Bridging the Gap between Industry and Academia for Networking Research


Bo Han


Jiasi Chen


Tian Guo


Sung-Ju Lee


Viswanathan Swaminathan

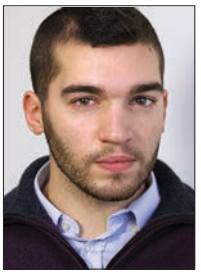
Stylianos I. Venieris

It has been widely acknowledged that there is a gap between the networking research conducted in industry and that performed by professors and students in universities. This gap is partially caused by the different goals of the two parties. Researchers in industry may focus more on the technology transfer side for offering better services and thus improving revenue for their companies. Academic researchers may tend to focus more on the intellectual challenges of either theoretical or practical problems, and advance the state of the art with novel algorithms, protocols, and architectures, without worrying about issues such as commercialization. Moreover, usually, there is limited access for academia to understand and appreciate the pain points of operating production networks and offering networking service at scale, and limited access to the non-public resources owned by private enterprises for more effective research (e.g., operational data of a network) due to the lack of broad collaborations between industry and academia. On the other hand, we have witnessed multiple successful cases where academia, government, and private enterprise worked together to commercialize some extraordinary ideas into real products. Our Internet may be arguably the most victorious story along this line, which was originally sponsored by the U.S. Department of Defense as ARPANET to connect a few computers at universities.

The goal of this Special Issue is to bridge the above-mentioned gap by promoting and maximizing the outcome of research collaborations between industry and academia, for example, through summer internships and funded research by companies conducted at universities. It is motivated by the fact that the networking industry has been actively bridging such gaps through numerous collaborative research and open source collaborations for improving the research-technology transfer. Thus, this Special Issue mainly targets research activities jointly contributed by academia and industry, which should include concrete evidence of the research collaboration. We welcomed technical papers addressing practical challenges faced by industry that are jointly tackled with complementary contributions from both sides, and papers illustrating use cases behind the technical issues that are under consideration by the industry. We encouraged authors who submitted papers to this Special Issue to carry out performance evaluations using platforms and infrastructures provided by industry. We especially encouraged submissions whose lead author is a student (either graduate or undergraduate) to disseminate and advertise experimental results of technical achievements during summer internships or through collaborations with industrial researchers. We also considered position papers from joint academia and industry authors that illustrate the existing gap and possible solutions to bridge it.

We received 18 high-quality submissions, of which 4 articles were eventually accepted after a rigorous review process. These articles cover several different aspects of networking research. In the first article, "Link layer Connectivity as a Service for Ad Hoc Microservice Platforms," authors from Universidad Carlos III de Madrid and Telefónica propose an enhancement to make the microservice platforms suitable to support telecommunication and vertical services beyond the network edge. Their solution to the current shortcomings provides a programmable data plane that enables the establishment of on-demand link layer connectivity between microservices on ad hoc networks. This solution has the flexibility to execute different algorithms to build traffic paths between microservices, as well as to react against temporary link breakdowns.

"Industry-Academia Research toward Future Network Intelligence: The NG-CDI Prosperity Partnership" presents the Next Generation Converged Digital Infrastructure (NG-CDI), an industry-academia partnership between several universities in the United Kingdom and British Telecom, to build and operate a future-proof network infrastructure and its autonomic management for fast and efficient service innovation. It demonstrates the benefits of a collaborative interdisciplinary approach for addressing networking challenges based on real-world problems and how academia and industry can work closely together to deliver positive impacts on business.

In the next article, "Spatial Models for Networks on Roads: Bridging the Gap between Industry and Academia," authors from Qualcomm and Virginia Tech look at a spatial model called the Poisson line Cox process (PLCP) for vehicular networks from an industry perspective. The PLCP, compared to the conventional regular grid models, can capture the randomness in both the road layouts and the vehicle's position. The article discusses some of the key performance metrics that are useful to quantify the reliability and latency of vehicular networks using the PLCP. Additionally, the article highlights the potential applications of the PLCP to transportation networks and Industrial Internet of Things (IoT) with its path distance characteristics.

In the last article, "Practically Deploying Multiple Vertical Services into 5G Networks with Network Slicing," authors from Telenor Research, University of Patras, and Simula Metropolitan Center for Digital Engineering focus on the scalable deployment of specialized vertical services (VSs) through 5G network slicing. Despite the broad use of slicing in 4G infrastructure, 5G networks pose a set of new challenges in terms of performance, functionality, and scalability requirements. In this context, the article presents a joint industry-academia effort to develop a principled framework for successfully deploying end-to-end 5G VSs, together with an underlying platform for empirical experi-

mentation across diverse use cases. Altogether, this work identifies and proposes solutions to important deployment issues and presents key takeaways for building next-generation VSs over 5G networks.

In conclusion, the articles presented in this Special Issue demonstrate the collaborative networking research between industry and academia and call for more active collaborations to bridge the gap. We would like to thank all the authors who submitted their original research work and all the reviewers for their efforts in providing timely and constructive reviews that significantly improve the quality of the articles. In addition, we are grateful to Dr. Chonggang Wang, Editor-in-Chief, and the *IEEE Network* staff for providing this opportunity and their invaluable support during the preparation of this Special Issue. Finally, we sincerely hope that readers will find the articles in this Special Issue interesting and helpful.

