

Editorial

Introduction to the Special Issue “Applications in Self-Aware Computing Systems and their Evaluation”

Christian Krupitzer ^{1,*}, Benedikt Eberhardinger ², Ilias Gerostathopoulos ³ and Claudia Raibulet ⁴

¹ Software Engineering Group, Department of Computer Science, University of Würzburg, 97070 Würzburg, Germany

² MHP—A Porsche Company, 71638 Ludwigsburg, Germany; benedikt.eberhardinger@mhp.com

³ Department of Computer Science, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, The Netherlands; i.g.gerostathopoulos@vu.nl

⁴ Dipartimento di Informatica, Sistemistica e Comunicazione, Università degli Studi di Milano-Bicocca, DISCo, 20126 Milano, Italy; raibulet@disco.unimib.it

* Correspondence: christian.krupitzer@uni-wuerzburg.de; Tel.: +49-931-31-83252

Received: 18 March 2020; Accepted: 18 March 2020; Published: 21 March 2020

Abstract: The joint 1st Workshop on Evaluations and Measurements in Self-Aware Computing Systems (EMSAC 2019) and Workshop on Self-Aware Computing (SeAC) was held as part of the FAS* conference alliance in conjunction with the 16th IEEE International Conference on Autonomic Computing (ICAC) and the 13th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO) in Umeå, Sweden on 20 June 2019. The goal of this one-day workshop was to bring together researchers and practitioners from academic environments and from the industry to share their solutions, ideas, visions, and doubts in self-aware computing systems in general and in the evaluation and measurements of such systems in particular. The workshop aimed to enable discussions, partnerships, and collaborations among the participants. This special issue follows the theme of the workshop. It contains extended versions of workshop presentations as well as additional contributions.

Keywords: self-aware computing systems; quality evaluation; measurements; quality assurance; autonomous; self-adaptive; self-managing systems.

1. Introduction

The joint workshops Evaluations and Measurements in Self-Aware Computing Systems (EMSAC <http://emsac2019.informatik.uni-wuerzburg.de/>) and the Workshop on Self-Aware Computing (SeAC <http://seac2019.informatik.uni-wuerzburg.de/>) were held as part of the FAS* conference alliance in conjunction with the 16th IEEE International Conference on Autonomic Computing (ICAC) and the 13th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO) in Umeå, Sweden at June 20, 2019. This was the first edition of the EMSAC workshop and the third of the SeAC workshop.

In the past decade, many different research communities have explored the aspects of self-awareness in computing systems, each from their own perspective. Relevant work can be found in different areas including autonomic computing [1], self-adaptive and self-organizing software and systems [2,3], machine learning, artificial intelligence and multi-agent systems [4], organic computing [5,6], context- [7] and situation-aware systems [8], reflective computing [9], model-predictive control, as well as in the models@run-time community [10].

Self-aware computing systems are understood as having two main properties [11]. They

- learn models, capturing knowledge about themselves and their environment (such as their structure, design, state, possible actions, and run-time behavior) on an ongoing basis;
- reason using the models (to predict, analyze, consider, or plan), which enables them to act based on their knowledge and reasoning (for example, to explore, explain, report, suggest, self-adapt, or impact their environment)

and do so in accordance with high-level goals, which can change. For this year's edition, we had a special theme on IoT topics.

The SeAC workshop on self-aware computing provides a forum to foster interaction and collaborations between the respective research communities, raising the awareness about related research efforts and synergies that can be exploited to advance the state of the art. The workshop was initiated by the 2015 Dagstuhl Seminar 15041 on model-driven algorithms and architectures for self-aware computing systems, which brought together 45 international experts.

The EMSAC workshop aims to address the issues and challenges concerning the evaluation of the quality (of the design and execution) of autonomous, self-aware and self-managing systems, as well as measurements concerning or related to the self-* or autonomous features of a system. A particular interest concerns the approaches which manage the trade-offs among various different quality attributes and focus on the autonomy of the system. Existing quality assurance and evaluation solutions should be adapted and improved to meet the evolving and dynamic requirements of self-aware and self-managing systems. Novel design and evaluation solutions are needed to face the complexity of autonomy and self-management, as well as to evaluate and measure the self-* mechanisms and the advantages they bring, also with respect to the effort spent for their development.

