Guest Editorial Special Issue on Green Pervasive and Ubiquitous Systems

PERVASIVE and ubiquitous computing is a promising paradigm in human's daily life, which enables computation, storage, and communications services to appear everywhere and anytime. It requires seamless integration and interaction of heterogeneous and complex systems, ranging from a variety of *local* smart yet resource-constrained devices (e.g., smart phones, tablets, laptops, and wireless mobile sensors) to *remote* powerful cloud platforms (e.g., datacenters, high-performance computing, parallel and distributed systems). Such a pervasive *ecosystem* not only depends on performance-oriented system infrastructures and ubiquitous computing/networking technologies, but also faces critical energy-efficiency issues, exemplified by controlling the tremendous power consumption, electricity cost, and carbon-emission of large-scale datacenter systems, as well as prolonging the limited battery lifetime of mobile devices in broad Internet applications. Therefore, new energy-aware architectures, emerging systems-of-systems and related techniques, such as mobile-cloud computing, virtualization, and scheduling of renewable energy resources, are increasingly drawing development efforts from both academia and industry, in order to shape an environmentally sustainable "Green" world for the next decade. Due to the challenging open issues, green pervasive and ubiquitous systems deserve academic attentions from the interdisciplinary aspects of information systems, computing and networking optimizations, energy control, and engineering technologies.

The paper "Energy-Aware Data Allocation with Hybrid Memory for Mobile Cloud Systems" by Meikang Qiu *et al*. focuses on resource scheduling in mobile cloud computing due to the constraints in memory, CPU, and bandwidth. High-energy consumption and low performance of memory accesses have become overwhelming obstacles for chip multiprocessor (CMP) systems used in cloud systems. In order to address the daunting "memory wall" problem, hybrid on-chip memory architecture has been widely investigated recently. Due to its advantages in size, real-time predictability, power, and software controllability, scratch-pad memory (SPM) is a promising technique to replace the hardware cache and bridge the processor–memory gap for CMP systems. However, if not well managed, frequent accesses to SPM will still be a major energy and computation time consumer. The authors present a novel hybrid on-chip SPM that consists of an SRAM, a magnetic RAM (MRAM), and a zero-capacitor RAM for CMP systems by fully taking advantages of the benefits of each type of memory. To reduce memory access latency, energy consumption, as well as the number of write operations to MRAM, they also propose a novel multidimensional dynamic programming data allocation algorithm to strategically allocate data blocks to each memory.

The paper "Energy Conservation in Progressive Decentralized Single-Hop Wireless Sensor Networks for Pervasive Computing Environment" by Chen Yu *et al*. discusses about energy efficiency issue that is very important in terms of prolonging the lifetime of the communication in pervasive computing environments. As a practical application in pervasive computing environment, wireless sensor networks (WSNs) consist of many sensors and several access points that make them work cooperatively to monitor/measure certain areas. Since the deployment of sensors in unknown sites impedes their recharging, thus exhausting their energy quite quickly, energy conservation becomes a critical issue. Unavoidably, the lavishness of both single-hop and multihop modes on energy conservation declines the system's life span severely. Therefore, a progressive decentralized single-hop method is conceived. This mechanism works with several phases, in each of which sensors may act in multihop or single-hop mode. The well-balanced energy consumption rate results in the extension of the system's life span. The method also adapts in general cases, and has been proven by mathematical demonstration to totally balance the energy consumption.

