be dynamically configured via a REST API exposed by the *Monitor* microservice itself or the admin dashboard. Once a control loop execution terminates, the *Executor* microservice notifies the *Monitor* so that the *Monitor* can invoke and request the *Analyze* microservice to start a new asynchronous adaptation loop with an up-to-date data collection for service snapshots.

The *Analyze* microservice exploits the updated metrics stored in the knowledge for the QoS values computations (e.g., availability and response time) adopting a *sliding window* approach, with configurable window dimensions. Data are collected in the last window of observation and a weighted average of the value of interest is then computed. The weights are values in $[0, 1]$ dynamically assigned to microservice instances by a load balancer. If the evaluated QoS is below the user-defined reference threshold, then a set of adaptation actions with different priority levels are selected and stored in the knowledge and the *Plan* microservice is triggered.

The *Plan* microservice compares the proposed adaptation actions and selects the one that best fits the user requirements. The benefit is calculated depending on the adaptation strategy and by estimating the QoS values obtained after applying the corresponding adaptation actions. For example, if the availability of a service does not satisfy the requirement, an *addInstance* adaptation action can be undertaken followed by a *changeConfiguration* action with the computation of the new weighted average availability of the current instances using the new weights. In each case, the best adaptation strategy is selected using, for example, the option with the associated highest estimated utility or the one that maximizes




Figure 3: Managed e-food microservice application

multiple objectives. The selected strategy becomes the adaptation plan in the knowledge and then the *Execute* microservice is invoked.

The *Execute* microservice retrieves the adaptation plan from the knowledge, if any. Then, for each adaptation action of the plan, it invokes the actuating API as exposed by the managed microservice application. All processed adaptations are persisted into the knowledge. Once the *Execute* microservice terminates, it notifies the completion of the current loop iteration to the *Monitor*.

3.4 Managed microservice application

Our use case is an e-food microservice application (see Figure 3) for ordering food from an online restaurant. It is made of a small set of microservices running in Docker containers, an API gateway that acts as a single entry point, infrastructure means for service management – Netflix Eureka for service discovery, and a configuration server for storing and serving hot configurations to microservices –, and service *circuit breakers* as supported by SPRING CLOUD in a native way. A load balancer (not shown in Figure 3) is also used to allocate requests to service instances according to a *roulette wheel selection* policy [8]. Such a module has been developed in-house to better oversight its internal behavior in a white box, thus avoiding external and not so transparent load balancing mechanisms as offered by management infrastructure solutions (e.g., Kubernetes) of containerized applications. To be fully working, some services (e.g., Restaurant service and Ordering Service) require some data to be persisted. To this end, a *database-per-service* pattern is applied [11], and independent MySQL databases are added. Finally, to make the managed microservice application also usable and inspectable by humans at the edge of the microservices, a reusable web application (also developed as a Spring Boot application) is provided. It serves as the front end of the managed microservice application and acts as a REST client that communicates with the API gateway.

3.5 HMI-dashboard

Once the managed application has been connected (via appropriate configuration property files) to the end-points of the probing/actuating APIs of RAMSES, a web-based HMI dashboard (developed using Spring Boot and the frontend engine THYMELEAF) allows an administrator for configuring (see the screenshot in Figure 4) and inspecting the control loop and the managed microservices at runtime (see the screenshot in Figure 5).

Figure 4a shows the control variables settable from the dashboard for the running managed microservices, including current

Service	Implementation Id	Preference	Trust	Penalty
RESTAURANT-SERVICE	restaurant-service	1.0	1	0
DELIVERY-PROXY-SERVICE	delivery-proxy-1-service	0.4	4	0

Configuration	Time Of Snapshot	Load Balancer Type
RESTAURANT-SERVICE	2022-11-21 16:23:29 UTC	WEIGHTED RANDOM
DELIVERY-PROXY-SERVICE	2022-11-21 16:23:36 UTC	WEIGHTED RANDOM

QoS	Value	Condition	Weight
Availability	92.00%	value > 90.00%	0.5
Average Response Time [ms]	29.4	value < 90.0	0.5

(a) Home

Managing System Dashboard

Adaptation Status: Disabled Start

Monitor Module configuration

Current Status: Not Running Start

Monitor scheduling period [ms]: 5000 Change

Attention! A quick failure detection capability of the managing system highly relies on a frequent monitor routine.

