EDITOR: Schahram Dustdar, dustdar@infosys.tuwien.ac.at

DEPARTMENT: INTERNET OF THINGS, PEOPLE AND PROCESSES

Service-Oriented Computing: A Trajectory for Research to 2030

Pierluigi Plebani, *Politecnico di Milano, 20133, Milano, Italy*

Stefan Schulte, *Hamburg University of Technology, 21079, Hamburg, Germany*

Damian Andrew Tamburri, *Politecnico di Milano, 20133, Milano, Italy*

Schahram Dustdar, *Technische Universität Wien, 1040 Wien, Austria*

Since service-oriented computing was introduced as a major topic in the industry and the research community, 20 years have passed. Today, service orientation has become a commodity in many areas. This has also changed the foci of the research community by a very large degree. In this article, we analyze the current state of the research in the field and give an outlook on future research topics.

Exactly 20 years ago, one of the seminal works in the research area defined service-oriented computing (SOC) "as the computing paradigm that utilizes services as fundamental elements for developing applications."[1]

Although, at that time, "services" meant Web services, this definition has passed the test of time, and it keeps capturing the essence of SOC. Over time, solutions moved from SOAP to Representational State Transfer or GraphQL, Web services were decomposed into microservices, and entire businesses were established around application programming interfaces (APIs). In any case, whenever we have a software component supporting functionalities offered by a provider to an audience of consumers, we keep calling this approach SOC.

In the early years of SOC research, a series of workshops at the International Conference on Very Large Data Bases (VLDB) named TES, i.e., *Technology of E-Services*,[2] paved the way to a research agenda that, during the next two decades, was finalized and repeatedly revised.[3,4,5] Since then, voices of criticism have emerged from the research community concerning the actual capacity of innovation of what SOC is proposing with respect to distributed computing approaches at the time.[6]

1089-7801 © 2024 IEEE
Digital Object Identifier 10.1109/MIC.2023.3338908
Date of current version 3 June 2024.

At this point, from a research perspective and after 20 years of honorable service, the voices of criticism are getting louder and louder. Therefore, we believe that it is important to understand whether SOC is still a relevant subject to be studied, from which perspective, and with what methodological approaches. To address these questions, this article takes a pragmatic approach to understanding at which point we are in the SOC research community, which directions have been taken so far, and if there is still space for high-quality research.

THE CURRENT STATE OF THE SOC COMMUNITY

To assess the current state of the SOC community, we have analyzed papers published at top conferences in the field, i.e., the IEEE International Conference on Web Services (ICWS) and the International Conference on Service-Oriented Computing (ICSOC). For this, we have identified a number of high-level research topics. We started by identifying which topics were highly relevant at ICSOC and ICWS in the years 2010–2012. Afterward, we had a look at the years 2020–2022 and checked how the topics evolved.

The results can be seen in Table 1. The table is not exhaustive; i.e., we are only showing the 10 most important categories. The "total" number of papers describes all papers published at a venue during a particular period, i.e., also from those categories not shown in the table.

TABLE 1. Publications at ICSOC and ICWS from 2010 to 2012 and 2020 to 2022.*

Category	ICSOC		ICWS	
	2010–2012	2020–2022	2010–2012	2020–2022
Service compositions	24	2	28	6
Service discovery/selection	13	3	26	5
Fog/edge	0	10	0	28
(Business) processes	14	10	8	1
IoT/smart systems	1	13	0	9
Security/trust	4	1	10	7
Cloud	5	9	0	6
Service/app recommendation	0	0	0	17
Microservices	0	11	0	5
Human-provided services/ crowdsourcing	3	3	0	5
Total	96	87	105	125

*ICSOC: International Conference on Service-Oriented Computing; ICWS: IEEE International Conference on Web Services; IoT: Internet of Things.