The paper "Energy Efficiency of Cloud Virtual Machines: From Traffic Pattern and CPU Affinity Perspectives" by Chi Xu *et al*. addresses networking and machine virtualization in the success of modern cloud computing. The energy consumption of physical machines has been carefully examined in the past, including the impact from network traffic. When it comes to virtual machines (VMs) in the cloud datacenters, the interplay between energy consumption and network traffic however becomes much more complicated. Through real-world measurement on both Xen- and KVM-based platforms, the paper shows that these state-of-the-art virtualization designs noticeably increase the demand of CPU resources when handling network transactions, generating excessive interrupt requests with ceaseless context switching, which in turn increases energy consumption. Even when a physical machine is in an idle state, its VM's network transactions will incur nontrivial energy consumption. More interestingly, the energy consumption varies significantly with traffic allocation strategies and virtual CPU affinity conditions, which was not seen in conventional physical machines. Looking closely into the virtualization architectures, they then

Digital Object Identifier 10.1109/JSYST.2017.2673218

pinpoint the root causes and provide initial solutions toward optimizing energy consumption in virtualized environments.

The paper "On Efficient Utilization of Green Energy in Heterogeneous Cellular Networks" by Bang Wang *et al*. deals with green energy issue in heterogeneous cellular networks. Efficient energy utilization to reduce carbon footprint has become a great concern in diverse domains. In cellular mobile communications, using green energy and adopting heterogeneous network have been recently promoted as a strategic shift for reducing greenhouse gas emission. Considering a green heterogeneous cellular network, this paper studies how to design user association algorithms for efficient utilization of green energy. The authors first formulate a constrained total energy cost minimization problem, and decompose the problem into two components: the user association problem and optimal bandwidth allocation problem, which are complementary to each other. Then, an optimal bandwidth allocation algorithm is proposed to minimize the power consumption of each BS under a given user association scheme. For the user association problem, the authors propose both a centralized and a distributed user association algorithm to find feasible and near-optimal solutions based on the optimal bandwidth allocation algorithm, making the efficient utilization of the green energy storage for energy cost minimization. Simulations are used to compare their algorithms with the maximum channel gain algorithm and the max-RSRP algorithm that also apply the optimal bandwidth allocation.

The paper "Energy and Delay Tradeoff for Application Offloading in Mobile Cloud Computing" by Xiumin Wang *et al*. focuses on application offloading in mobile could computing. Recent work shows that offloading mobile application from mobile devices to cloud servers can significantly reduce the energy consumption of mobile devices, and thus extend the lifetime of mobile devices. However, the previous work only considers the energy saving of mobile devices, while ignoring the execution delay of mobile applications. To reduce the energy consumption of mobile devices, one may offload as many mobile applications as possible. However, offloading to cloud servers may incur large execution delay, because of waiting time at servers or the communication delay from mobile devices to the servers. Thus, to balance the tradeoff between energy consumption and execution delay of mobile applications, it is necessary to determine whether the mobile application should be offloaded to the cloud server or run locally at the mobile devices. The authors first formulate a joint optimization problem, which minimizes both the energy consumption at mobile devices and execution delay of mobile applications. They prove that the proposed problem is NP-hard. For a special case with unlimited residual energy at mobile device and the same size of resources required by each mobile application, they propose a polynomial-time optimal solution and an efficient heuristic algorithm to solve the general case of the problem.

The paper "A Green TDMA Scheduling Algorithm for Prolonging Lifetime in Wireless Sensor Network" by Jun Long *et al*. discusses about time-division multiple access (TDMA) scheduling issues. In WSNs, sensor nodes are energy constrained since a battery, which is hardly replaced once deployed, powers them. The sensor nodes detect a specific phenomenon in the environment and route the sensed data to a relatively small number of central data processing nodes, called sinks. That causes imbalance in energy consumption among the nodes such that nodes closer to the sinks consume more energy than others. According to typical network parameters, the authors theoretically analyze a formula of energy consumption for general k-hop networks and find optimal k such as making a network lifetime the longest. On this basis, they propose a TDMA scheduling algorithm for the k-hop networks and show the upper bound of time slots required in the k-hop networks. Theoretical analysis and simulation results demonstrate the accuracy and effectiveness of the proposed algorithm.

