Special Issue on The ACM SIGMETRICS Workshop on Measurements for Self-Driving Networks

Arpit Gupta
University of California
arpitgupta@ucsb.edu

Ramakrishnan Durairajan
University of Oregon
ram@cs.uoregon.edu

Walter Willinger
NIKSUN, Inc.
wwillinger@niksun.com

The design and implementation of autonomous or "self-driving networks" represent some of today's most significant challenges in networking research. The vision for these networks is that they will be able to make management and control decisions in real time, typically without human intervention. Recent technological advancements, like SDN and 5G networks, along with scientific innovations such as XAI and transformers, have paved the way for this vision. Key innovations include: (1) fully programmable, protocol-independent data planes and the languages to program them; (2) scalable platforms capable of processing distributed streaming data, bolstered by the latest tools and software for data analysis and machine learning (ML).

A particularly promising development is the fusion of programmable control capabilities in the data plane with advanced ML-based inference techniques. This combination offers unprecedented opportunities for querying the network's state on a vast scale, providing the essential data for the many network management and control tasks that self-driving networks must autonomously perform.

However, the path toward realizing self-driving networks is strewn with obstacles. Practical, deployable system designs that are scalable and robust are scarce. Likewise, many ML-based inference tools available today are not production-ready; they typically lack generalizability, trustworthiness, or assurance of system safety. Realizing the vision of practical self-driving networks will require scalable system designs that employ closed-loop feedback at multiple levels to ensure their robustness with respect to the uncertainties of their environments. Moreover, a shift in perspective will be necessary for developing ML-based inference solutions. The success of the learning models that drive these solutions will have to be gauged by their explainability, trustworthiness, and safety rather than just traditional concerns like accuracy.

To assess the current level of interest and activity in this area, we organized the 1st Workshop on Measurements for Self-Driving Networks that took place in Orlando, Florida, USA, on June 19, 2023. The workshop was sponsored by NSF, organized by ACM SIGMETRICS, and held in conjunction with ACM SIGMETRICS 2023/FCRC 2023. This workshop served as a platform for researchers to present and discuss their latest research on technologies poised to make practical, deployable self-driving networks a reality. We sought contributions from experts in fields such as networking, applied as well as theoretical machine learning, network security, control theory, distributed systems, computer architecture, and data science, all united by their enthusiasm to realize the vision of self-driving networks.

The workshop featured presentations from invited speakers that represented 11 universities and included a diverse mix of senior and early-career researchers, as well as graduate students. All speakers were invited to submit a three-page paper on the topic of their presentation. The papers in this Special Issue are a testimony to the exciting ongoing developments in this area of research and address topics such as traffic monitoring, approximate querying, and decision making at data plane speeds; deployability and engineering challenges for self-driving networks; explainable network controllers; learning-assisted QoE enhancements; and the need for a paradigm shift in how ML-based solutions for networking problems in general and self-driving networks in particular ought to be developed and evaluated in the future so that they can be deployed and used in practice.

This collection of papers also highlights four fundamental limitations faced by researchers pursuing self-driving networks. These include: (1) a lack of capabilities to label network datasets at scale; (2) an urgent need for frameworks that facilitate privacy-preserving collaboration among researchers; (3) difficulties in developing provably generalizable ML artifacts; and (4) uncertainties about creating feasible pathways for safely road-testing ML models.

Technical discussions between speakers and participants during the workshop hinted at a potential path forward. In particular, one possible way to overcome these limitations involves building a community-wide infrastructure designed specifically to: (1) facilitate flexible, high-quality data generation and collection efforts that can be easily replicated across different networks; (2) offer an innovative framework for collaborative and privacy-preserving knowledge sharing, such as labeling functions, model specifications, and data features; (3) bolster a principled approach to developing generalizable learning models for networking problems; and (4) establish a strategy for deploying ML-based solutions in production networks.

We extend our gratitude to all the speakers, co-authors, and attendees who contributed with their presentations and actively participated in the workshop. We would also like to acknowledge the support of ACM SIGMETRICS and, in particular, the workshop co-chairs Leana Golubchik and Daniel Sadoc, and the PER editor Zhenhua Liu for their continuous guidance and assistance in producing this special issue of PER.

Workshop on Measurements for Self-Driving Networks Orlando, FL, USA