Guest Editorial Introduction to the Special Issue on Adaptive Learning Systems in Communication Networks

THE LAST ten years have seen an explosive growth in the progress and adoption of communication networks for data and various multimedia applications. In particular, the emergence of the Internet as a new medium for business transactions, government services, information acquisition, computing and communication has created a vast array of problems unforeseen just a few years ago. As the capabilities of the available networking infrastructure improve, the convergence of data, voice and video transport over the Internet or its derivatives is becoming a reality.

Communication networks and internetworks, and in particular the Internet, have been characterized as the ultimate data-rich environments, dynamically evolving and expanding practically without any centralized control. Such data-rich, unstructured environments present a particular challenge for traditional methods of analysis and design. Adaptive learning methods, in general, including adaptive signal processing, neural networks, fuzzy logic and other data-driven methods and algorithms are in the unique position to offer credible alternatives. Such approaches have the potential for solving and improving the available solutions for some of the toughest problems faced in this newly emerging set of interrelated information technologies.

The goal of this special issue is two-fold:

- to highlight the ongoing research in the field of adaptive learning systems, and in particular adaptive signal processing, neural networks, fuzzy logic, etc., as it is applicable to computer and communication networks;
- to present to the neural networks community and to others interested in adaptive learning systems, in general, a variety of new and challenging problems and their proposed solutions, originating from the rapidly expanding universe of computer and communication networks.

As the use of networking technologies spreads, numerous modeling, estimation, control, classification, clustering and signal processing problems are emerging. Many of these problems currently have no satisfactory solutions and some have been addressed with *ad hoc* solutions with much room remaining for improvements. A common underlying theme of these problems is that they are the following:

- very data-rich;
- represent a dynamically changing environment where the lack of valid mathematical models is predominant;
- representative of systems with minimal or no centralized control.

The problems resulting from the evolution of modern communication networks appear amenable to data-driven approaches and algorithms, such as adaptive learning methods. To a large extent, the focus of the papers in this special issue is on data-driven or the so-called measurement-based methods and algorithms, rooted in the general areas of adaptive learning methods.

The papers appearing in this special issue are grouped into four broad categories based on the network application being pursued by each individual paper. They are: 1) fault detection and diagnosis, 2) estimation and prediction, 3) stability analysis and control, and 4) network management. Even though traditionally fault detection and diagnosis has been considered part of network management, many of the papers in the special issue address problems that go beyond the traditional fault detection and diagnosis topics. As such a separate category has been created. The lead paper(s) for each of the four categories of papers were solicited directly from the authors because of their prominence in their respective research communities.

The group of fault detection and diagnosis papers is led by three solicited papers. Convergence of data, voice and video within the same network infrastructure is raising a number of rather interesting technical issues. One such issue is the characterization of the load and the detection of anomalies present in voice over IP (VoIP) traffic. The paper by Mandjes *et al.* derives new general formulae for the variance of the cumulative traffic over a fixed time interval and show how the derived analytical expression simplifies for the case of VoIP traffic. The paper also proposes simple anomaly detection tests including detection of over/underload. The paper concludes with the application of the scheme to field data from an operational network.

Most Internet services (e-commerce, search engines, etc.) suffer faults that lead to downtime. Downtime, or lost opportunity, costs for such services could be very significant. Quickly detecting these faults can be the largest bottleneck in improving system availability. The paper by Kiciman and Fox presents a methodology for automating fault detection in Internet services. Without requiring any *a priori* application-specific information, experiments performed the authors correctly detected 89%-96% of major failures, as compared with 20%-70% detected by current application-generic techniques.

Probabilistic frameworks for fault detection and diagnosis present another alternative for addressing such problems. The

Digital Object Identifier 10.1109/TNN.2005.857081

paper by Athanasopoulou and Hadjicostis presents such a framework for fault detection in communication networks, using finite state machines. The authors analyze ways to optimally choose detection system parameters and develop adaptive strategies that achieve a low probability of aliasing. An application of these ideas to the problem of protocol testing/classification is presented.

The paper by Hajji addresses the problem of normal operation baselining for automatic detection of network anomalies. An analytical expression of false alarm rate allows the author to choose the detection threshold automatically. Extensive experimental results on a real network showed that the proposed monitoring agent is able to detect unusual changes in the characteristics of network traffic, adapt to diurnal traffic patterns, while maintaining a low false alarm rate.

Barreto *et al.* develop an unsupervised approach to condition monitoring of cellular networks using competitive neural algorithms. The authors perform anomaly detection tests via percentile-based confidence intervals computed over the global and local normality profiles. The performance of four competitive algorithms is compared and the results suggest that the joint use of global and local normality profiles is more efficient and more robust than currently used single-threshold methods.