The paper "Energy-Efficient Web Server Load Balancing" by Jorg Lenhardt *et al*. deals with web server load balancing. Rising energy costs and negative environmental impact resulting from the generation of electricity, especially when relying on fossil fuels, increase the need for energy-efficient computing. Besides the optimization of hardware devices, the use of sophisticated software solutions for reducing energy consumption is an important area of current research. Distributing load among servers for performance optimization is a well-researched area. There is a large potential to reduce energy consumption with similar techniques. In this paper, the authors examine the application of energy-efficient load balancing (or unbalancing) strategies for web server requests send to a web server farm. Single requests usually lead to small fractions of performance need. They are quite suited for strategies relying on the assumption of a fully divisible load. The power consumption is continuously reduced for replays of Wikipedia access traces. Eventually, the energy consumption in web server farms is reduced significantly.

The paper "Data-Aware Task Dispatching for Batch Queuing System" by Xieming Li and Osamu Tatebe describes a scheduling method focusing on exploiting the local access of nonuniform filesystem. In their approach, the cost of file accessing is calculated, and combined with CPU load average to a comprehensive value. This value will be used as the standard for scheduling. They evaluated approach with three benchmarks in comparison with the original Torque scheduler: thput-gfpio, readgf, and BLAST Benchmark. In thput-gfpio, the read throughput showed 78% boost up, while in readgf, the total execution time is reduced to about 1/6. Finally, with the BLAST Benchmark, the total execution time is reduced by 50%.

The paper "A Comprehensive Evaluation of Scheduling Methods of Virtual Machine Migration for Energy Conservation" by Dancheng Li and Osamu Tatebe. addresses the scheduling methods of VM migration for energy conservation. A number of scheduling methods of VM migration have been proposed and/or improved in several research works. However, none of the scheduling methods was evaluated from a comprehensive viewpoint and there is no useful reference for practices of energy conservation in various data centers. In order to provide a useful reference for best practices of energy conservation in various data centers, the authors present the first comprehensive evaluation of scheduling methods of VM migration. After giving an overview of major optimization problems, they propose a new set of evaluation metrics, which can be used in the evaluation of various scheduling methods of VM migration from different aspects, present an evaluation environment they constructed according to their evaluation metrics, and discuss

evaluation results of those proposed scheduling methods on their evaluation environment.

The paper "Game Theoretic Market Driven Smart Home Scheduling Considering Energy Balancing" by Yang Liu *et al*. focuses on smart home scheduling. In a smart community infrastructure that consists of multiple smart homes, smart controllers schedule various home appliances to optimize the energy balancing and reduce the electricity bills of customers. This process has an implicit effect of reducing the power generation facilities, thus mitigating the potential pollution due to excessive power generation. In this paper, the impact of the smart home scheduling to the electricity market is analyzed using a proposed smart home-driven bilevel market model. In this model, the customers schedule home appliances for bill reduction at the community level, while aggregators minimize the energy purchasing expense from utilities at the market level, both of which consider the smart home scheduling impacts. A game theoretic algorithm is proposed to solve this formulation, which handles the bidirectional influence between the community level and the market level.

The paper "Incentive Load Scheduling Schemes for PHEV Battery Exchange Stations in Smart Grid" by Feng Ye *et al*. addresses load scheduling in smart grid. The authors study load-scheduling schemes for plug-in hybrid electric vehicle battery exchange stations (BESs) in smart grid. Since each BES stores a relatively large amount of batteries, it can significantly contribute to the demand side management (DSM) system in smart grid by selling back the electricity to the grid during peak hours. By doing so, the peak-to-average ratio (PAR) of the grid can be further be reduced on top of existing DSMs for other applications. In order to achieve that, the authors propose several load scheduling schemes for BESs. One is to minimize the PAR, followed by an incentive scheme so that the BESs will be motivated to participate. They also propose a game theoretical scheme so that the load scheduling for each BES can be done locally with limited information exchange. In the simulations, they show that BESs contribute to DSM and further smooth the load of the power grid although the total load is increased due to large amount of PHEVs. They also analyze the impact of total amount of battery storage at BESs on PAR. Moreover, they analyze the impact of charging ports number in each BESs on the total amount of battery storage as well as PAR.

