Guest Editorial
Introduction to the Special Section on Next-Generation Traffic Measurement With Network-Wide Perspective and Artificial Intelligence

TRAFFIC measurement is the bedrock of the next-generation network systems. While it plays a crucial role in bringing fundamental data and support to core network functions, it also confronts the challenge of meeting the diverse demands of new network traffic characteristics and emerging applications. The network-wide measurement has received more and more attention. Given that big network data is distributed in nature, it is essential to aggregate the views of multiple measurement points to build a network-wide perception of traffic. Another latest trend involves artificial intelligence technologies that allow seamless aggregation of multifaceted network traffic data to advance traffic data analysis and support related applications. Nonetheless, a gap remains in existing methodologies, which often fail to fully address the diverse demands of network traffic measurement in this evolving landscape.

This special section focuses on the network-wide and AI-powered traffic measurement and related applications. Thanks to the extensive efforts of the reviewers and the great support from the Editor-in-Chief, Dr. Jianwei Huang, we were able to accept 18 contributed articles covering several important topics, from the sketch-based passive measurement [A1], [A2], [A3], [A4], [A5] to the INT-based active measurement [A6], [A7], to task offload and privacy protection regarding network-wide measurement [A8], [A9], [A10], to AI-powered traffic analysis [A11], [A12], [A13], [A14] and other related applications [A15], [A16], [A17], [A18]. A brief review follows:

Li et al., in "CS-Sketch: Compressive Sensing Enhanced Sketch for Full Traffic Measurement" [A1], proposed a lightweight framework CS-Sketch to perform accurate measurements of the elephant and mouse flows. Based on compressive sensing, CS-Sketch views the flow size vector of all flows as a signal vector and compresses it to a measurement vector through a sensing matrix. The measurement system consists of multiple switch nodes and a data center side, where the switch node constructs a sparse 0–1 sensing matrix to track the flow size with lightweight update operations, and the data center side accurately estimates the flow size of each flow through the measurement vector at high speed.

Ma et al., in their paper "From CountMin to Super kJoin Sketches for Flow Spread Estimation" [A2], pointed out that most per-flow spread estimation inherits a similar design from a flow size estimation solution - CountMin that conducts minimum estimation, which indeed restricts the estimation accuracy. The authors further exploit the internal structure of plug-ins and replace the min operation with new position-aware operations to achieve better estimation accuracy. After that, super kJoin sketches are proposed to achieve more accurate removal of inter-flow noise and reduce the worst-case errors.

In "Unbiased Real-time Traffic Sketching" [A3], Wu et al. proposed Unbiased Cleaning Sketch (UC sketch), enabling unbiased per-flow size measurements within a sliding window model. First, the paper partitions the time window into d time segments and deploys d+1 counters within each estimator in the UC sketch. These counters are allocated to record the flow size in the preceding d time segments and the current time segment, respectively. Then, the authors introduce the linear scaling technique, facilitating median-unbiased flow size estimation and significantly reducing the estimation variance of the UC sketch. Lastly, they employ the Column Randomizing technique to mitigate errors resulting from hash collisions.

In "Micro-burst Aware ECN in Multi-Queue Data Centers: Algorithm and Implementation" [A4], Kang et al. proposed Micro-Burst aware Explicit Congestion Notification (MBECN+) to cope with the mismarking problem raised by micro-burst traffic. MBECN+ first adopts an ECN threshold setting mechanism that can find appropriate lower bounds for each queue based on steady-state analysis. Next, Dequeue with Slope ECN Marking (DSEM), an ECN marking scheme, is designed to eliminate the influence of backlog. It uses two K values to determine whether the growth of backlog causes congestion and avoids spurious congestion caused by micro-burst flows.

Xiao et al., in "Multi-resolution Odd Sketch for Mining Extended Jaccard Similarity of Dynamic Streaming Sets" [A5], proposed a multi-resolution odd sketch (MROS) that combines the odd sketch with a multi-resolution sampling structure to compress a dynamic streaming set. The MROS summary allows for both the insertion and deletion of elements, and it is mergeable for estimating the Jaccard similarity of two users' itemsets with equal memory size and the same hash functions. After that, the authors extend their algorithm to the scenario with more users by using the set expression cardinalities as unknown variables to establish a linear equation system.