The EMSAC workshop interests concern solutions proposed in various application domains which include cloud computing, smart systems, IoT, cyber-physical systems, systems-of-systems, robotics, autonomous driving, and social networks. EMSAC aims to discuss the principles of quality assurance and evaluation/measurement for autonomous and self-managing systems and mechanisms, the current trends, the future issues and challenges to be addressed at design and run-time. The EMSAC workshop invites researchers and practitioners from industry and academic environments to share their solutions, ideas, visions, and doubts in the quality assurance and evaluation in autonomic computing. During the EMSAC workshop, we enabled discussions and collaborations among the software engineers interested in these themes.

2. Structure of the Workshop

The workshop program included four sessions. The first session was dedicated to the invited talk, the following two sessions to the presentation of the peer-reviewed accepted long and short papers of the EMSAC and the SeAC workshops, respectively. In the last session, the presentations of EMSAC and SeAC were complemented by three presentations of the *1st IEEE International Workshop on Self-Protecting Systems* (SPS). SPS (<https://sites.google.com/view/sps19/home>) is a forum for researchers working at the intersection of autonomic computing and cyber-security. The papers covered different aspects of self-protection: assuring security compliance in self-adaptive systems, autonomous defense in Industrial Internet of Things, and performance evaluation of deep-learning for intrusion response. The detailed program of the event is available online on the workshop Web page (<http://emsac2019.informatik.uni-wuerzburg.de/program/>). The following sections of this report briefly present the contribution of the invited talk and the presented papers of the EMSAC and the SeAC workshops.

2.1. Session 1—Invited Keynote

The workshop started with a keynote given by Prof Betty H.C. Cheng from the Michigan State University on *A Requirements-Driven and Context-Aware Approach to Assurance of Autonomous Systems*. The keynote presented current developments towards a requirements-driven and context-aware autonomous system and overviewed a framework that supports run-time monitoring and adaptation

of tests for evaluating whether an autonomous system satisfies, or is even capable of satisfying, its requirements given its current execution context. Furthermore, Prof Betty H.C. Cheng described specific techniques that instantiate this framework, which apply a multidisciplinary approach to support requirements-based adaptive testing of an autonomous system at run-time. The inspiring keynote raised fruitful discussion on the current state of the art in testing and assurance under uncertainty in autonomous systems.

2.2. Session 2—EMSAC

The session of EMSAC presentation was started by the PhD student Sona Ghahremani from the Hasso Plattner Institute at the University of Potsdam in Germany. She presented a joint work with Holger Giese entitled *Performance Evaluation for Self-Healing Systems: Current Practice & Open Issues*. In their work [12], the authors provide a systematic literature review of approaches concerning the evaluation of self-healing systems. One of the main findings, which also raised discussions, was the miss of using real systems and demonstrators for evaluating the degree of self-healing in systems rather than simulation.

The second EMSAC workshop presentation was given by the PhD student André Bauer from the University of Würzburg, Germany, entitled *Systematic Search for Optimal Resource Configurations of Distributed Applications*. It was authored by André Bauer, Simon Eismann, Johannes Grohmann, Nikolas Herbst, and Samuel Kounev. The authors proposed and implemented a solution based on an adapted hill-climbing algorithm for finding all optimal configurations in a feasible time in micro-architecture approaches [13].

The third EMSAC presentation by Prof Sven Tomforde from the University of Kiel, Germany raised discussions on the trade-off between stability and benefits of adaptation decisions for adaptive systems. In his presentation entitled *From “normal” to “abnormal” self-adaptation: A concept for determining expected adaptation efforts*, Sven proposed a framework for the measurement of self-adaptivity and self-organization in complex systems based on the idea of identifying what is a normal and what is an abnormal adaptation [14].

2.3. Session 3—SeAC

Thereafter, the three presentations of the SeAC 2019 workshop followed. The first presentation hold by the PhD student Vladimir Podolskiy from the Technical University of Munich, Germany focused on *Metrics for Self-Adaptive Queuing in Middleware for Internet of Things* co-authored with Peeranut Chindanonda and Michael Gerndt. The paper presented several metrics for automating the scaling of message queuing subsystems and evaluated them on CPU-intensive and blocking I/O-intensive tasks [15].