Naturally, it was not always possible to find an absolute truth on which topic a particular paper belongs to, e.g., because a paper aimed at two distinct categories. In this case, the paper was classified into the topic that it focused on. Importantly, we primarily considered the research object to do the classification and less the method or other aspects applied. For instance, machine learning is, these days, very often used as a method in SOC. However, it is rather seldom the case that the focus of a paper is on machine learning itself, i.e., that machine learning is the research object. Only full/regular papers from the main tracks of ICWS and ICSOC have been regarded. Invited (if identifiable), industry, and short papers are not considered.

Of course, the information provided in the table can only provide a coarse-grained overview of the developments of SOC research in the past decade. Nevertheless, it provides some interesting insights. As can be seen from the table, some research topics are "evergreens," i.e., were being covered a decade ago and are also still in consideration today. This is the case, e.g., for business processes; cloud computing; and human-provided services, where the focus has, however, mostly shifted to crowdsourcing. Of course, the foci of these research categories have widely evolved in the last 10 years. Not surprisingly, novel technologies and paradigms, like microservices or fog and edge computing, which were not topics at the beginning of the 2010s, have risen to prominence.

It can also be seen that basic SOC functionalities like service compositions and service discovery are not so much in the focus any longer. The reasons for this may range from the insight that a particular basic functionality has, indeed, not played a very big role in practice or that this functionality is, these days, common. Notably, the topic of service discovery has, to some degree, evolved into service and API recommendations, which are more proactive approaches to provide potential service consumers with information about fitting software entities.

As a second insight, the most prominent research categories from 2020 to 2022 are even more technology-driven than was the case from 2010 to 2012; i.e., research is focusing on solving particular issues for particular technologies instead of fundamental SOC research topics. Last but not least, the heterogeneity of topics published at the discussed venues has become much greater. In recent years, many smaller topics have been presented, ranging from service applications in video streaming and big data to papers on SPARQL queries and private data management, which do not have a direct link to SOC. Nevertheless, the community deems these topics as fitting for its top venues.

In our opinion, these insights show that fundamental SOC functionalities have become a widely accepted commodity, as has SOC itself. Basic capabilities of service-oriented systems, which were already named at the beginning of SOC research—e.g., encapsulation of (business) functions, loose coupling of services, service autonomy, or service reusability[7]—are, today, applied in many types of software systems.

RESEARCH ROAD MAP

After about a quarter of a century, it is a fact that SOC has deeply affected the way in which software-based systems have been—and are still—designed, deployed, and executed. If there is now the impression that SOC is mature enough to be considered relevant for practitioners and has a significant impact on the industry, the question remains if there is any space for the SOC research community to keep this influence.

To answer this question, we have started with the "Service Computing Manifesto,"[4] published in 2017, celebrating the 10 years from the publication of the original SOC manifesto.[8] In 2017, a road map was proposed around four main topics: service design, service composition, crowdsourcing-based reputation, and the Internet of Things (IoT). As emerged in the previous section, these topics have been considered differently in the community—in some cases, e.g., composition, with a decreasing interest and, in some other cases, e.g., the IoT, confirming the expectations.

With the aim of building a road map that can suggest research topics for the upcoming years, we start from some megatrends that are already influencing—and will do even more in the future—the design and development of software based on services: notably, the rise of quantum computing, the convergence into the cloud–edge compute continuum, the ever-increasing adoption of machine learning and AI, trustworthiness as a key requirement, the sustainability of SOC, and the transition from good dominance to service dominance logic.

Quantum Services Computing

Quantum computing is an early-stage yet rapidly evolving field that holds great promise for solving complex problems that are beyond the reach of classical computers. At the forefront of quantum computing, there are several key areas of research and development that relate to SOC. One prominent example is Shor's algorithm, which demonstrates the potential for quantum computers to factor large numbers exponentially faster than classical computers. Services around Shor's and similar algorithms might lead to service-oriented architectures with completely different service-level profiles to be studied and experimented upon further. Similarly, other algorithms, such as Grover's algorithm for service-data search, may be used to improve service operations at a large scale, thereby improving existing service solutions on a global scale.