The paper "Hybrid DVFS Scheduling for Real-Time Systems Based on Reinforcement Learning" by Fakhruddin Muhammad Mahbub ul Islam and Man Lin discusses about hybrid DVFS scheduling in real-time systems. Power consumption is one of the most challenging issues in the design of modern computing systems. In any computational devices, processor consumes significant amount of power comparing with other components. In order to reduce power consumption and control temperature, dynamic voltage and frequency scaling (DVFS) has been commonly used in modern processors. In recent years, there have been many research works on real-time DVFS techniques. They work with different strategies and perform well for a given set of conditions. However, a single algorithm is not optimal for all operating conditions. One algorithm may outperform another under different workloads, dynamic slacks, and power settings. For example, some approaches are energy efficient in low workload conditions, whereas some are at higher. Some are appropriate when dynamic slack is higher. Furthermore, the variation of device power configuration also affects the performance of a given DVFS algorithm. Therefore, aiming for adaptability, the authors propose a novel reinforcement learning-based approach which takes a set of well-known existing techniques, specialized to handle different conditions, and switches to the most suitable one in various situations.

The paper "Subtasks Scheduling for Distributed Robots in Cloud Manufacturing" by Wenxiang Li *et al*. addresses robot scheduling in cloud manufacturing (CMF). CMF provides a sharing and cooperation platform for efficient utilization of distributed manufacturing resources, e.g., industrial robots. However, effective scheduling of tasks or subtasks to these resources is a challenging problem in CMF. Based on the analysis on the elements and procedure of task processing, this paper proposes a CMF system model for efficiently exploiting the available manufacturing resources, so manufacturing robots of different locations and functions can cooperatively handle a batch of tasks. Specifically, this paper considers the performance of four distributed robots deployment methods, including random deployment, robot-balanced deployment, function-balanced deployment, and location-aware deployment. Furthermore, based on genetic algorithm, three subtask-scheduling strategies are derived from this model for three optimization objectives, including load-balance of robots, minimizing overall cost, and minimizing overall processing time. Simulation results demonstrate that each strategy can achieve the relevant optimization objective, respectively.

The paper "Energy-Aware Resource and Revenue Management in Federated Cloud: A Game Theoretic Approach" by Mohammad Mehedi Hassan *et al*. discusses about resource and revenue management in clouds. Reduction of energy expenditure is becoming an important issue for a cloud provider (CP) when providing cloud services over the Internet. Federation among CPs, whereby a set of CPs cooperating together to provide VM instances requested by users, can be a promising solution to address this issue. This paper presents an efficient resource and revenue sharing mechanism in a cloud federation that can lead to a global energy sustainability policy for the federation and motivates them to cooperate. The proposed mechanism models the interactions among the CPs in a federation as a coalition game. However, unlike the existing approaches, the game model aims at finding a set of low-energy cost CPs in a federation, such that the social welfare is maximized and providing a fare and suitable revenue for them. In addition, the authors consider the demand variations of internal users of a CP as part of the VM sharing decision. Moreover, they present a comprehensive analysis of the related costs and revenue associated with the various decisions of the CPs in a federation. Simulation results demonstrated that the proposed mechanism could satisfy the fairness and stability properties, maximize the social welfare of the CPs, and achieve cost effectiveness of resource.

The paper "Dynamic Power Management with Optimal Time-Out Policies" by Hiroyuki Okamura and Tadashi Dohi addresses the issue to switch a high-power consuming state to a low-power

consuming one at a suitable timing in the dynamic power management (DPM). This problem has been formulated as Markov decision processes with state-dependent control in the past literature. However, this approach may not be often feasible in many practical situations, because the state-dependent policy requires that all the states on a request arrival process must be observed through an online monitoring. The authors develop a simple time-out policy in the DPM, which can be regarded as the optimal timing to take a GO-SLEEP action during an idle period of the transaction system. They derive the optimal time-out policy minimizing the expected power consumption per unit time in the steady state analytically under the assumption that the request arrival process is given by a Markovian arrival process with an arbitrary number of phases.

