Preface: Special Issue on Recent Advances in Networks and Distributed Systems

MATHIAS FISCHER, Universität Hamburg, Germany
WINFRIED LAMERSDORF, Universität Hamburg, Germany
JÖRG LIEBEHERR, University of Toronto, Canada
MAX MÜHLHÄUSER, TU Darmstadt, Germany

ACM Reference Format:
Mathias Fischer, Winfried Lamersdorf, Jörg Liebeherr, and Max Mühlhäuser. 2021. Preface: Special Issue on Recent Advances in Networks and Distributed Systems. *J. ACM* 37, 4, Article 111 (August 2021), 3 pages. https://doi.org/10.1145/1122445.1122456

The history of research and development for the Internet has alternated between time periods of grand new visions and time periods concerned with integrating new advances in information and communication technology. As we entered the 21st century, many voices called for a redesign of the Internet architecture. The Internet was said to be *ossified* in the sense that its ability to incorporate major upgrades had been largely lost. Viewed as being in a state beyond repair, the era saw calls for and many efforts on a clean-slate design of the Internet architecture. A decade on – and with the Internet architecture largely unchanged – the appetite for big solutions has largely waned, and been replaced by efforts that seek to adapt networks and distributed systems to challenges posed by the advent of the Internet-of-Things, network softwarization, mission-critical applications, and artificial intelligence:

IoT The Internet-of-Things (IoT) will boost the number of Internet nodes well beyond the billions, which will lead to new traffic patterns (e.g., high-frequency low-volume) and resource requirements.

SDN++ Softwarization is 'conquering' the net. Technologies such as software-defined networking (SDN) and network function virtualization (NFV) will continue to deliver increased flexibility for providing network services. As this softwarization significantly increases the complexity of today's networks, new security problems will emerge.

MCA Mission-critical applications (MCA), which used to be confined to dedicated real-time systems and networks, such as time-sensitive networks (TSN), industrial Ethernet, and so on, are migrating to public and wide area networks. This impacts technologies and protocols, such as ultra-reliable low latency communications (URLLC) in 5G networks. Moreover, delay requirements of real-time applications create a need to move cloud functionality closer to the action scene, e.g., via fog and edge computing.

AI The resurgence of artificial intelligence (AI), evoked by stunning successes of machine learning, boosts the need for computing resources that cannot be embedded in IoT devices and

Authors' addresses: Mathias Fischer, Universität Hamburg, Germany, mathias.fischer@uni-hamburg.de; Winfried Lamersdorf, Universität Hamburg, Germany, winfried.lamersdorf@uni-hamburg.de; Jörg Liebeherr, University of Toronto, Canada, jorg@ece.utoronto.ca; Max Mühlhäuser, TU Darmstadt, Germany, max@informatik.tu-darmstadt.de.

0004-5411/2021/8-ART111 $15.00
https://doi.org/10.1145/1122445.1122456

energy-constrained handhelds like smartphones. This and the need for joint processing requires offloading approaches that shuffle data from the action scene to backend resources.

The objective of this TOIT special issue was to bring together contributions to the above four and additional challenges and shed light on their inter-dependencies. We received in total 12 submissions and 5 high-quality manuscripts were selected for publication in this issue. Our accepted papers reflect the diversity of currently addressed research problems. Four papers are concerned with topic areas envisioned by the guest editors: one paper caters to the growing importance of mission-critical applications (MCA) that require real-time networking, another one addresses QoS and reliability challenges in safety-critical avionic wireless sensor networks, a third paper contributes to edge computing and hence to MCA, but also to other fields where two-tier mobile-cloud settings can be reasonably extended to a three-tier mobile-edge-cloud architecture, the fourth paper proposes a novel algorithm for congestion control for the Tor network and that is based on modern control theory. Paper five addresses additional challenges that were not in direct scope of this SI, namely the resource allocation in data center networks where the accelerated use of tensor processing units (TPUs) and graphical processing units (GPUs) creates new bottlenecks.

1. The paper on real-time networking by Polachan et al. uses the term Tactile Cyber-Physical Systems (TCPS) instead of MCA. It assumes an underlying layer-2 network that supports IEEE 802.1 time sensitive networking (TSN) and a corresponding time-aware shaper that assists in assuring time-bounded packet latencies. The paper caters to an important missing link between such MCA-supporting underlying infrastructure and MCA-demanding applications: a dynamic scheduling protocol that can be easily plugged in by applications, allowing them to turn their MCA requirements into corresponding slot reservations on the TSN level in a decentralized way that reflects concurrent multi-tenant demands.

2. The paper on QoS for wireless sensors networks used in avionics by Shudrenko et al. proposes a cross-layer protocol extension of the protocol *IpV6 over the TSCH mode of IEEE 802.15.4 (6TiSCH)* that the authors call *6TiSCH-CLX* This extension can help to decrease latency and increase reliability in so-called wireless avionics intra-communication (WAIC) that will be used in the future for safety-critical avionic applications, e.g., for smoke detection. Hence, the paper fits the Iot and MCA areas of the CfP. The authors evaluate their approach thoroughly via simulations and an analytical model and the results clearly indicate that even in huge safety-critical communication scenarios the approach outperforms existing solutions in terms of keeping low-latency requirements and packet delivery ratio.

3. The paper on Fog Computing Platforms for Smart City Applications by Silva et al. presents a survey on fog computing platforms for cloud-centric Smart City IoT applications with rather stringent requirements on latency and high-volume data processing as a consequence of high numbers of distributed sensor and many other devices. Motivated by the high potential of the Fog Computing paradigm, this contribution provides a comprehensive overview of what has been investigated on the use of such a paradigm in platforms for Smart Cities and also addresses open future research and development issues. The paper reports on a systematic mapping study and presents a comprehensive view of using the Fog Computing paradigm in Smart City platforms, also including an overview of the current state of research on this topic.

4. The paper by Doepmann et al. addresses a drawback of Tor as most popular anonymization network, namely its lack of adequate congestion control. For that, the authors propose a novel approach for congestion control, called PredicTor, that minimizes the latency in a congested network and achieves max-min fairness. PredicTor uses the concept of distributed model predictive control and allows multiple controllers to make local decisions and to achieve global control goals by communicating their decisions to each other. The authors evaluate their approach via a large-scale simulation study and their results indicate that PredicTor can reduce significantly reduce network latencies and realize fair rate allocations at the same time. Moreover, and according to the authors, their paper is the first of its kind that applies modern control theory in networking for congestion control. Thus, it might induce the development of further (generic) congestion control algorithms.

5. The paper on IDM: Applying Intent-based Cloud Management to Desktop as a Service (DaaS) to Enhance users' QoE by Wu et al. addresses challenges of the growing market of distributed Desktop as a Service (DaaS) applications. For that it proposes a so-called Intent-driven Management (IDM) framework for a respective cloud data management in data center networks. It specifically allows for autonomous decisions on cloud resource configurations to reduce design effects, time and human costs for such applications. The paper reports in detail on an analysis of five main challenges and restrictions of applying this machine learning-based DaaS resource allocation framework to practical services, including identification of the resource design objective, quantification of DaaS QoE, low data availability, model designing, and low resource variations in the log data. The final goal of this framework is to enhance the Quality of Experience (QoE) of DaaS service usage as one of the most critical success factors for DaaS.

We thank all authors for submitting to this SI and all reviewers for their thorough and helpful reviews. We also appreciate the support and help from the editorial staff and the Editor-in-Chief, Ling Liu.