Decentralized Systems

Samuel H. Christie V, *North Carolina State University, Raleigh, NC, 27695-8206, USA*
Lalana Kagal, *Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*
Alessandro Ricci, *University of Bologna, 47522, Cesena, Italy*
Munindar P. Singh, *North Carolina State University, Raleigh, NC, 27695-8206, USA*

Our world is filled with decentralized systems, human society and organizations being the most familiar examples. In such systems, we see multiple loci of ownership or control, representing different parties or administrative domains. That is, the members of a decentralized system are, belong to, or represent different people or organizations.

We use the term decentralized system to mean any system formed of autonomous entities, with a special interest in cases where the entities are heterogeneous. Autonomy here refers to the decision-making capacity of an entity, meaning that it decides for itself. In other words, autonomy reflects the freedom to act. Heterogeneity here refers to the design, construction, and configuration of an entity, meaning that it is potentially built on distinct grounds from any other entity. In other words, heterogeneity reflects the freedom of a designer to apply any reasoning method on any available information.

Traditional computing is not able to deal well with decentralization because it has inadequate support for the interactions of autonomous and heterogeneous parties, and its abstractions are not suited to the kinds of concepts that are prominent in decentralized systems. As computing systems become more integrated into our lives and grow beyond the scale of a single organization, we will need better tools and methods for designing and implementing such systems. Specifically, we need approaches that give first-class status to the concepts hinted at above, including communications on semantic grounds, social norms (including laws, regulations, and established practices), accountability, and emergent behaviors.

Relevant research into decentralized systems is scattered across multiple (though partially overlapping) communities, including those focused on Web programming, multiagent systems (MAS), distributed computing, self-adaptive systems, and the semantic Web. This issue consolidates the state-of-the-art in decentralized systems.

In "Towards Decentralized Cloud Storage with IPFS: Opportunities, Challenges, and Future Considerations," Trinh Viet Doan, Yiannis Psaras, Jörg Ott, and Vaibhav Bajpai gave an overview of the "Interplanetary File System," a popular decentralized storage protocol and network. They give updates on the adoption and performance of IPFS, describe how the major components work, present several projects built on top of it including the FileCoin incentivized storage system, and close with challenges and open questions for the future of decentralized storage. Although this paper does not develop new techniques or software, it provides both an accessible introduction for new readers unfamiliar with IPFS and enough new information and thought-provoking questions to interest the moderately informed.

In "Dynamic Decentralization Domains for the Internet of Things," Gianluca Aguzzi, Roberto Casadei, Danilo Pianini, and Mirko Viroli described an approach to model dynamic cyber-physical systems based on a notion of concurrent distributed computational processes (DCPs). Their approach is designed for decentralized settings, such as those based on the Internet of Things, where there are many sensors and compute notes that are owned or controlled by different parties. Aguzzi et al.'s approach supported self-organization wherein the DCPs figure out the interesting regions (in physical but possibly in virtual) space, continually monitor those regions, respond to events, and adapt their responses. Aguzzi et al. provided a methodology leading up to programming abstractions through which DCPs can be readily specified to function in such an environment. Their approach makes possible ways to structure the interactions of DCPs such that the emergent behaviors of the system align with system objectives.

1089-7801 © 2022 IEEE
Digital Object Identifier 10.1109/MIC.2022.3216093
Date of current version 23 December 2022.

In "Interoperable AI: Evolutionary Race Towards Sustainable Knowledge Sharing," Stefan Sarkadi, Andrea G. B. Tettamanzi, and Fabien Gandon described a model simulation to study the sustainability of different strategies for cooperation in decentralized systems focusing on knowledge sharing among MAS. The authors suggest that the cost of knowledge sharing is greatly impacted by the ontologies being used by the agents and thereby the strategy for the alignment. The authors have identified seven strategies that range from using individual ontologies and translation services to using a common ontology. In their game-theoretic simulation, cooperating agents use different strategies to compete to be the best at sharing knowledge. The authors studied various hypotheses and found that, under the assumption of cooperation and the possibility of decentralization, translation strategies are more sustainable than shared ontology strategies for knowledge sharing between autonomous agents. However, currently there is a lack of translation services of this kind and the authors use the results of their study to motivate translation-based approaches for knowledge sharing in decentralized systems.

In "Exception Handling as a Social Concern," Matteo Baldoni, Cristina Baroglio, Roberto Micalizio, and Stefano Tedeschi described an approach to model and manage exceptions in intelligent distributed systems modeled as MAS. The approach frames exceptions at an organizational level, providing a social modeling based on the notion of responsibility. Accordingly, exception handling is more than a mechanism for coping with runtime problems: it is understood as a means of distributed composition, a software engineering tool that can be used to develop loosely coupled distributed systems with highly cohesive components. Baldoni et al. showed the approach at work by means of an extension developed upon an existing MAS platform called JaCaMo, which provides first-class abstractions for modeling agent organizations. In this approach, when agents play organizational roles they will also take on responsibility for those goals related to exception handling, specified through notification policies, integrated in the normative system. Accordingly, the exception handling mechanism is endowed with a normative power, yielding the obligations that agents receive for raising and handling goals.

SAMUEL H. CHRISTIE V is at North Carolina State University, Raleigh, NC, 27695, USA. He was a research associate at the School of Computing and Communications, Lancaster University, U.K. V received his Ph.D. degree in computer science from NC State University. Contact him at shcv@sdf.org.

LALANA KAGAL is a principal research scientist at the Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA, where she leads the Decentralized Information Group that focuses on topics at the intersection of technology, society, and policy. Her research interests include privacy preserving analytics, frameworks for federated learning, and fairness in machine learning. Contact her at lkagal@csail.mit.edu.

ALESSANDRO RICCI is an associate professor of Department of Computer Science and Engineering, Alma Mater Studiorum Università di Bologna, 47522, Cesena, Italy. His research interests include agents and multiagent systems as a paradigm for engineering software systems, and agent-oriented programming and multiagent programming. Contact him at a.ricci@unibo.it.

MUNINDAR P. SINGH is a professor of computer science and a co-director of the Science of Security Lablet, North Carolina State University, Raleigh, NC, 27695, USA. His research interests include sociotechnical systems. Contact him at singh@ncsu.edu.