Date of current version 30 April 2024.
Digital Object Identifier 10.1109/TNSE.2024.3389428

Liu et al., in “SFANTL: A SRv6-based flexible and active network telemetry scheme in programming data plane” [A6], proposed an active network telemetry scheme based on P4, introducing the SRv6 mechanism with custom probes. By dynamically generating probe packets with both INT and SRv6 headers, SFANTL enables the flexible specification of monitoring targets, ranges, and information categories, effectively covering telemetry requirements with minimal path cost and bandwidth usage. The evaluation results demonstrate that SFANTL can provide efficient and accurate network telemetry and suggest future directions to explore the sending rate of probes, optimizing the balance between real-time monitoring and resource consumption.

Zhang et al., in “An AI-Augmented Kalman Filter Approach to Monitoring Network Traffic Matrix” [A7], proposed an AI-augmented Kalman filter to address the problem of network traffic matrix estimation. The filter leverages the ConvGRU recurrent neural network to learn and model the underlying characteristics of the traffic matrix, capturing features such as spatiotemporal correlations and traffic dynamics to establish mapping relationships for state transition and state difference reasoning. Subsequently, indirect traffic measurements (i.e., link loads) are collected to estimate the missing knowledge of the traffic matrix. The proposed approach effectively combines the advantages of both direct and indirect measurements to achieve superior traffic matrix estimation accuracy.

Fan et al., in “Joint Optimization of Measurement Point Intelligent Selection and End-to-End Network Traffic Calculation in Datacenters” [A8], proposed an innovative end-to-end traffic inference algorithm named LLS-TC for data center networks, which leverages network tomography and Simple Network Management Protocol (SNMP) data to rapidly and accurately calculate network traffic. LLS-TC addresses the challenges of high software and hardware costs by employing an intelligent selection of measurement points based on node weighting and a network tomography method tailored for cloud computing. By modeling the algorithm problem into a linear state space model, LLS-TC can significantly improve the accuracy of traffic calculation while maintaining computational efficiency.

To avoid measurement redundancy, Yao et al., in “Distributed Strategy for Collaborative Traffic Measurement in a Multi-controller SDN” [A9], proposed a novel distributed iterative strategy for collaborative traffic measurement in a multi-controller SDN. The strategy enables each controller to get the proper sampling probability of any flow at any switch by local flow information and a little message exchange with neighboring controller(s). Moreover, a dedicated packet processing mechanism is designed in the data plane to support the multi-controller scenarios. The extensive simulations demonstrate that the proposed strategy achieves a near-optimal performance in terms of measurement load and communication load.

In “Differentially Private Top-k Flows Estimation Mechanism in Network Traffic” [A10], Zhu et al. considered the privacy leakage problem in aggregating traffic data to find top-k flows and proposed a local differential privacy (LDP) method to protect sensitive information. Their method utilizes the sparsity property of network flow; namely, the number of flows met by all clients is much less than the flow domain size. By representing the flows of each client as a sparse vector, they first present a high-utility LDP traffic aggregation scheme based on HyperLogLog. Further, they utilize multi-iteration approximation to reduce the computation cost to find the top-k flows.

To leverage the advantages of serverless platforms, Zhao et al., in “faaShark: an End-to-End Network Traffic Analysis System atop Serverless Computing” [A11], proposed faaShark, an end-to-end traffic network analysis system based on the serverless computing platform. FaaShark employed distributed training to leverage the lightweight virtualized runtime and fair scalability of serverless platforms over PaaS (Platform as a Service) platforms and provides cloud-native, cost-effective, and convenient training and deployment services for network traffic analysis models. Additionally, they proposed a gradient-based cold start optimization algorithm to minimize cold start hit rates when serving pre-trained models to handle network analysis requests.