The second SeAC presentation was held by the PhD student Johannes Grohman from the University of Würzburg, Germany about *Utilizing Clustering to Optimize Resource Demand Estimation Approaches*. In his joint work with Simon Eismann, André Bauer, Marwin Züfle, Nikolas Herbst, and Samuel Kounev, Johannes proposes an approach for resource demand estimation which uses automated clustering for grouping training sets into groups of similar optimization behavior to tailor optimization towards certain groups of resource traces in a self-aware manner [16].

The presentation session of the SeAC workshop was closed by Prof Cristina Abad from the Escuela Superior Politécnica del Litoral, Ecuador. The presentation entitled *Optimizing Cloud Caches For Free: A Case for Autonomic Systems with a Serverless Computing Approach*, authored by Xavier Andrade, Jorge Cedeno, Edwin F. Boza, Harold Aragón, Cristina L. Abad, and Jorge Murillo, discusses that serverless computing platforms can be leveraged to solve complex optimization problems that arise during self-tuning loops, and thus can be used to optimize resources in cloud caches, for free [17].

3. Contributions to this Special Issue

Kaiser et al. propose in their paper *Towards Self-Aware Multirotor Formations* [18] a framework to combine self-aware computing with multirotor formations for ordering drone traffic based on the principle of platooning. Platooning arranges vehicles in a string behind the lead vehicle. A modification for the LRA-M self-awareness loop is proposed. The implemented framework is able to offer two flight modes that enable waypoint following and the self-awareness module to find a path through scenarios, where obstacles are present on the way, onto a goal position. The evaluation shows that the proposed framework can use self-awareness to learn about its environment, avoid obstacles, and can successfully move a platoon of drones through multiple scenarios.

The paper *A Taxonomy of Techniques for SLO Failure Prediction in Software Systems* [19] by Grohmann et al. describe a taxonomy for the prediction of Service Level Objective (SLO) failures. The taxonomy classifies related work along the dimensions of the prediction target (e.g., anomaly detection, performance prediction, or failure prediction), the time horizon (e.g., detection or prediction, online or offline application), and the applied modeling type (e.g., time series forecasting, machine learning, or queuing theory). The classification is derived based on a systematic mapping of relevant papers.

The paper *Self-Adaptive Data Processing to Improve SLOs for Dynamic IoT Workloads* [20] by Chindanonda, Podolskiy, and Gerndt addresses the challenge of adapting the capacity of the data processing part of IoT pipeline in response to dynamic workloads for centralized IoT scenarios where the quality of user experience matters, e.g., interactivity and media streaming as well as the predictive maintenance for multiple moving vehicles, centralized analytics for wearable devices and smartphones. In their paper, the authors propose augmentations to the computation schemes of data processing component's desired replicas count from the previous work; these augmentations aim to repurpose original sets of metrics to tackle the task of SLO violations minimization for dynamic workloads instead of minimizing the cost of deployment in terms of instance seconds.

Boza et al. show in their contribution *On Implementing Autonomic Systems with a Serverless Computing Approach: The Case of Self-Partitioning Cloud Caches* [21] that serverless computing cloud services can be leveraged to solve the complex optimization problems that arise during self-tuning loops and can be used to optimize cloud caches for free. To illustrate that our approach is feasible and useful, we implement a modified version of Redis that optimizes memory management in the multi-instance Redis scenario. A cost analysis shows that the serverless computing approach can lead to significant cost savings.

In their paper *Evaluation of Self-Healing Systems: An Analysis of the State of the Art and Required Improvements* [22], the authors Ghahremani and Giese study the state of the art for evaluating the performance of self-healing systems by means of a systematic literature review. They provide a classification of different input types for such systems and analyze the limitations of each input type. A main finding is that the employed inputs are often not sophisticated regarding the considered characteristics for failure occurrences. To further study the impact of the identified limitations, the authors present experiments demonstrating that wrong assumptions regarding the characteristics of the failure occurrences can result in large performance prediction errors, disadvantageous design-time decisions concerning the selection of alternative self-healing approaches, and disadvantageous deployment-time decisions concerning parameter tuning.