In mobile *ad hoc* networks (MONET), nodes act both as terminals and information relays, and they participate in a common routing protocol. The network is vulnerable to routing misbehavior, due to faulty or malicious nodes. Misbehavior detection systems aim at removing this vulnerability. In their paper, Sarafijanovic and Le Boudec investigate the use of an artificial immune system (AIS) to detect node misbehavior in a MONET. The approach is inspired by the natural immune system of vertebrates. The goal of the paper is to build a system that, like its natural counterpart, automatically learns and detects new misbehavior.

Real-time problem diagnosis in large distributed computer systems and networks is a challenging task that requires fast and accurate inferences from potentially huge data volumes. Rish *et al.* propose a cost-efficient, adaptive diagnostic technique called active probing. Active probing uses probabilistic reasoning techniques combined with an information-theoretic approach. It enables a fast online inference about the current system state via active selection of a small number of most-informative tests. Empirical tests demonstrate that the active probing scheme greatly reduces both the number of probes (from 60% to 75%), and the time needed for localizing the problem when compared with nonadaptive (preplanned) probing schemes. Some theoretical results on the complexity of probe selection are provided.

The second group of papers is focused on estimation and prediction (or forecasting) applications. The lead paper for the group is on Internet backbone traffic forecasting. The work by Papagiannaki *et al.* introduces a methodology to predict when and where link additions/upgrades have to take place in an IP backbone network. The authors show that IP backbone traffic exhibits visible long term trends, strong periodicities, and variability at multiple time scales. The largest amount of variability in the original signal is due to its fluctuations at the 12 hour time scale. The paper shows that forecasting the long term trend and the fluctuations of the traffic at the 12 hour time scale yields accurate estimates for at least six months in the future.

One-way delay (OWD) traces are important measurements for analyzing end-to-end performance on the Internet. It is still a great challenge to provide a scalable solution for large-scale OWD measurement. Because the clocks at the end systems are usually not synchronized, the OWD measurements are often inaccurate. The paper by Lin *et al.* uses fuzzy-clustering analysis to propose a new algorithm for estimation and removal the clock skews and resets from measurement results. The implementation of the proposed algorithm on several Internet path measurements is presented. The authors' experiments demonstrate that the proposed approach is more accurate and robust than existing techniques. The computational complexity of the proposed method is similar to that of the convex-hull method.

Real-time prediction of video source traffic is an important step in many network management tasks, such as dynamic bandwidth allocation and end-to-end quality of service (QoS) control strategies. In the paper by Atiya *et al.* an adaptive prediction model for MPEG-coded traffic is presented. A novel technology is used, first developed in the signal processing community called sparse basis selection. The method is based on selecting a small subset of inputs (basis) from among a large dictionary of possible inputs. When a new measurement is received, the proposed algorithm updates the selected inputs in a recursive manner. The authors apply the algorithm for single-step-ahead prediction of MPEG-coded video source traffic. The developed method achieves improved results, as compared to published results in the literature.

The lead paper for the stability analysis and control group is on adaptive bandwidth control (ABC). The paper by Siripongwutikorn *et al.* presents the use of ABC for a quantitative packet loss rate guarantee to aggregate traffic in packet switched networks. The proposed ABC starts with some initial amount of bandwidth allocated to a queue and adjust it over time based on online measurements of system states to ensure that the allocated bandwidth is just enough to attain the specified loss requirement. The authors present an extensive simulation study based on both theoretical traffic models and real traffic traces under a wide range of system configurations. They demonstrate that the proposed control itself is highly robust, yields high-bandwidth utilization, and it is a viable alternative to static bandwidth allocation and existing adaptive bandwidth allocation schemes. Additionally, the authors develop a simple and efficient measurement-based admission control procedure which limits the amount of input traffic to maintain the performance of the controller at an acceptable level.

Buffer management in queuing systems plays an important role in addressing the tradeoff between efficiency, measured in terms of overall packet loss, and fairness, measured in terms of individual source packet loss. Ease of implementation is the key issue when determining the practicality of a dynamic buffer management technique. In the paper by Yousefi'zadeh and Jonckheere, two novel dynamic buffer management techniques

are introduced for queuing systems accommodating self-similar traffic patterns. The techniques take advantage of the adaptive learning power of neural networks. Relying on the water-filling approach, the techniques proposed by the authors are capable of coping with the tradeoff between packet loss and fairness issues. Computer simulations reveal that both of the proposed techniques enjoy great efficiency and fairness.