Wang et al. adopted a federated learning framework to capture the unique properties of real-time traffic with higher efficiency in “FedStream: A Federated Learning Framework on Heterogeneous Streaming Data for Next-generation Traffic Analysis” [A12]. The authors consider the heterogeneity in data distribution and arrival patterns of traffic and propose H_strSAGA for local optimization and FedStream to tackle the dual heterogeneity challenge. Also, they introduce an asynchronous aggregation algorithm to deal with increasing device heterogeneity. This work mitigates the negative impact of dual heterogeneity on global model performance and enhances straggler tolerance.

Yang et al., in “Transforms-based Bayesian Tensor Completion Method for Network Traffic Measurement Data Recovery” [A13], introduced the transforms-based Bayesian tensor completion (TBTC) method to infer network-wide traffic from incomplete data. They represent heterogeneous traffic measurements as observation tensors, preserving data structure and correlation. By diagonalizing the block matrix, they devise an efficient variational Bayesian inference algorithm in arbitrary invertible linear transforms domain, reducing complexity and enhancing missing data recovery. TBTC is adaptable to various linear transforms, enhancing network traffic data quality for monitoring, analysis, and intrusion detection, ensuring network traffic data integrity.

Existing tensor-based anomaly pursuit methods are hindered by overly ideological assumptions, failing to address structured anomalies and sparse corruption. In “Structured-Anomaly Pursuit of Network Traffic via Corruption-Robust Low-Rank Tensor Decomposition” [A14], Zeng et al. proposed the Corruption-Robust Low-Rank Tensor Decomposition (Cr-LTD) method. This novel approach incorporates $l_{\{2,1\}}$ and l_1-norms to characterize structured anomalies and robustness to sparse corruption effectively. Cr-LTD introduces the tensor tubal rank to capture the low-rank property of network traffic and employs a novel tensor nuclear norm to relax it, circumventing NP-hard problems. By leveraging the alternating direction method of multipliers and acceleration mechanism, Cr-LTD achieves an efficient structured-anomaly pursuit of network traffic.

Tang et al. introduced FTOP, an effective system to counter flow table overflow issues, in “FTOP: An Efficient Flow Table Overflow Preventing System for Switches in SDN” [A15]. It tackles Low-rate Flow Table Overflow (LFTO) attacks and Flash Crowds (FCs) with an eviction-based online approach. FTOP comprises Predictor, Detector, Mitigator, and Preventer modules, minimizing evictions to prevent overflow. It integrates Kalman filtering for flow count prediction and Random Forest classifiers for attack detection. This system promptly responds to overflow risks, enhancing SDN network security and reliability by mitigating associated risks. FTOP effectively safeguards SDN networks from attack impacts by preventing flow table overflow.

Xiao et al. examined traffic-aware resource allocation for UAV communications using RSMA in "Traffic-Aware Energy-Efficient Resource Allocation for RSMA Based UAV Communications" [A16]. They maximize UAV energy efficiency by jointly optimizing UAV deployment, beamforming, rate allocation, and subcarrier allocation based on user needs. To address non-convexity, they propose a joint optimization approach, including a heuristic UAV 3D location method and RSMA parameter optimization using the successive convex approximation method. Moreover, this paper formulates the subcarrier allocation problem as a many-to-one and two-sided matching game.

Hu et al. proposed a privacy-preserving few-shot traffic detection (PFTD) method against Advanced Persistent Threats (APT) in "Privacy-preserving Few-shot Traffic Detection against Advanced Persistent Threats via Federated Meta Learning" [A17]. They introduced federated meta-learning (FML) into the algorithm design and treated APT detection as a model generalization optimization process, transferring learned knowledge to identify local unknown samples. On the client side, Model-Agnostic Meta-Learning (MAML) enables personalized adjustments for quick adaptation and accurate classification. On the server side, aggregation of global knowledge through federated learning ensures information security and privacy of edge devices.