Also, quantum systems are prone to errors caused by environmental disturbances and imperfections in hardware. A question emerges here, namely, could service continuity principles and technologies play a role in addressing such circumstances?

Furthermore, achieving quantum supremacy means demonstrating a quantum computer's ability to solve a problem that is infeasible for classical computers. Showing the supremacy in service-oriented systems—for example, thinking of large-scale supply chain management—and blending it with current service operations is worthy of future research.

Cloud–Edge Compute Continuum

The trend toward distributed services has gained a lot of traction, as can be seen by the large number of papers on fog and edge computing published in recent years. More recently, this has evolved into a discussion of the cloud–edge compute continuum, which considers cloud and edge computing models as part of a continuous compute model. Despite this, a seamless transfer of services from one device to another one, considering the specific technical and nontechnical context, has still not been achieved. The number of research questions to be tackled ranges from lightweight virtualization techniques over resource- and energy-aware resource allocation approaches to security and privacy solutions.

Machine Learning and AI

As has already been outlined, machine learning is, these days, a very common method in SOC, used, e.g., to optimize service topologies. Also, machine learning is already provided today as a cloud service by, e.g., Amazon AWS or IBM. With the recent advent of novel machine learning approaches like federated learning, the notion of machine learning in SOC will further evolve. Especially, the basic principles of SOC can help to establish common interfaces for federated learning, allowing us to easily replace single clients and even model aggregators, which are, these days, very often located in the cloud. Another recent trend is edge intelligence, i.e., to deploy AI close to the data sources with or without support by cloud resources. Here, the research questions span from approaches that are able to consider the capabilities of beyond-5G mobile networks, integration with the compute continuum, and novel machine learning algorithms, which can, e.g., self-adapt to trade off resource demands and accuracy.

Trustworthiness

With SOC being built on a relationship between service providers and service consumers, especially when these two roles are held by entities for which no agreement has already been established, it is fundamental to define mechanisms that define and strengthen trust among them. This requires balancing between the observability of a service in terms of its structure and behavior to increase the transparency for the consumer and the need to protect the intellectual property embedded in the service that constitutes the business for the provider. From another angle, more related to the semantics of the service provisioning, trust also implies having a correct service that provides the right data in response to consumer requests.

In this area, among others, first results have been obtained by the adoption of solutions based on zero-knowledge proof protocols that allow the verification of credentials while limiting data leakage or blockchain-based solutions that allow getting rid of mediators to ensure trust. In the future, these and other approaches can be followed to establish trustworthy, service-based relations.

Sustainability

As in many IT-related domains, for SOC it is, today, also fundamental to consider the need to reach the Sustainable Development Goals defined by the United Nations. Among them, SOC research can especially hold a significant role in contributing to climate actions and the reduction of resource usage.

As is happening for software development in general, it is now important to consider the amount of resources along the whole service lifecycle, from the design and development phases—which have to consider more than just the cost of the resources as the metric for reducing their usage—to the execution phase, where it is very difficult to exactly monitor the amount of energy consumed.

Service Dominance Logic

This was probably just by chance, but, in the years in which SOC was establishing itself, the theorization of the so-called service dominance logic was also taking shape in the managerial literature,[9] where a service is defined as "the application of specialized competences (knowledge and skills) through deeds, processes, and performances for the benefit of another entity or the entity itself." Also, on top of this theory, information systems literature has put the basis for service science where technological aspects of services are linked with organizational ones.

Although SOC is more related to the technical and technological side, the resulting services are components that will live in a complex environment with the presence of software systems and human beings. In this sense, we think that SOC research should consider more the interrelation between providers and consumers to allow the envisioned cocreation of service value. This also implies reducing accessibility barriers to the services, making them easy to discover and use. Thus, the service design and the service lifecycle in general must address the complexity of interactions with human beings.