The paper "Performance Monitoring Based Traffic-Aware Virtual Machine Deployment on NUMA Systems" by Yuxia Cheng *et al*. discusses about the performance-monitoring issue in nonuniform memory access (NUMA) systems. Virtualization technology enables multiple VMs to share a single physical server. Commercial servers increasingly use the NUMA architecture due to its scalable memory performance. However, multiple VMs running on a NUMA physical server will cause performance overheads such as remote memory access latency and shared microarchitectural resources contention, which makes the VM performance less efficient and stable. These performance overheads are mainly caused by memory traffic from data intensive workloads. The authors propose a traffic-aware VM optimization (TAVO) scheme on NUMA systems. Based on the performance monitoring of data traffic and CPU/memory resource usages in the system, TAVO addresses VM memory access locality and shared resource contention problems via automatic VM initial placement and NUMA-aware VM online scheduling.

The paper "Enabling Technologies for Green Internet of Things" by Faisal Karim Shaikh *et al*. covers the recent technological advances, which have led to an increase in the carbon footprint. Energy efficiency in the Internet of Things (IoT) has been attracting a lot of attention from researchers and designers over the last couple of years paving the way to an emerging area called Green IoT. There are various aspects (such as key enablers, communications, services, applications) of IoT, where efficient utilization of energy is needed to enable a Green IoT environment. The authors explore and discuss how the various enabling technologies (such as the Internet, smart objects, sensors, etc.) can be efficiently deployed to achieve a Green IoT. Furthermore, they also review various Green IoT applications and standardization efforts currently under way. Finally, they identify some of the emerging challenges that need to be addressed in the future to enable a Green IoT.

The paper "Practical Performance of MANETs under Limited Buffer and Packet Lifetime" by Yujian Fang *et al*. works on the performance of mobile ad hoc networks (MANETs). While scaling law results on the performance of MANETs have been extensively reported in the literature, the exact performance of such networks, in particular, their real achievable performance under practical network constraints, is still largely unexplored. As one step toward practical performance study for MANETs, this paper considers a MANET with constraints on both buffer size and packet lifetime and explores the impacts of such constraints on network performance. The authors first provide analysis on the exact throughput capacity of the network to reveal its maximum possible and input rate-independent throughput performance. With the help of the embedded Markov chain theory, a complete theoretical framework is then developed, which enables the achievable and input rate-dependent throughput and packet loss ratio to be derived in closed form under any exogenous rate. Based on the M/G/1/K queuing theory, the packet end-to-end delay under any exogenous rate is further studied to give a relatively whole picture on how the buffer size and packet lifetime impact the network throughput, packet loss ratio, and packet delay.

The paper "A Fairness-Aware Pricing Methodology for Revenue Maximization in Service Cloud Infrastructure" by Yuanfang Chi *et al*. discusses about fairness-aware pricing methodology in service cloud infrastructure. Infrastructure as a Service CPs sell their resources as different types of VMs called instances. Cloud tenants run their applications or other computational tasks on the VMs they have requested and the number of VMs that they need varies from time to time. Any resource is limited in capacity, so is the resource in cloud. Development of pricing models that can be used as a tool not only for CPs to cover their costs and realize a profit, but also to encourage cloud tenants to use cloud resources efficiently and to help CPs better estimate resource demand level and achieve high utilization has become a key research problem. The authors propose a new pricing methodology that encourages cloud tenants whose requested VMs can be allocated easily and fairly to use more cloud service by offering them lower prices, while discouraging cloud tenants whose requested VMs are difficult to allocate to use cloud service by charging them higher prices, and maximizes the revenue a CP receives. They perform a case study of the proposed pricing methodology with a multiresource allocation fairness algorithm, the dominant resource sharing algorithm, and show that the proposed pricing model achieves personal and social fairness.