Santos et al., in "Random Access based on LSTM for Mixed Traffic IoT Networks" [A18], focused on the co-existence of massive machine-type communication (mMTC) packets and mission-critical ultra-reliable low latency (URLLC) packets in IoT traffic, which require different network coordination and resource allocation strategies. The authors proposed a hybrid random access protocol designed for mixed URLLC-mMTC scenarios, utilizing a long short-term memory neural network for traffic forecasting. Also, a resource slicing scheme is developed to assign channels to the upcoming traffic at every frame.

In summary, the collected articles offer innovative traffic measurement scenarios and shed light on the underlying principles of traffic measurement design for next-generation networks. We hope this special section will trigger more future work in the emerging area.

HE HUANG, *Guest Editor*
School of Computer Science and Technology
Soochow University
Suzhou 215123, China
huangh@suda.edu.cn

SHIGANG CHEN, *Guest Editor*
Department of Computer and Information
of Science and Engineering
University of Florida
Gainesville, FL 32611 USA
sgchen@cise.ufl.edu

RAN BEN BASAT, *Guest Editor*
Department of Computer Science
University College London
London, NW1 2AE U.K.
r.benbasat@cs.ucl.ac.uk

HAIPENG DAI, *Guest Editor*
State Key Laboratory for Novel Software Technology
Nanjing University
Nanjing 210023, China
haipengdai@nju.edu.cn

AMIRHOSEIN TAHERKORDI
Department of Informatics
University of Oslo
Oslo 0316, Norway
amirhost@ifi.uio.no

JUN XU, *Guest Editor*
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280 USA
jx@cc.gatech.edu

APPENDIX
RELATED ARTICLES

[A1] L. Li, K. Xie, S. Pei, J. Wen, W. Liang, and G. Xie, "CS-sketch: Compressive sensing enhanced sketch for full traffic measurement," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2338–2352, May/Jun. 2024.

[A2] C. Ma, O. O. Odegbile, D. Melissourgos, H. Wang, and S. Chen, "From CountMin to super kJoin sketches for flow spread estimation," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2353–2370, May/Jun. 2024.

[A3] Y. Wu et al., "Unbiased real-time traffic sketching," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2371–2383, May/Jun. 2024.

[A4] J. Zhang et al., "Micro-burst aware ECN in multi-queue data centers: Algorithm and implementation," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2384–2398, May/Jun. 2024.

[A5] Q. Xiao, S. Yang, P. Li, K. Li, and L. Wen, "Multi-resolution odd sketch for mining extended jaccard similarity of dynamic streaming sets," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2399–2414, May/Jun. 2024.

[A6] Y. Liu, Y. Xia, W. Zhang, W. Jia, and J. Wu, "SFANT: A SRv6-based flexible and active network telemetry scheme in programming Data plane," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2415–2425, May/Jun. 2024.

[A7] Q. Zhang and S. Pan, "An AI-augmented Kalman filter approach to monitoring network traffic matrix," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2426–2437, May/Jun. 2024.

[A8] W. Fan, F. Xiao, L. Han, X. He, and J. Wang, "Joint optimization of measurement point intelligent selection and end-to-end network traffic calculation in datacenters," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2438–2449, May/Jun. 2024.

[A9] D. Yao, Q. Ma, H. Wang, M. Chen, and H. Xu, "Distributed strategy for collaborative traffic measurement in a multi-controller SDN," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2450–2461, May/Jun. 2024.

[A10] Y. Zhu, Q. Song, and Y. Luo, "Differentially private top-k flows estimation mechanism in network traffic," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2462–2472, May/Jun. 2024.

[A11] H. Zhao et al., "faaShark: An end-to-end network traffic analysis system atop Serverless Computing," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2473–2484, May/Jun. 2024.

[A12] N. Wang, X. Li, Z. Guan, and S. Yuan, "FedStream: A federated learning framework on heterogeneous streaming data for next-generation traffic analysis," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2485–2496, May/Jun. 2024.

[A13] Z. Yang, L. T. Yang, L. Yi, X. Deng, C. Zhu, and Y. Ruan, "Transforms-based Bayesian tensor completion method for network traffic measurement data recovery," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2497–2509, May/Jun. 2024.