Biographies

Bo Han (bohan@gmu.edu) received his Ph.D. degree from the University of Maryland, his M.Phil. degree from the City University of Hong Kong, and his Bachelor's degree from Tsinghua University, all in computer science. He is an associate professor in the Department of Computer Science, George Mason University. Before that, he was a principal inventive scientist at AT&T Labs-Research. His research interests are in the areas of networked systems, mobile computing, and wireless networking. His recent research focuses on immersive video streaming, augmented, virtual, and mixed reality, and their applications in various domains such as education, training, healthcare, data visualization and analytics, and digital twin. He has published 70+ papers in prestigious international journals including *IEEE/ACM Transactions on Networking* and *IEEE Transactions on Mobile Computing* and conference proceedings such as ACM MobiCom, USENIX NSDI, and IEEE INFOCOM. He is the co-inventor of 40+ U.S. patents. He received the best paper award from ACM CoNEXT 2016 and the DASH-IF Excellence Award from ACM MMSys 2019. He serves on the Editorial Board of *IEEE Network* and is the co-founder of the ACM SIGCOMM Workshop on 5G Measurements, Modeling, and Use Cases (5G-MeMU).

Jiasi Chen is an associate professor in the Department of Computer Science & Engineering at the University of California, Riverside. She received her Ph.D. from Princeton University and her B.S. from Columbia University. Her research interests include edge computing, wireless/mobile networks, and multimedia networking, with a recent focus on machine learning and AR/VR. She is a recipient of the NSF CAREER award and the Hellman Fellowship.

Tian Guo is currently an assistant professor of Computer Science at Worcester Polytechnic Institute. She earned her Ph.D. degree from the University of Massachusetts Amherst. She is an empirical system researcher by training. She is interested in designing systems mechanisms and policies to handle trade-offs in cost, performance, and efficiency for emerging applications. Specifically, she has worked on projects at the intersection of systems and machine learning, secure machine learning, cloud/edge resource management, big data frameworks, deep learning inference, distributed training, neural architecture search, and AR/VR. Her recent work has a strong focus on improving system support for deep learning and on the practical applications of deep learning in AR/VR. She is the recipient of the National Science Foundation CRII award and MMSys 2020 best paper award.

Sung-Ju Lee [F] received his Ph.D. in computer science from the University of California, Los Angeles in 2000. He spent 12.5 years at Hewlett-Packard Company as a principal research scientist and distinguished mobility architect. He was then a principal member of technical staff at the CTO Office of Narus, Inc. (now part of Symantec). In 2015, he joined Kores Advanced Institute of Science and Technology (KAIST), where he is a KAIST Endowed Chair Professor. His publications are well cited, reaching a total of over 15,000 citations, and his h-index is 53 according to Google Scholar. In addition, he has 51 granted U.S. patents. He has published in top conference venues such as MobiCom, MobiSys, CHI, INFOCOM, CoNext, SenSys, UbiComp, UIST, CSCW, ICDCS, NDSS, and IMC. He won the HP CEO Innovation Award in 2010, which recognizes the people behind the most innovative products that HP has brought to market. He is also the winner of the test-of-time paper award at ACM WiNTECH 2016, the best paper awards at IEEE ICDCS 2015 and ACM CSCW 2021, and the methods recognition award at ACM CSCW 2021. In addition, he received the Technology Innovations Award from KAIST. He was the General Chair of ACM MobiCom 2014 and co-TPC Chair of IEEE INFOCOM 2016 and ACM MobiCom 2021. He is xan ACM Distinguished Scientist.

Viswanathan Swaminathan is a senior principal scientist at Adobe Research working at the intersection of insights from behavioral data and multimedia content. His areas of research include next-generation video and immersive experiences, video streaming, and, in general, data-driven content and marketing technologies. His research work has substantially influenced various technologies in Adobe's video delivery, advertisement, and recommendations products including the guts of HTTP Dynamic Streaming, which won the Best Streaming Innovation of 2011 Streaming Media Readers Choice Award. He has received several awards including four best papers, 2017 Distinguished alumnus from Utah State University ECE Department, and three ISO certificates of appreciation for MPEG Standards editing contributions. Prior to joining Adobe, he was a senior researcher at Sun Labs. He received his M.S. and Ph.D. in electrical engineering from Utah State University. He received his B.E degree in ECE from the College of Engineering, Guindy, Anna University, Chennai, India. He has authored several papers, articles, RFCs, and book chapters, has about 75+ issued patents, and volunteers in organizing IEEE and ACM conferences.

Stylianos I. Venieris is currently a senior researcher at Samsung AI, Cambridge, United Kingdom, where he leads the Distributed AI group. He received his Ph.D. in reconfigurable computing and deep learning from Imperial College London in 2018 and his M.Eng. degree in electrical and electronic engineering from Imperial College London in 2014. His research interests include principled methodologies for the mapping of deep learning algorithms on distributed and mobile platforms, the design of novel end-to-end deep learning systems that robustly meet multi-objective performance requirements, and the design of next-generation hardware accelerators for the high-performance, energy-efficient deployment of deep neural networks.