

Introduction to the Special Section on Advances in Internet-based Collaborative Technologies

SCHAHRAM DUSTDAR, TU Wien, Austria

SURYA NEPAL, CSIRO Data61, Australia

JAMES JOSHI, University of Pittsburgh, USA

Individuals, organizations, and government agencies are increasingly relying on Internet-enabled collaboration among distributed teams of humans, computer applications, and autonomous entities such as robots to develop products and deliver services. Technology trends in areas such as networking, data analytics, and distributed systems have significantly shifted the landscape of Internet-based collaborative tools and services. This particular special issue contains articles describing novel and innovative Internet-based collaborative technologies that leverage emerging technologies and enable seamless collaboration.

CCS Concepts: • **Information systems** → **Information systems applications**;

Additional Key Words and Phrases: Collaboration technology

ACM Reference format:

Schahram Dustdar, Surya Nepal, and James Joshi. 2019. Introduction to the Special Section on Advances in Internet-based Collaborative Technologies. *ACM Trans. Internet Technol.* 19, 3, Article 37e (November 2019), 4 pages.

<https://doi.org/10.1145/3361071>

1 BACKGROUND

The Internet is the single most important innovation of our time that has significant impacts on society. The use of the Internet has grown rapidly since the introduction of the World Wide Web (WWW) in 1989, more specifically from the beginning of this century. The evolution of Internet technology is unprecedented. The Internet now enables interactions and collaborations among electronic documents, physical “things,” and human beings seamlessly within the complex Cyber-Social-Physical ecosystem. This is possible due to the development of the growing number of collaboration-enabling technologies. Individuals, organizations, and government agencies are increasingly relying on Internet-enabled collaboration among distributed teams of humans, computer applications, and/or autonomous robots to develop products and deliver services that would have been impossible to develop without the contributions of collaborators.

The emerging technology trend, more specifically, social computing, Internet-of-Things, cloud computing, mobile computing, edge computing, Big Data, and automation, has significantly shifted the landscape of Internet-based collaborative tools and services. They have also brought several

Authors' addresses: S. Dustdar, TU Wien, Distributed Systems Group, Argentinierstraße 8, A-1040 Wien; S. Nepal, CSIRO Data61, Australia; J. Joshi, University of Pittsburgh, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author

© 2019 Copyright held by the owner/author(s).

1533-5399/2019/11-ART37e

<https://doi.org/10.1145/3361071>

newer challenges. This special issue is aimed toward shedding insights into these challenges and discussing innovative solutions. The special issue contains seven articles covering the following four areas: social network, web search, security, and software engineering. Below, we briefly summarize their key contributions:

2 ARTICLES IN THIS SPECIAL ISSUE

Twitter is the most popular social network platform. A large number of services have been developed using Twitter contents such as in the area of disaster management and syndromic surveillance. One of the challenges faced by such applications/services is how to identify relevant Tweets for each application, or more specifically, unseen events. To address this challenge, Ning et al. have presented an article titled "[Source-aware Crisis-relevant Tweet Identification and Key Information Summarization](#)." In this article, the authors have proposed a new correlative conventional neural network that uses a shared hidden layer to learn effective representations of multi-faceted features such as linguistic, sentimental, and emotional in tweets coming from a variety of sources. The authors have empirically demonstrated the advantages of such approaches.

One of the criticisms of social media is that it has created an echo chamber. People are interacting with like-minded people, and they are not provided with opportunities to hear opposing views on the topics. The echo chamber has been exploited by state or individual to create a more considerable influence on members of online social network communities. Some of the problems reside on the inherent features and architecture of these online social network platforms. To address this challenge, Rafael et al. have presented an article titled "[Universal Social Network Bus: Towards the Federation of Heterogeneous Online Social Network Services](#)." In this article, the authors have presented a universal social network bus, borrowing the idea of service bus from service-oriented computing paradigm, where an individual user of an online social network can connect to people in another online social network, enabling them to make friends beyond their closed circle.

Since the advent of the WWW in 1989, there have been many collaborative platforms, such as online forums, that have been developed and widely used in answering questions from community members. One of the problems with such forums is that they are nowhere close to real time in getting answers to the questions. Many of our queries need real-time answers with real-time data. For example, if a person is planning to go to a stadium to watch tennis and he/she wants to know if there is a ticketing queue. Getting answers to such questions via existing platforms is hard, since they are built on offline, asynchronous protocols. With the broader availability and use of smartphones, such platforms are becoming capable of near real-time response. To address this challenge, Christian et al. have presented an article titled "[CloseUp: Community-driven Live Online Search](#)." In this article, the authors have presented a system called CloseUp, which enables a community-driven live search. The system uses a deep-learning pipeline to analyze the submitted queries and convert them into questions. CloseUp is a standalone mobile application for submitting, browsing, and replying to questions. The authors have evaluated the system using a field study and presented the results.