[A14] J. Zeng, L. T. Yang, C. Wang, Y. Ruan, and C. Zhu, "Structured-anomaly pursuit of network traffic via corruption-robust low-rank tensor decomposition," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2510–2523, May/Jun. 2024.

[A15] D. Tang, Z. Zheng, K. Li, C. Yin, W. Liang, and J. Zhang, "FTOP: An efficient flow table overflow preventing system for switches in SDN," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2524–2536, May/Jun. 2024.

[A16] M. Xiao, H. Cui, D. Huang, Z. Zhao, X. Cao, and D. O. Wu, "Traffic-aware energy-efficient resource allocation for RSMA based UAV communications," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2537–2548, May/Jun. 2024.

[A17] Y. Hu, J. Wu, G. Li, J. Li, and J. Cheng, "Privacy-preserving few-shot traffic detection against advanced persistent threats via Federated Meta Learning," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2549–2560, May/Jun. 2024.

[A18] H. L. d. Santos, J. C. Marinello, C. M. Panazio, and T. Abrão, "Random access based on LSTM for mixed traffic IoT networks," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2561–2573, May/Jun. 2024.

He Huang (Senior Member, IEEE) received the Ph.D. degree from the Department of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His advisor is Prof. Guoliang Chen. From 2019 to 2020, he was a Visiting Research Scholar with Florida University, Gainesville, FL, USA. He is currently a Professor and an Associate Dean of the School of Computer Science and Technology, Soochow University, Suzhou, China, where he also directs the Institutes of Network Science and Engineering. He has authored or coauthored more than 100 peer-reviewed journal/conference papers and has 13 Chinese patents. His research interests include traffic measurement for high-speed Internet, software defined networks, mobile computing, privacy and security, cyber-physical systems, and algorithms. He was the recipient of the Changjiang (Yangtze River) Youth Scholar Award in 2021. He is a member of the Association for Computing Machinery (ACM). His students and he was the recipient of three best paper awards (Bigcom 2016, IEEE MSN 2018, IEEE Bigcom 2018). He is currently an Associate Editor for two international journals, including INTERNET OF THINGS and *Cyber-Physical Systems Journal*, and *Frontiers in the Internet of Things*. He was with various capacities (publicity chair, publication chair, and technical program committee) in a number of conferences, such as, Publicity Co-Chair of ACM MobiHoc-MSCC 2016, Publication Chair of the Second IEEE INFOCOM Workshop on Networking Algorithms (WNA) in conjunction with IEEE INFOCOM 2022, TPC member of IEEE INFOCOM, IEEE MASS, IEEE ICC, and IEEE Globecom.

Shigang Chen (Fellow, IEEE) received the B.S. degree in computer science from the University of Science and Technology of China, Hefei, China, in 1993, and the M.S. and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign (UIUC), Champaign–Urbana, IL, USA, in 1996 and 1999, respectively. After graduating from UIUC, he was with Cisco Systems on network security for three years and helped start a network security company, Protego Networks. He joined the University of Florida, Gainesville, FL, USA, as an Assistant Professor in 2002, and was promoted to an Associate Professor in 2008 and to a Professor in 2013. He has authored or coauthored more than 200 peer-reviewed journal/conference papers and has 13 U.S. patents. He held the University of Florida Research Foundation Professorship during 2017–2020 and University of Florida Term Professorship during 2017–2020. He was the recipient of IEEE Communications Society Best Tutorial Paper Award in 1999, NSF CAREER Award in 2007, and Cisco University Research Award in 2007, 2012. He is an ACM Distinguished Member. He is currently an Associate Editor for IEEE TRANSACTIONS ON MOBILE COMPUTING and was an Editor of IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY , and CN. He was also with various capacities in a number of conferences, such as, Area TPC Chair for IEEE INFOCOM 2014–2022, General Chair of IEEE BIGCOM 2019, and Co-Chair of NSF Workshop on Edge Networking 2018.