Playout delay adaptation algorithms are often used in real-time voice communication over packet-switched networks to counteract the effects of network jitter at the receiver. More recently developed algorithms strive to achieve better quality by allowing for playout delay adaptation within a talkspurt (intratalkspurt adaptation). The adaptation algorithms rely on short-term estimations of the characteristics of network delay that would be experienced by up-coming voice packets. The paper by Ranganathan and Kilmartin uses neural networks and fuzzy systems as estimators of network delay characteristics. The authors present the design of a fuzzy trend analyzer system for network delay trend analysis and its usage in intratalkspurt playout delay adaptation. The performance of the proposed mechanism is analyzed using measured Internet delays.

The paper by Baglietto, *et al.* considers the problem of resource allocation for a satellite network, where variations of fading conditions are added to those of traffic load. In practice, a good approximation of the optimal solution could be obtained through the adoption of a closed-form expression of the performance measure in steady-state conditions. The authors address two novel optimization approaches. The first is based on the minimization over the discrete constraint set using an estimate of the gradient. The second is based on an open-loop feedback control strategy. The optimization approach leads to a functional optimization problem, and the adoption of a neural network based technique is investigated. The authors obtain near-optimal reallocation strategies with a small computational effort.

Controlling transmitted power in a wireless network is critical for maintaining QoS, maximizing channel utilization and minimizing near-far effect for suboptimal receivers. Paul *et al.* study a general proportional-integral derivative (PID) type algorithm for controlling transmitted power in wireless networks. The authors propose a systematic way to adapt or tune the parameters of the controller in a distributed fashion. The algorithm is data driven and can distinguish between stabilizing and destabilizing controller gains as well as rank the stabilizing controllers based on their performance. Simulation results indicate that the proposed scheme performs better than several candidate controllers, including a well known distributed power control algorithm.

Liu *et al.* develop a call admission control scheme for CDMA cellular networks that handle both voice and data services that can learn from the network environment and user behavior. The idea is built upon a novel learning control architecture with only a single module instead of the two or three modules of the adaptive critic designs (ACDs). The proposed call admission controller can perform learning in real-time as well as in offline environments and the controller improves its performance as it gains more experience. The authors demonstrate the performance of their control algorithm through computer simulation and comparisons to existing algorithms.

The paper by Alpcan *et al.* initiates a study toward developing and applying randomized algorithms for stability of high-speed communication networks. The focus is on congestion and delay based flow controllers for sources which are "utility maximizers" for individual users. In the general case, when closed-form derivation is not possible, the authors construct specific randomized algorithms which provide a probabilistic estimate of the local stability of the network. In particular, the authors use Monte Carlo as well as quasi-Monte Carlo techniques. The results presented provide a complete analysis of congestion control algorithms for internet style networks with a single bottleneck node, as well as for networks with general random topologies.

The last group of papers on network management has a lead article on multiprotocol label switching (MPLS). The combined use of the differentiated services and MPLS technologies is envisioned to provide guaranteed QoS for multimedia traffic in IP networks. These networks need to be managed adaptively to cope with the changing network conditions and provide satisfactory QoS. In this paper, Anjali *et al.* discuss three aspects of the management of a layered MPLS network. In particular, the authors present an optimal technique for the setup of label switched paths (LSPs), capacity allocation of the LSPs and LSP routing. The presented techniques are based on measurements of the network state, and adaptation of the network configuration to changing traffic conditions.

Freeman and Yin present a new method for content management and knowledge discovery using a topology-preserving neural network. The method, termed topological organization of content (TOC), can generate taxonomy of topics from a set of unannotated, unstructured documents. The approach proposed by the authors has been tested and compared with several existing methods on real world web page datasets. The authors have clearly demonstrated the advantages and efficiency of the proposed method in content organization in terms of computational cost and representation. The TOC can be easily adapted for large-scale applications. The topology provides a unique, additional feature for retrieving related topics and confining the search space.

The paper by Nakano and Suda proposes a novel framework for developing adaptive and scalable network services. In the proposed framework, a network service is implemented as a group of autonomous agents that interact in the network environment. The authors use an evolutionary adaptation mechanism using genetic algorithms for agents to evolve their behaviors and improve their fitness values to the environment. The authors evaluate the proposed framework through simulations, demonstrating the ability of autonomous agents to adapt to the network environment. The proposed framework may be suitable for disseminating network services in dynamic and large-scale networks where a large number of data and services need to be replicated, moved and deleted in a decentralized manner.