The paper "Attribute-Based Hash Proof System Under Learning-With-Errors Assumption in Obfuscator-Free and Leakage-Resilient Environments" by Mingwu Zhang *et al*. addresses attributed-based hash proof systems (AB-HPSs). Node attributes such as MAC and IP address, and even GPS position can be considered as exclusive identity in the distributed networks such as cloud computing platform, wireless body area networks, and IoT. Nodes can exchange or transmit some important information in the networks. However, as the openness and exposure of node in the networks, the communications between the nodes are facing lots of security issues. Especially, sensitive information may be leaked to the attackers in the presence of side-channel attacks, memory leakages, and time attacks, etc. In this work, the authors present a new notion of AB-HPS in the bounded key-leakage model, to be resistant to the possible quantum attackers. They also give the construction of AB-HPS in lattices, and prove the security of indistinguishability of valid and invalid ciphertext and leakage smoothness under the decisional LWE assumption. They also provide the general

leakage resilient attribute-based encryption construction using AB-HPS as the primitive. Finally, they discuss some extensions to improve the schemes in larger space for the message, larger alphabet for the attribute, and arbitrary access structure for the policy, respectively.

The paper "Rotating Directional Sensors to Mend Barrier Gaps in Line-Based Deployed Directional Sensor Network" by Jiaoyan Chen *et al*. discusses about the barrier coverage of WSN in IoT. Barrier coverage guarantees that a chain of connected sensors detects all intruders traversing the protected regions. However, when the sensors are randomly deployed, the barrier gaps may occur due to the deployment randomness or insufficient sensors. How to locate the barrier gaps and mend them is an important aspect in the network. The authors study the barrier gap problem in weak and strong barrier coverage, which consists of directional sensors, and the sensors are deployed by line-based deployment strategy. A gap finding algorithm is proposed to find subbarriers and barrier gaps. Two gap mending algorithms are devised to mend barrier gaps in the network: One is a simple rotation algorithm that only rotates two critical sensors in two subbarriers to fix the gap; the other one is the chain-reaction rotation algorithm that rotates sensors in the subbarrier in a chain-reaction manner to mend the gap.

The paper "Lifetime Optimizing Clustering Structure Using Archimedes' Spiral-Based Deployment in WSNs" by Amrita Ghosal and Subir Halder intends to optimize the clustering structure. WSNs are receiving significant attention due to their potential applications in environmental monitoring and surveillance domains. In WSNs, preserving energy requires utmost attention, as they are highly resource constrained. Clustering is commonly considered as one of the efficient energy conservation technique. In clustered WSN, due to the nature of operation, cluster heads (CHs) near the sink bear the major share of data forwarding than the CHs far away from the sink, resulting in energy hole problem. First, the authors have analyzed the optimization of network lifetime by balancing the energy consumption among different CHs. To meet the requirement of optimization of network lifetime, they have devised a routing aware optimal clustering strategy. Furthermore, keeping the network model in mind, they identify Archimedes' spiral, based on which a deployment function is proposed for distributing MN and CH. Performance of the optimal clustering strategy is evaluated in terms of energy balance and network lifetime.

The paper "A Hybrid eBusiness Software Metrics Framework for Decision Making in Cloud Computing Environment" by Feng Zhao *et al*. introduces an eBusiness software metrics framework for clouds. Developing high-quality software is essential for eBusiness organizations to cope with drastic market competition. With the development of cloud computing technologies, eBusiness systems and applications pay more attention on open-endedness. In cloud computing environments, eBusiness systems have the ability to provision IT resources on demand. Traditional software metric methods in distributed systems and applications are technical and project-driven, that the market demand and internal practical operation cannot be perfectly balanced within cloud computing-based eBusiness corporation. To address this issue, this paper presents a hybrid framework based on the Goal/Question/Metric paradigm to evaluate the quality and efficiency of previous software products, projects, and development organizations in cloud computing environment. In their approach, for supporting decision-making at the project and organization level, three-angular metrics are used: project metrics, product metrics, and organization metrics. Furthermore, an improved radial basis function-based model is also provided to manage current projects and scheme future projects.