Analyse Module configuration

Metrics Window Size: 4 Change

Analysis Window Size: 5 Change

Failure Rate Threshold: 0.1 Change

Unreachable Rate Threshold: 0.35 Change

Parameters Satisfaction Rate: 0.6 Change

(b) Parameter Configuration

Figure 4: HMI dashboard

implementations, active service instances and their quality values and QoS requirements. Configuration parameters of the control loop (see Figure 4b) are also visible and settable: the current state of the loop (running or not running), the monitor sampling period, and specific parameters for the analysis. These latter include the *metric window size*, i.e. the number of metric collected and buffered before triggering the analysis; the *analysis window size*, i.e. the number of metric values used to estimate the QoS level of each microservice instance; and the *shutdown threshold*, i.e. the minimum amount of requests that each microservice instance must be able to process to stay alive. The dashboard also allows us to control the execution of the MAPE-K loop, for example, to stop the autonomic manager and simply monitor the managed microservice application.

4 EXEMPLAR EVALUATION

This section describes preliminary results about the efficacy of the RAMSES autonomic manager. We refer the reader to our publicly available package for a comprehensive guideline on utilizing the artifact and conducting independent experimental campaigns with RAMSES (see Section 6).

4.1 Setup of the testbed infrastructure

The experiments have been executed on two physical machines: a *driver node* N_1 and a *subject node* N_2 equipped with an Apple M1



Figure 5: Screenshot of the dashboard showing QoS data of the managed microservices at runtime.

(8-core) processor, 16GB RAM LPDDR4, and 256GB NVMe SSD. To run our experiments, we used the *swarm mode* of Docker to manage a cluster of two Docker daemons: DockerManager deployed onto N_1 , and a DockerWorker deployed onto N_2 . The RAMSES managing subsystem and a software module ScenarioRunner for conducting experiments have been deployed into the driver node N_1 . The DockerManager and DockerWorker are in charge of the deployment/undeployment of the containerized microservices of the managed subsystem onto the subject node N_2 .

For the sake of simplicity, the artifact comes with instructions to set up and run RAMSES as well as a set of pre-configured scenarios in a single machine. To promote portability and usability, our package includes also instructions for other mainstream platforms other than MacOS (i.e., Windows, and Linux).

The ScenarioRunner software module runs as a stand-alone microservice and has been developed to facilitate and automate the setup of the experimental campaign. It takes as input the specification of the experiments in terms of all configurable factors required to create and reproduce a certain test scenario: the number of instances per microservice and boot time, the observation period T of the experimental campaign, workload (service requests per second to issue to the managed subsystem), synthetic delays in given time intervals (following a given Normal distribution), failures (synthetic exceptions to generate in a given time interval according to a certain failure rate), and network issues (i.e., time intervals of service unreachability) to inject. The specification is provided to ScenarioRunner through a configuration file where the user specifies the aforementioned factors through pre-defined variables as follows:

```
FAILURE_INJECTION = Y
FAILURE_INJECTION_1_START = 180
ID_OF_INSTANCE_TO_FAIL = restaurant:58085
```

In this example, the option failure injection is set to “Yes”. Thus, a failure is injected at time 180s targeting the microservice instance with ID `restaurant:58085`.

Given the specification, ScenarioRunner runs all the experiments following a simple pipeline for each one of them. Such a pipeline is made of three subsequent phases: *setup*, *execution & monitoring*, and *teardown*. During the *setup* phase the RAMSES autonomic manager and the managed microservice application are deployed following the given configuration. The database images are then loaded onto the managed microservices, and the control loop’s knowledge is reset to its initial state to make sure that the adaptation decisions undertaken are not influenced by prior knowledge collected from past experiments. During the *execution* phase, the ScenarioRunner replicates the desired operating conditions in N_2 by generating a workload for the managed microservices and by injecting specific issues/failures, as specified in the experimental campaign. After a ramp-up period, the *monitoring* sub-phase starts and the ScenarioRunner computes the metrics of interest to quantify the effectiveness of RAMSES in achieving the adaptation goals. When the observation period T ends, the *teardown* phase terminates and undeploys the managing and the managed subsystems.

In our preliminary experiments, the effects of adaptations applied by the autonomic manager have been measured over the observation period of $T = 20$ minutes for all experiments. Concerning the other configuration parameters of the managing layer, we adopted the following default values (see Figure 4b): metric and analysis window size equal to 5, shutdown threshold 40%, and monitor scheduling period of 5 seconds.

4.2 Preliminary results

We here focus on the concern C4 of Table 1 about achieving QoS requirements, and report some outcomes of the comparative analysis of the capability of the system to maintain QoS properties with and without adaptation (see Figure 6 and Figure 7).

