



Cloud Computing

Schahram Dustdar, TU Wien

This theme issue explores compelling new trends and advances in cloud computing.

ABOUT THE AUTHOR

SCHAHRAM DUSTDAR is a full professor of computer science and head of the Distributed Systems Group at TU Wien in Vienna. His research interests include the Internet of Things, cloud computing, and distributed systems. Dustdar received a PhD in business informatics from the University of Linz and a Habilitation in computer science from TU Wien. He is an area editor of cloud computing for *Computer*, an editorial board member of *IEEE Internet Computing*, an IEEE Fellow, an ACM Distinguished Scientist, and a member of Academia Europea. Contact him at dustdar@dsg.tuwien.ac.at.

In the past few years, cloud computing has transformed the IT landscape for both individuals and enterprises—from the way we access, store, and share information to how we communicate, collaborate, and process data. This has led to unprecedented levels of R&D and spawned numerous academic and industry conferences. As the cloud computing paradigm has matured, research has expanded to include virtually all aspects of modern distributed systems including the development, management, deployment, and orchestration of networking, storage, and computing resources.

This theme issue of *Computer*—curated from recently accepted articles submitted for general consideration as research features, perspectives, and computing practices—provides an overview of current efforts to realize the full potential of cloud computing technology and to address some of the ongoing challenges to wide-scale adoption, including concerns over loss of control, vendor lock-in, data security and privacy, and unpredictable cost structures. The four articles selected reflect some of the most important areas of current research.

In “Extending Amdahl’s Law for the Cloud Computing Era,” Fernando Díaz-del-Río, Javier Salmerón-García, and José Luis Sevillano describe the current shift from centralized to distributed application execution, as well as various process migration-related challenges. Such challenges include

managing the overlap between communication and execution times and ensuring efficient and secure data transmission. The authors argue that extending Amdahl’s law to cloud computing would improve our understanding of the process factors involved in offloading applications to the cloud, including time and energy speedups, and thus which apps might or might not be suitable for migration. They also hint at what lies ahead in the cloud computing era.

In “Dimensions for Evaluating Cloud Resource Orchestration Frameworks,” Alireza Khoshkbarforoushha, Meisong Wang, Rajiv Ranjan, Lizhe Wang, Leila Alem, Samee U. Khan, and Boualem Benatallah discuss issues related to CROFs, which are designed to facilitate the migration of in-house applications to cloud environments. Research data to aid software engineers, DevOps managers, and infrastructure administrators in understanding CROFs’ features is lacking. To address this problem, the authors present a set of generic technical dimensions for analyzing and comparing the capabilities of the most prominent CROFs, along with a concise survey of recent research.

“Managing Trust in the Cloud: State of the Art and Research Challenges,” by Talal H. Noor, Quan Z. Sheng, Zakaria Maamar, and Sherali Zeadally, addresses key trust management challenges—including identification, privacy, personalization, integration, security, and scalability—hindering

cloud computing adoption. The authors describe different trust-management perspectives and techniques, identify the trust characteristics of several major cloud service providers, and propose a generic framework with numerous criteria for assessing trust management systems. They also discuss open research challenges revealed by a comprehensive analysis of existing systems.

Finally, in “Energy-Efficient Workload Placement in Enterprise Datacenters,” Quan Zhang and Weisong Shi describe a new way to improve datacenter energy efficiency, an increasingly important research area given significantly rising electricity costs. Previous research has focused on reducing energy consumption in IT equipment, largely ignoring power losses caused by uninterruptible power supplies (UPSs). However, a UPS could account for 15 percent of a large datacenter’s total energy cost. The authors’ proposed technique, which uses a rack-level power model to map workload directly to power dissipation and then optimizes workload allocation to minimize equipment power consumption and power loss from UPSs, could dramatically improve energy efficiency.

hope that you enjoy reading these articles on cutting-edge cloud computing research, and that they inspire novel ideas in this exciting and rapidly evolving field. ■