Ran Ben Basat received the B.Sc., M.Sc., and Ph.D. degrees from the Computer Science Department, Technion, Haifa, Israel. He was a Postdoctoral Fellow with Harvard University, Cambridge, MA, USA, working with Minlan Yu and Michael Mitzenmacher. He is currently an Assistant Professor with the Computer Science Department, University College London, London, U.K. Before Harvard, he was the Ph.D. Candidate, under the supervision of Roy Friedman. His research focuses on algorithms for networking systems and measurement. He is currently a Reviewer of IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, IEEE/ACM TRANSACTIONS ON NETWORKING , *ACM Transactions on Storage* (Distinguished Reviewer), IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING, *Journal of Parallel and Distributed Computing*, IEEE COMMUNICATION LETTERS, AND IEEE NETWORKING LETTERS. He was also a TPC Member of a number of conferences, such as, ACM SIGMETRICS, APNET, ACM CoNEXT, ACM/IEEE ANCS, USENIX ATC, IEEE INFOCOM Workshop on Networking Algorithms, ACM SOSR, and EuroP4.

Haipeng Dai (Senior Member, IEEE) received the Ph.D. degree in computer science and technology from Nanjing University, Nanjing, China, in 2014. He is currently an Associate Professor with the Department of Computer Science and Technology, Nanjing University. His research interests include data mining, mobile computing, wireless power transfer technology, wireless sensor networks, and wireless networks , convex optimization, combinatorial optimization, probability analysis, coding theory, network security, game theory, and computational geometry. He is the Young Chang Jiang Scholar. He was the recipient of the Jiangsu Computer Society Youth Science and Technology Award in 2020, and IEEE SAGC 2021 Best Paper Award in 2021. He is currently the Section Editor of *Sensors*, an Associate Editor for *Hans Journal of Data Mining*, and an Associate Editor for *Frontiers in Communications and Networks*. He was also with various capacities in a number of conferences, such as, TPC Vice-Chair for IEEE HPCC 2021, TPC Chair of iThing 2021, Track Chair of IEEE ICPADS 2021, and TPC member of IEEE INFOCOM, VLDB, IJCAI, IEEE ICNP, and IEEE ICDCS.

Amirhosein Taherkordi (Member, IEEE) received the Ph.D. degree from the Informatics Department, University of Oslo (UiO), Oslo, Norway, in 2011. After completing the Ph.D. studies, he joined Sonitor Technologies as a Senior Embedded Software Engineer. From 2013 to 2018, he was a Researcher with the Networks and Distributed Systems (ND) Group, Department of Informatics, UiO. He is currently an Associate Professor with the Department of Informatics, UiO. He has authored or coauthored several articles in high-ranked conferences and journals, and he has experience in several national (Norwegian Research Council) and international (European Research Funding Agencies) research projects. His research interests include broadly on resource-efficiency, scalability, adaptability, dependability, mobility, and data-intensiveness of distributed systems designed for emerging computing technologies, such as IoT, fog/edge/cloud computing, and cyber-physical systems. He is currently an Associate Editor for IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING.

Jun Xu (Senior Member, IEEE) received the B.S. degree in computer science from the Illinois Institute of Technology, Chicago, IL, USA, in 1995, and the Ph.D. degree in computer and information science from The Ohio State University, Columbus, OH, USA, in 2000. He is currently a Professor with the School of Computer Science at the Georgia Institute of Technology. He has authored or coauthored more than 800 peer-reviewed journal/conference papers. His research interests include designing high-speed Internet routers, firewalls, and measurement devices. He was the recipient of 2006 and 2008 IBM Faculty Awards for making fundamental contributions to the development of system/network performance evaluation methodologies, such as new data streaming algorithms and large deviation techniques, the NSF CAREER Award in 2003, and ACM Sigmetrics 2004 Best Student Paper Award, and Ameritech fellowship for outstanding research in telecommunications. He has been an ACM Distinguished Scientist since 2010. He was a Guest Editor of "Special Issue on Quality-of-Service" in IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT. He was also a TPC Member of a number of conferences, such as, ACM SIGMETRICS, ACM IMC, IEEE INFOCOM, IEEE ICNP, IEEE ICDCS, and IEEE ICC.