To evaluate the system degradation, we introduce the quantitative metric termed the *QoS Degradation Area* (QoSDA). The QoSDA is calculated by measuring the area between the QoS threshold and

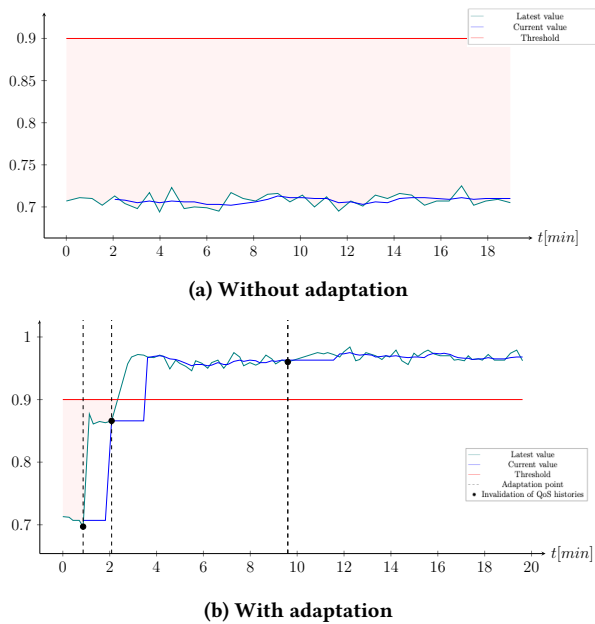


Figure 6: Restaurant service availability

the actual QoS during specific time intervals within the observation period T . These intervals correspond to instances where the operational state is degraded, indicating that the actual QoS falls below the prescribed QoS requirement, such as a target availability level. In essence, the adaptation actions performed by RAMSES aim to minimize the QoSDA. Smaller QoSDA indicates a more effective mitigation of degradation.

Figure 6 shows the results of the Restaurant service availability by running the e-food application without and with the RAMSES autonomous manager under a uniform workload intensity of 100 requests per second generated by 50 concurrent users. The application itself is not able to meet the QoS requirement $availability > 0.9$ as illustrated in Figure 6a, with a QoSDA equal to 3.65×10^3 (red area). With the introduction of the autonomous manager, RAMSES plans and then actuates two adaptations to satisfy the QoS requirement (adaptation points corresponding to $t = 2$ and $t = 3$). The first action corresponds to the change of the weights of the load balancer to penalize the bad-performing instances. This improves the availability up to ~ 0.8 at time $t = 3$, but the requirement is still not satisfied. Another adaptation action is then performed: two instances out of three are shut down, due to their poor performances. The microservice finally achieves a stable and desired QoS value from $t = 4$ on. The QoSDA is 4.28×10^2 (red area) in Figure 6b. The adaptation options reduce the QoSDA by 88%.

Similar results can be observed in Figure 7 considering the Ordering service response time. In this case, the QoS requirement is $response\ time < 800\ milliseconds$, which is not fulfilled without introducing the managing layer for adaptation. Figure 7a shows a QoSDA of 8.22×10^3 . The results of the introduction of the autonomous manager are shown in Figure 7b, where the addition of a new instance and the change of the weights of the load balancer

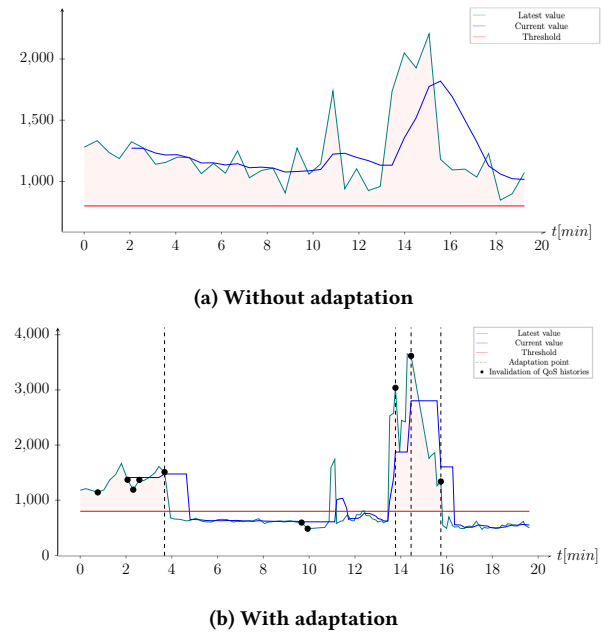


Figure 7: Ordering service response time

yield QoS requirement satisfaction. QoSDA in this last plot is equal to 6.39×10^3 which leads to a 22% reduction.