Tomforde and Goller argue in their paper *To Adapt or Not to Adapt: A Quantification Technique for Measuring an Expected Degree of Self-Adaptation* [23] that too much adaptation is in conflict with goals such as stability and user acceptance. Consequently, a kind of situation-dependent degree of adaptation is desired, which defines the amount and severity of tolerated adaptations in certain situations. As a first step into this direction, their paper presents a quantification approach for measuring the current adaptation behavior based on generative, probabilistic models. The authors analyze the behavior of their approach in three application scenarios: urban traffic control, the swidden farming model, and data communication protocols. Furthermore, the authors define a research roadmap with challenges for an overall measurement framework.

4. Conclusions

The participants of the successful first edition of the EMSAC workshop joint with the SeAC workshop series were highly satisfied with the high quality of the keynote and the presentations and benefited from interesting discussion on future directions in the field. The next edition of the SeAC workshop will be a part of the new *International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)* which results from the merge of the IEEE International Conference on Autonomic Computing (ICAC) and the IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO) taken place at Washington, DC, from August 17 until August 21, 2020.

Acknowledgments: We would like to thank Ingrid Nunes and Robert Birke, the workshop co-chairs of ICAC/SASO 2019, for their collaboration in the organization of the EMSAC and SeAC 2019 workshops. We thank the authors of the submitted papers, as well as the program committee members for their contributions. A special thank goes to Betty H.C. Cheng for her interesting and inspiring keynote. We further thank André Bauer for his effort in publicity management of the workshops. Furthermore, we thank all presenters for their valuable contribution and the audience for interesting discussions. Additionally, we thank all authors for contributing to this special issue. Lastly, we thank Florian Erhard for preparing this report as well as Blink Yu as the assigned editor of the special issue for his cooperation and support.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ACSOS	International Conference on Autonomic Computing and Self-Organizing Systems
EMSAC	Workshop on Evaluations and Measurements in Self-Aware Computing Systems
ICAC	IEEE International Conference on Autonomic Computing
SASO	IEEE International Conference on Self-Adaptive and Self-Organizing Systems
SeAC	Workshop on Self-Aware Computing
SLO	Service Level Objective
SPS	IEEE International Workshop on Self-Protecting Systems

References

1. Kephart, J.O.; Chess, D.M. The Vision of Autonomic Computing. *Computer* **2003**, *36*, 41–50. [[CrossRef](#)]
2. Cheng, B.H.; de Lemos, R.; Giese, H.; Inverardi, P.; Magee, J.; Andersson, J.; Becker, B.; Bencomo, N.; Brun, Y.; Cukic, B.; others. Software engineering for self-adaptive systems: A research roadmap. In *Software Engineering for Self-Adaptive Systems*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 1–26.
3. Krupitzer, C.; Roth, F.M.; VanSyckel, S.; Schiele, G.; Becker, C. A survey on engineering approaches for self-adaptive systems. *Pervasive Mob. Comput.* **2015**, *17*, 184–206. [[CrossRef](#)]
4. Wooldridge, M. *An Introduction to MultiAgent Systems*, 2nd ed.; Google-Books-ID: X3ZQ7yeDn2IC; John Wiley & Sons: Hoboken, NY, USA, 2009.
5. Allrutz, R.; Cap, C.; Eilers, S.; Fey, D.; Haase, H.; Hochberger, C.; Karl, W.; Kolpatzik, B.; Krebs, J.; Langhammer, F.; et al. *Organic Computing: Computer-und Systemarchitektur im Jahr 2010*; Technical Report, VDE/ITG/GI Position Paper; VDE/ITG/GI: Frankfurt am Main, Germany, 2003.
6. Tomforde, S.; Prothmann, H.; Branke, J.; Hähner, J.; Mnif, M.; Müller-Schloer, C.; Richter, U.; Schmeck, H. Observation and Control of Organic Systems. In *Organic Computing—A Paradigm Shift for Complex Systems*, Springer: Berlin/Heidelberg, Germany, 2011; pp. 325–338.
7. Schilit, B.; Adams, N.; Want, R. Context-Aware Computing Applications. In Proceedings of the 1994 First Workshop on Mobile Computing Systems and Applications, Washington, DC, USA, 8–9 December 1994; pp. 85–90. [[CrossRef](#)]
8. Endsley, M.R. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors* **1995**, *37*, 32–64. [[CrossRef](#)]
9. Smith, B.C. Procedural Reflection in Programming Languages. PhD thesis, Laboratory for Computer Science, Massachusetts Institute of Technology: Cambridge, MA, USA, 1982.
10. Blair, G.; Bencomo, N.; France, R.B. Models@ run.time. *Computer* **2009**, *42*, 22–27. [[CrossRef](#)]