Many competing organizations build a transient collaboration in many industries, more specifically, in a supply chain environment. For example, two mining companies may share the same railway lines and ports. Besides, they may contribute their company resources for mutual benefits. In such an environment, they need to give access to each others' resources without revealing their business secrets. Each organization has its access control policies. Such transient collaboration requires a standard access control policy that meets the access control requirements of all the participating organizations. Attribute-based Access Control (ABAC) has emerged as a promising option for addressing such requirements. However, translating participating organizations' access policy to a standard access policy requires a significant amount of manual effort. Even then, it

is challenging to get it right. To address this challenge, Das et al. have presented an article titled [“Policy Adaptation in Hierarchical Attribute-based Access Control Systems.”](#) In this article, the authors present a heuristic technique to automatically determine an optimal assignment of attribute values to subjects for enabling the desired accesses to be granted while minimizing the number of ABAC rules used by each subject.

One of the applications that often uses transient collaboration is e-health. Several health service providers may be involved while treating a patient. Each service provider may play multiple roles at different times. One of the challenges of using ABAC in such circumstances is resolving access disputes and making service providers accountable for their actions. To address this challenge, Puseewalage and Oleshchuk have presented an article titled [“An Anonymous Delegatable Attribute-based Credential Scheme for a Collaborative E-health Environment.”](#) In this article, the authors have proposed an efficient anonymous, attribute-based credential scheme capable of provisioning multi-level credential delegations. The proposed scheme is integrated with a mechanism to revoke the anonymity of credentials, for resolving access disputes and making users accountable for their actions.

Industry 4.0 applications, where a collaboration between cyber and physical environment is needed, is driven by the Internet-of-Things (IoT). IoT introduces a huge attack surface. Hence, there is a growing number of threats in IoT-driven cyber-physical systems. The threats are ever increasing and have not shown any signs of subsiding. This demands a need for a threat management framework. To address this challenge, Felemban et al. have presented an article titled [“Threat Management in Data-centric IoT-based Collaborative Systems.”](#) In this article, the authors have proposed a Threat Management System for Data-centric Internet-of-Things-based Collaborative Systems with a specific focus on tampering attacks that target shared databases.

Most of today’s collaborative systems are deployed in the cloud as a software as a service (SaaS). To address emerging threats, such a deployed SaaS needs to be patched regularly. However, regular updates of SaaS need to be done without having an impact on business continuity. How to perform the cost-effective rolling updates of SaaS without any interruptions on business activities remains a challenging problem. To address this challenge, Sun et al. have presented an article titled [“Multi-objective Optimisation of Online Distributed Software Update for DevOps in Clouds.”](#) In this article, the authors have presented an approach for synchronous online distributed software update. This means software versions are upgraded in virtual machine instances even when various failures may occur. The authors have proposed a stochastic model and a novel optimization method and have presented their results through both experiments in Amazon Web Service and simulations.

3 RESEARCH CHALLENGES

Even though the articles in this special issue propose timely and exciting solutions to advances in Internet-based Collaborative Technologies, they certainly do not cover all the aspects. Much more work needs to be done along several directions: security, privacy, and trust; interoperability and automation; and cyber-physical-social systems. We briefly describe a few key challenges below.

Security, Privacy, and Trust: Collaborative applications share data that contains personally identifiable information as well as business secrets. Secure multi-party computation and differential privacy have been recently promoted to tackle this problem. However, security and privacy have been proven to be ephemeral, as evident from many recent data breaches. Emerging technologies in this area include distributed ledger, federated machine learning, and a model-to-data (instead of traditional data-to-model) paradigm. Distributed ledger can help to build accountability of actions without having a central trusted authority. Federated machine learning approaches show promise

in helping build models on aggregated data without releasing the data. Similarly, a model-to-data paradigm can help to keep the data in a secure vault all the time.

Interoperability and Automation: Collaborative applications are built over existing legacy systems for many reasons, the most important one being the cost of developing a new application for a transient purpose. Thus, it is important that collaborative platforms can be integrated with existing systems. This demands a look at the Application Programming Interface–based architectural solutions. Besides, integration of applications should be automated to reduce the massive amount of manual efforts needed. To address this challenge, one possible direction to explore is the orchestration and automation platforms for collaboration, where the integration of different applications is automated, as much as possible, while minimizing manual interventions.

Cyber, Physical, and Social: The evolution of social media in the past decade has given rise to social media–based collaborative platforms. The ubiquitous nature of smartphones and IoT devices have provided a thrust of cyber-physical collaboration platforms. The next generation of collaborative platforms is expected to be cyber-physical-social systems, where the Internet-of-Things meets with the Internet-of-Humans. For example, in a collaborative precision health application, all participants share their data and expertise with the expectation that an individual’s expertise, as well as privacy, are respected and valued for the benefit of society. Ethics and culture become an integrated feature of the applications and will be reflected as the expected value of society.

ACKNOWLEDGMENTS

The authors thank all reviewers who have provided reviews on time and ACM TOIT editorial board and administration, specifically former Editor-in-Chief Prof. Munindar P. Singh and current Editor-in-Chief Prof. Ling Liu. Without their help and support, this special issue would not have been possible.