CONCLUSION

This article recaps the state of the art in SOC as reflected in top conferences addressing the topic. The picture we elaborate on is projecting a somewhat conflicting perspective. On the one hand, SOC is nowhere near dead but, rather, elaborating an inclusion and openness—in the style of its foundational open-world assumption—toward other related disciplines, such as blockchain, machine learning operations, quantum computing, and more. On the other hand, the very nature of services seems to be no longer of interest to the community, with a few sporadic and stovepiped exceptions to be found over recent years. Could this mean a lack of interest in the basic foundations of SOC and how they are changing to cope with the different contexts, domains, and technologies they reflect? Could this mean those foundations are left to emerge in continuity with the state of practice rather than being tackled from a foundational, more theoretical perspective?

We conclude that the road ahead can only reflect a proactive self-reflection of the communities and disciplines in question to find the right strategy and approach, reflective of a sound way forward. We believe this article will serve as a basic and rudimentary first step in that respect.

REFERENCES

1. M. P. Papazoglou and D. Georgakopoulos, "Introduction: Service-oriented computing," *Commun. ACM*, vol. 46, no. 10, pp. 24–28, 2003, doi: 10.1145/944217.944233.
2. F. Casati, U. Dayal, and M.-C. Shan, "Report on the VLDB workshop on technologies for E-services (TES)," *ACM SIGMOD Rec.*, vol. 29, no. 4, pp. 11–15, 2000, doi: 10.1145/369275.369276.
3. M. P. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, "Service-oriented computing: A research roadmap," *Int. J. Cooperative Inf. Syst.*, vol. 17, no. 2, pp. 223–255, 2008, doi: 10.1142/S0218843008001816.
4. A. Bouguettaya et al., "A service computing manifesto: The next 10 years," *Commun. ACM*, vol. 60, no. 4, pp. 64–72, 2017, doi: 10.1145/2983528.
5. D. Bermbach et al., "A research perspective on fog computing," in *Proc. Service-Oriented Comput.*

–*Workshops (ICSOC)*, 2018, pp. 198–210, doi: 10.1007/978-3-319-91764-1_16.

6. F. Arbab, "Will the real service oriented computing please stand up?" in *Formal Aspects of Component Software*, L. S. Barbosa and M. Lumpe, Eds., Berlin, Germany: Springer-Verlag, 2012, pp. 277–285.
7. M. P. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, "Service-oriented computing: State of the art and research challenges," *Computer*, vol. 40, no. 11, pp. 38–45, Nov. 2007, doi: 10.1109/MC.2007.400.
8. M. Papazoglou et al., "Service-oriented computing," The S.O.C. Manifesto, 2003. [Online]. Available: https://www.csd.uoc.gr/~hy565/docs/pdfs/papers/soc_manifesto.doc
9. S. L. Vargo and R. F. Lusch, "Service-dominant logic 2025," *Int. J. Res. Marketing*, vol. 34, no. 1, pp. 46–67, 2017, doi: 10.1016/j.ijresmar.2016.11.001.

PIERLUIGI PLEBANI is an associate professor at Politecnico di Milano, Italy. Contact him at pierluigi.plebani@polimi.it.

STEFAN SCHULTE is full professor for data engineering at Technische Universität (TU) Hamburg and leads the Christian Doppler Laboratory Blockchain Technologies for the Internet of Things. Contact him at stefan.schulte@tuhh.de.

DAMIAN ANDREW TAMBURRI is an associate professor at the Eindhoven University of Technology, Jheronimus Academy of Data Science, s'Hertogenbosch, The Netherlands, with a double-affiliation at Politecnico di Milano, Italy, as well. Contact him at damianandrew.tamburri@polimi.it.

SCHAHRAM DUSTDAR is full professor of computer science heading the Research Division of Distributed Systems, TU Wien, Austria. Contact him at dustdar@dsg.tuwien.ac.at.