The study of fair allocation of bandwidth in multicast networks with multirate capabilities forms the basis of the paper

by Sarkar and Tassiulas. In multirate transmission, each source encodes its signal in layers. The lowest layer contains the most important information and all receivers of a session must receive it. The authors present a computationally simple, decentralized scheduling policy that attains the max-min fair rates without using any knowledge of traffic statistics and layer bandwidths. This policy learns the congestion level from the queue lengths at the nodes, and adapts the packet transmissions accordingly. Analytical and simulation results are presented by the authors that guarantee the max-min fairness of the resulting rate allocation, and upper bound the packet loss rates for different layers.

Network security is an important task of network management. One threat to network security is malware (malicious software) propagation. One type of malware is called topological scanning that spreads based on topology information. The focus of the paper by Chen and Ji is on modeling the spread of topological malwares. The model is motivated by probabilistic graphs, widely investigated in machine learning. The authors conduct both theoretical analysis and extensive simulations on large networks using both real measurements and synthesized topologies to test the performance of the proposed models. They show that the model can capture temporal dependence and detailed topology information and, thus, outperforms the previous models in characterizing the behavior of malware propagation.

ACKNOWLEDGMENT

The guest editors would like to thank the authors for their contributions to the special issue, and the referees for their great help in improving and selecting the papers of this volume. The editors feel that this special issue will motivate and attract readers of the Transactions to participate in problems of interest to the networking community.

ALEXANDER G. PARLOS, *Guest Editor (Coordinator)*
Texas A&M University
Department of Mechanical Engineering
College Station, TX 77483

CHUANYI JI, *Guest Editor*
Georgia Institute of Technology
Department of Electrical and Computer Engineering
Atlanta, GA 30332

THOMAS PARISINI, *Guest Editor*
University of Trieste
Department of Electrial, Electronic, and Computer Engineering
Trieste, Italy

MARCO BAGLIETTO, *Guest Editor*
University of Genova
Department of Communications, Computer, and Systems Science
Genova, Italy

AMIR F. ATIYA, *Guest Editor*
Cairo University
Department of Computer Engineering
Cairo, Egypt

K. CLAFFY, *Guest Editor*
University of California at San Diego
Cooperative Association for Internet Data Aanlysis, San Diego Supercomputer Center
La Jolla, CA 92093

Alexander G. Parlos (S'81–M'86–SM'92) received the B.S. degree in nuclear engineering from Texas A&M University, College Station, and the S.M. degree in nuclear engineering, the S.M. degree in mechanical engineering, and the Sc.D. degree in automatic control and systems engineering, all from the Massachusetts Institute of Technology (MIT), Cambridge.

He is currently a Professor of Mechanical Engineering at Texas A&M University, holding joint appointments with the Department of Nuclear Engineering, and, by courtesy, the Department of Electrical Engineering. He has established and is heading the Networked and Intelligent Machines Laboratory (NIML), where research on wireless sensor/actuator networks focuses on enhancing sensor/actuator performance and on machine-to-machine communication over standard IP networks. He holds numerous patents involving "sensorless" machine condition assessment technologies and algorithms for training neural networks for use in estimation and control—many of which have been licensed for commercial use. Since his arrival to Texas A&M University he has supervised the research of over 50 graduate students. He has published over 160 research papers. He is a member of the editorial board of the *Transactions on Control, Automation and Systems Engineering*.

Dr. Parlos is currently a member of the editorial board of the IEEE TRANSACTIONS ON NEURAL NETWORKS. He is also the recipient of multiple awards and honors.

Chuanyi Ji (S'85–M'91) received the B.S. degree (Honor) from Tsinghua University, Beijing, China in 1983, the M.S. degree from University of Pennsylvania, Philadelphia, in 1986, and the Ph.D. degree from the California Institute of Technology, Pasadena, in 1992, all in electrical engineering.

She is currently an Associate Professor in the Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta. She was a Faculty Member at Rensselaer Polytechnic Institute, Troy, NY, from 1991 to 2001. She spent her sabbatical at Bell-Labs Lucent in 1999, and was a visiting faculty at the Massachusetts Institute of Technology, Cambridge, MA, in 2000. Her research lies in the areas of network management and control, and adaptive learning systems. Her research interests are in understanding and managing complex networks, applications of adaptive learning systems to network management, learning algorithms, statistics and information theory.

Dr. Ji received the CAREER award from the National Science Foundation in 1995 and the Early Career Award from Rensselaer Polytechnic Institute in 2000.