11. Kounev, S.; Lewis, P.R.; Bellman, K.L.; Bencomo, N.; Cámara, J.; Diaconescu, A.; Esterle, L.; Geihs, K.; Giese, H.; Götz, S.; Inverardi, P.; Kephart, J.O.; Zisman, A. The Notion of Self-aware Computing. In *Self-Aware Computing Systems*; Springer International Publishing: Cham, Switzerland, 2017; pp. 3–16. [[CrossRef](#)]
12. Ghahremani, S.; Giese, H. Performance Evaluation for Self-Healing Systems: Current Practice & Open Issues. In Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 116–119. [[CrossRef](#)]
13. Bauer, A.; Eismann, S.; Grohmann, J.; Herbst, N.; Kounev, S. Systematic Search for Optimal Resource Configurations of Distributed Applications. IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 120–125. [[CrossRef](#)]
14. Tomforde, S. From "Normal" to "Abnormal": A Concept for Determining Expected Self-Adaptation Behaviour. In Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 126–129. [[CrossRef](#)]
15. Chindanonda, P.; Podolskiy, V.; Gerndt, M. Metrics for Self-Adaptive Queuing in Middleware for Internet of Things. In Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 130–133. [[CrossRef](#)]
16. Grohmann, J.; Eismann, S.; Bauer, A.; Züfle, M.; Herbst, N.; Kounev, S. Utilizing Clustering to Optimize Resource Demand Estimation Approaches. In Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 134–139. [[CrossRef](#)]
17. Andrade, X.; Cedeno, J.; Boza, E.F.; Aragon, H.; Abad, C.; Murillo, J.R. Optimizing Cloud Caches For Free: A Case for Autonomic Systems with a Serverless Computing Approach. In Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, 16–20 June 2019; pp. 140–145. [[CrossRef](#)]
18. Kaiser, D.; Lesch, V.; Rothe, J.; Strohmeier, M.; Spieß, F.; Krupitzer, C.; Montenegro, S.; Kounev, S. Towards Self-Aware Multirotor Formations. *Computers* **2020**, *9*, 7. [[CrossRef](#)]
19. Grohmann, J.; Herbst, N.; Chalbani, A.; Arian, Y.; Peretz, N.; Kounev, S. A Taxonomy of Techniques for SLO Failure Prediction in Software Systems. *Computers* **2020**, *9*, 10. [[CrossRef](#)]
20. Chindanonda, P.; Podolskiy, V.; Gerndt, M. Self-Adaptive Data Processing to Improve SLOs for Dynamic IoT Workloads. *Computers* **2020**, *9*, 12. [[CrossRef](#)]
21. Boza, E.F.; Andrade, X.; Cedeno, J.; Murillo, J.; Aragon, H.; Abad, C.L.; Abad, A.G. On Implementing Autonomic Systems with a Serverless Computing Approach: The Case of Self-Partitioning Cloud Caches. *Computers* **2020**, *9*, 14. [[CrossRef](#)]
22. Ghahremani, S.; Giese, H. Evaluation of Self-Healing Systems: An Analysis of the State-of-the-Art and Required Improvements. *Computers* **2020**, *9*, 16. [[CrossRef](#)]
23. Tomforde, S.; Goller, M. To Adapt or Not to Adapt: A Quantification Technique for Measuring an Expected Degree of Self-Adaptation. *Computers* **2020**, *9*, 21. [[CrossRef](#)]

