Distributed Computing Continuum Systems – Opportunities and Research Challenges

Victor Casamayor Pujol(✉), Praveen Kumar Donta, Andrea Morichetta, Ilir Murturi, and Schahram Dustdar

Distributed Systems Group, TU Wien, Vienna, Austria
{casamayor,donta,morichetta,murturi,dustdar}@dsg.tuwien.ac.at

Keywords: Distributed Computing Continuum Systems · Human body · SLO · Self-adaptation · AI · Monitoring · Knowledge · Security

1 Vision

Internet distributed systems are subjected to a new transformation thanks to the success of Cloud Computing. The scientific community is looking beyond Cloud Computing and tackling its limitations with new paradigms such as Edge Computing [11]. In this tutorial, we aim at showing that we need to look towards a Distributed Computing Continuum Systems (DCCS) paradigm; further, we discuss that the complexity of such a paradigm requires new methods and techniques [2]. Some of these can be inspired by other complex systems that we have in our surroundings. In this regard, we aim at showing that the human body can be a useful analogy for the DCCS, hence, we aim to use it to encourage the general audience and the scientific community in this research.

A key aspect of complex systems, such as the human body, is that they are homeostatic, keeping their vital signs within their expected range, such as the body temperature. In Cloud systems, Service Level Indicators (SLIs) are metrics of the infrastructure, such as CPU usage, that are monitored to be within the ranges specified by the Service Level Objectives (SLOs). Hence, if the SLI is not in the range, an elastic strategy is triggered, and the SLI returns within its range. DCCS need this homeostatic behavior, however, the heterogeneity, distribution, scale, and interconnection of the infrastructure challenge the current approach. Further, the seamless integration between the infrastructure and the application on DDCS demands focus on the infrastructure.

Interestingly, when the human body has an issue, we feel first a general malaise, and then we can trace it down to the specific part that is not right. Following this idea, we assume that DCCS also need a hierarchical structure for its definition, from abstract and general metrics such as cost, quality, and performance [3], to the aforementioned CPU usage of a component. Last, each of these defining metrics would be an ontological Markov Blanket [5], having the capacity to sense and act on its environment, further, this would allow us to

J. Troya et al. (Eds.): ICSOC 2022 Workshops, LNCS 13821, pp. 405–407, 2023.
https://doi.org/10.1007/978-3-031-26507-5_41

actually define Markov Blankets at the needed granularity [9] enabling a scale-free methodology. Our road map for this vision works on four parallel tracks.

2 Research Lines

Self-adaptation. We propose an artifact called DeepSLO, which is a hierarchically structured set of SLOs that relate causally and purposefully. Hence, we could capture in a single artifact both the state of the infrastructure and the application running on top. This would instantiate our vision of the homeostatic DCCS including elastic strategies associated with each SLO. Nonetheless, the development of DeepSLOs is still in progress as many challenges arise once you scratch their surface, in this regard, we are currently investigating how to model them and their relations, or how to analyze it.

Inter-relations. A captivating idea derived from biology [6] and neural science [4] is to have an architecture composed of independent, autonomous models. These models represent mechanisms that - in the causality sense [8] - perform specific functions and *cause* the accomplishment of an action, e.g., clicking on the switch to turn on the light. These modules learn from the data to specialize in distinct tasks. These mechanisms adapt to various environments through reciprocal interaction, guaranteeing better generalization and adjusting to various applications and SLOs [7].

Monitoring & Knowledge. A healthy body and mind result in a happy life; similarly, healthy computing devices and efficient processing improve the quality of service. The context of the DCCS's health monitoring includes continuous diagnosis, prediction, and cure of problems to maintain the system healthy. Representation learning is a solution to monitor the system and give notifications and predictions timely to the administrators. We discuss possible solutions using recent technologies as future research directions [1].

Security. In the human body, the immune system distinguishes pathogens from the organism's healthy tissues via protein markers in each cell. The immune system is a safelist; through protein distinguishment, it will eliminate any organic entity that fails to present that specific body protein marker. In the context of DCCS, threats and malicious activities should be identified and eliminated like in the human body. Protecting resources and identifying malicious activities are among the core challenges that should be addressed. Therefore, we investigate the current trend toward protecting distributed resources with Zero Trust [10] concepts, its particular characteristics, and challenges to implement and execute on DCCS.

3 Conclusion

In this tutorial, we have conveyed key ideas for the future generation of Internet distributed systems. Leveraging the human body as an analogy for DCCS allows a clearer definition of the requirements for the systems, as well as inspiration for the research. Moreover, it helps convey the challenges to a broader audience increasing the impact of the research and, hopefully, involving a larger community in the future. We have presented our main research lines with their current objectives and challenges to spark discussion among the scientific community and boost this topic.

References

1. Donta, P.K., Dustdar, S.: The promising role of representation learning for distributed computing continuum systems. In: 2022 IEEE International Conference on Service-Oriented System Engineering (SOSE), pp. 126–132. IEEE (2022)
2. Dustdar, S., Casamajor Pujol, V., Donta, P.K.: On distributed computing continuum systems. IEEE Trans. Knowl. Data Eng. IEEE Computer Society (2022). https://doi.org/10.1109/TKDE.2022.3142856
3. Dustdar, S., Guo, Y., Satzger, B., Truong, H.L.: Principles of elastic processes. IEEE Internet Comput. **15**(5), 66–71 (2011). https://doi.org/10.1109/MIC.2011.121
4. Goyal, A., et al.: Coordination among neural modules through a shared global workspace. ArXiv abs/2103.01197 (2022)
5. Kirchhoff, M., Parr, T., Palacios, E., Friston, K., Kiverstein, J.: The Markov blankets of life: autonomy, active inference and the free energy principle. J. Royal Soc. Interface **15**(138), 20170792. The Royal Society (2018). https://doi.org/10.1098/RSIF.2017.0792, http://royalsocietypublishing.org/doi/abs/10.1098/rsif.2017.0792
6. Kozachkov, L., Ennis, M., Slotine, J.J.E.: RNNs of RNNs: recursive construction of stable assemblies of recurrent neural networks (2021)
7. Morichetta, A., Pujol, V.C., Dustdar, S.: A roadmap on learning and reasoning for distributed computing continuum ecosystems. In: 2021 IEEE International Conference on Edge Computing (EDGE), pp. 25–31 (2021). https://doi.org/10.1109/EDGE53862.2021.00021
8. Parascandolo, G., Rojas-Carulla, M., Kilbertus, N., Schölkopf, B.: Learning independent causal mechanisms. In: ICML (2018)
9. Pujol, V.C., Raith, P., Dustdar, S.: Towards a new paradigm for managing computing continuum applications. In: IEEE 3rd International Conference on Cognitive Machine Intelligence, CogMI 2021, pp. 180–188. Institute of Electrical and Electronics Engineers Inc. (2021). https://doi.org/10.1109/COGMI52975.2021.00032
10. Rose, S., Borchert, O., Mitchell, S., Connelly, S.: Zero trust architecture. Technical report. National Institute of Standards and Technology (2020)
11. Shi, W., Dustdar, S.: The promise of edge computing. Computer **49**(5), 78–81. IEEE Computer Society (2016). https://doi.org/10.1109/MC.2016.145