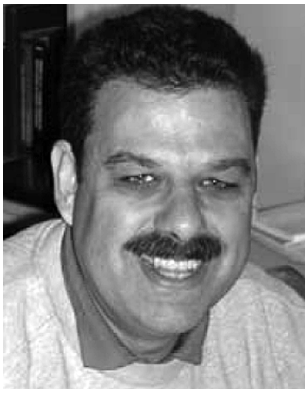

Thomas Parisini (S'87–M'88–SM'99) was born in Genoa, Italy, in 1963. He received the laurea degree (cum laude and printing honors) in electronic engineering and the Ph.D. degree in electronic engineering and computer science, both from the University of Genoa in 1988 and 1993, respectively.

From 1988 to 1995, he was with the Department of Communications, Computer, and Systems Sciences (DIST), University of Genoa. In 1995, he joined the Department of Electrical, Electronic, and Computer Engineering (DEEI), University of Trieste, Trieste, Italy, as an Assistant Professor. In 1998, he joined the Department of Electronic and Information Engineering (DEI), Politecnico di Milano, Milam, Italy, as an Associate Professor. In 2001, he was appointed a Full Professor and Danieli Endowed Chair of Automation Engineering in the Department of Electrical, Electronic, and Computer Engineering (DEEI), University of Trieste. He is currently serving as an Associate Editor of *Automatica*, the *International Journal of Control*, as the Subject Editor of Intelligent Control of the *International Journal of Adaptive Control and Signal Processing*. His research interests include neural-network approximations for optimal control and filtering problems, fault diagnosis for nonlinear systems, and hybrid control systems.

Dr. Parisini is the present Chair of the IEEE Control Systems Society Conference Editorial Board. He was the Chairman of the Technical Committee on Intelligent Control and an appointed member of the Board of Governors of the IEEE Control Systems Society. He is a senior member of the IEEE and a Distinguished Lecturer of the IEEE Control Systems Society. He has served as an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL and the IEEE TRANSACTIONS ON NEURAL NETWORKS. He was involved in the organization and in the technical program committees of several IEEE CSS sponsored conferences including the IEEE Conference on Decision and Control and the IEEE Conference on Control Applications.

Marco Baglietto (M'04) was born in Savona, Italy, in 1970. He received the Laurea degree in electronic engineering and the Ph.D. degree in electronic engineering and computer science from the University of Genoa, Genoa, Italy, in 1995 and 1999, respectively.

Since 1999, he has been an Assistant Professor of Automatic Control in the Department of Communications, Computer and Systems Science (DIST), University of Genoa. His research interests include neural approximations, linear and nonlinear estimation, distributed-information control systems.

Amir F. Atiya (S'86–M'90–SM'97) was born in Cairo, Egypt, in 1960. He received the B.S. degree in electrical engineering from Cairo University, Cairo, Egypt, in 1982 and the M.S. and Ph.D. degrees in electrical engineering from Caltech, Pasadena, CA, in 1986 and 1991, respectively.

He is currently an Associate Professor in the Department of Computer Engineering, Cairo University. He held various appointments in industry, such as in QANTXX, Houston, TX; Simplex Technology, Hong Kong; Countrywide, California; and Dunn Capital Management, Florida. From 1997 to 2001, he was a Visiting Associate in Caltech. His research interests are in the areas of neural networks, learning theory, pattern recognition, Monte Carlo methods, data compression, and optimization theory. His most recent interests are the application of learning theory and computational methods to financial prediction and to communications networks prediction problems.

Dr. Atiya received the highly regarded Egyptian State Prize for Best Research in Science and Engineering in 1994. He also received the Young Investigator Award from the International Neural Network Society, in 1996. He has been an Associate Editor for the IEEE TRANSACTIONS ON NEURAL NETWORKS since 1998. He was a Guest Co-editor of the special issue of the IEEE TRANSACTIONS ON NEURAL NETWORKS on "Neural Networks in Financial Engineering," July 2001. He served on the organizing and program committees of several conferences, including being Program Co-Chair for the IEEE Conference on Computational Intelligence in Financial Engineering (CIFER-03), March 2003, Hong Kong.

K. Claffy is with the Cooperative Association for Internet Data Analysis (CAIDA), based at the San Diego Supercomputer Center, University of California at San Diego, La Jolla. Her research focuses on collection, analysis and visualization of wide-area Internet data on topology, workload, performance, and routing. CAIDA seeks, through the development of neutral, open tools and methodologies, to promote cooperation among internet service providers (ISPs) in face of commercialization and competition, and to meet engineering and traffic analysis requirements of the commercial Internet community.