5 CONCLUSION AND FUTURE DIRECTIONS

This work describes RAMSES a microservice-based autonomous manager tailored to microservice applications. RAMSES is reusable given that the managed system implements the probing/actuating APIs and the provided load-balancing mechanism. The e-food microservice application developed as the managed application is itself reusable by other applications to experiment with different kinds of adaptation means.

In future work, RAMSES may be extended with: more metrics (e.g., resource usage and circuit breakers metrics) for a more in-depth evaluation; more adaptation actions and scenarios to deal with additional quality like security and with more complex situations; more complex decision-making approaches taking into account also the costs and the risks derived from the application of an adaptation option.

6 ARTIFACT AVAILABILITY

The package of RAMSES, including the sources and the instructions to set up and run it is publicly available at <https://zenodo.org/doi/10.5281/zenodo.10400820>. We refer the reader to the README file included in the package for further details.

ACKNOWLEDGMENTS

This work has been partially funded from the topic Engineering Secure Systems of the Helmholtz Association (HGF) and by KASTEL Security Research Labs.

REFERENCES

- [1] C. Barna, H. Ghanbari, M. Litoiu, and M. Shtern. 2015. Hognu: A Platform for Self-Adaptive Applications in Cloud Environments. In *SEAMS*. 83–87.
- [2] R. R. Filho, E. Alberts, I. Gerostathopoulos, B. Porter, and F. M. Costa. 2022. Emergent Web Server: An Exemplar to Explore Online Learning in Compositional Self-Adaptive Systems. In *SEAMS*. 36–42.
- [3] Luca Florio and Elisabetta Di Nitto. 2016. Gru: An Approach to Introduce Decentralized Autonomic Behavior in Microservices Architectures. In *IEEE ICAC*. 357–362.
- [4] Jeffrey O. Kephart and David M. Chess. 2003. The Vision of Autonomic Computing. *IEEE Computer* 36, 1 (Jan. 2003).
- [5] Sara Mahdavi-Hezavehi, Vinicius H.S. Durelli, Danny Weyns, and Paris Avgeriou. 2017. A systematic literature review on methods that handle multiple quality attributes in architecture-based self-adaptive systems. *Information and Software Technology* 90 (2017), 1–26. <https://doi.org/10.1016/j.infsof.2017.03.013>
- [6] N. C. Mendonca, P. Jamshidi, D. Garlan, and C. Pahl. 2021. Developing Self-Adaptive Microservice Systems: Challenges and Directions. *IEEE Software* 38, 2 (2021), 70–79.
- [7] Nabor C. Mendonça, David Garlan, Bradley Schmerl, and Javier Cámara. 2018. Generality vs. Reusability in Architecture-Based Self-Adaptation: The Case for Self-Adaptive Microservices. In *Proceedings of the 12th European Conference on Software Architecture: Companion Proceedings (Madrid, Spain) (ECSA '18)*. Association for Computing Machinery, New York, NY, USA, Article 18, 6 pages.
- [8] M. Mitchell. 1999. *An Introduction to Genetic Algorithms*. The MIT Press, 124–125.
- [9] G. A. Moreno, B. Schmerl, and D. Garlan. 2018. SWIM: An Exemplar for Evaluation and Comparison of Self-Adaptation Approaches for Web Applications (*SEAMS '18*). ACM, 137–143.
- [10] F. Quin and D. Weyns. 2022. SEABYTE: A Self-adaptive Micro-service System Artifact for Automating A/B Testing. In *SEAMS*. 77–83.
- [11] C. Richardson. 2022. Pattern: Microservice Architecture. <https://microservices.io/patterns/microservices.html>
- [12] H. Samin, L. H. G. Paucar, N. Bencomo, C. M. C. Hurtado, and E. M. Fredericks. 2021. RDMSim: An Exemplar for Evaluation and Comparison of Decision-Making Techniques for Self-Adaptation. In *SEAMS*. 238–244.
- [13] D. Weyns and R. Calinescu. 2015. Tele Assistance: A Self-Adaptive Service-Based System Exemplar. In *SEAMS*. 88–92.
- [14] Shuai Zhang, Mingjiang Zhang, Lin Ni, and Peini Liu. 2019. A Multi-Level Self-Adaptation Approach For Microservice Systems. In *2019 IEEE 4th International Conference on Cloud Computing and Big Data Analysis (ICCCBDA)*. 498–502.