The paper "Modeling User Activity Patterns for Next Place Prediction" by Chen Yu *et al*. intends to predict location in pervasive computing systems. Beyond the current location, knowing the individual's next location in advance can also enable many novel mobile applications and services such as targeted advertising and smooth handover between two separate networks. Although extensive studies about location prediction have been carried out, the existing prediction methods either encounter "cold start" problems when the individual's trajectory data are sparse, or perform erratically when the individual performs activities in a new region. In this paper, the authors propose a novel approach based on the activity pattern for location prediction. Instead of directly predicting the individual's next location, they first infer the individual's next activity by modeling user activity patterns, and then predict his/her next location on the basis of the inferred next activity. Using the real-life trajectory data, they demonstrate that the proposed approach can realize the smooth upgrade of prediction performance and perform robustly.

We would like to express our sincere appreciation to all the reviewers, who helped us in selecting the best papers for this special issue. Last but not least, we are grateful to all the authors for their valuable contributions.

HAI JIANG, *Guest Editor*
Department of Computer Science
Arkansas State University
Jonesboro, AR 72401 USA

FANGWING LIU, *Guest Editor*
School of Computer Science and Technology
Huazhong University of Science and Technology
Wuhan, China

RUPPA K. THULASIRAM, *Guest Editor*
Computational Finance Laboratory
Department of Computer Science
University of Manitoba
Winnipeg, MB R3T 2N2, Canada

EROL GELENBE, *Guest Editor*
Department of Electrical and Electronic Engineering
Imperial College London
London SW7 2AZ, U.K.

Hai Jiang (S'02–M'03) received the B.S. degree from the Beijing University of Posts and Telecommunications, Beijing, China, and the M.A. and Ph.D. degrees from Wayne State University, Detroit, MI, USA.

He is a Professor with the Department of Computer Science, Arkansas State University, Jonesboro, AR, USA. Before joining Arkansas State University, he has spent four years with the State Key Laboratory of Switching Technology and Telecommunication Networks, Beijing, China, as a Research Fellow and five years in industry with Ford Motor Company, Dearborn, MI, USA, where he held various technical positions. He has authored 2 books and more than 100 papers in refereed journals, conference proceedings, and book chapters. His current research interests include parallel and distributed systems, computer and network security, high-performance computing and communication, and modeling and simulation.

Dr. Jiang is a professional member of the ACM and IEEE Computer Society. He has been involved in 100 conferences and workshops as a program/workshop chair or as a Program Committee member. He serves as an Editor for the *International Journal of High Performance Computing and Networking*, *International Journal of Computational Science and Engineering*, and *International Journal of Embedded Systems*, as well as a Guest Editor for the IEEE SYSTEMS JOURNAL, *Cluster Computing*, and *Concurrency and Computation: Practice and Experience*. He also serves on the Editorial Board of the *International Journal of Big Data Intelligence*, *GSTF Journal on Social Computing*, *Open Journal of Internet of Things*, and *The Scientific World Journal*.

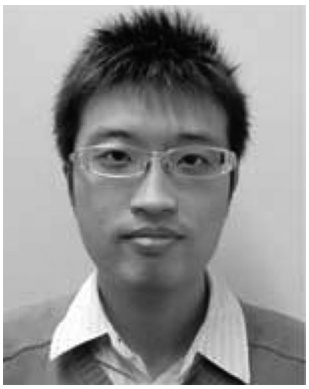

Fangwing Liu (S'08–M'10–SM'16) received the B.Eng. degree from Tsinghua University, Beijing, China, and the Ph.D. degree from the Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong.

He is a Full Professor with the Huazhong University of Science and Technology, Wuhan, China. His research interests include cloud computing/storage and datacenter, mobile cloud and D2D/P2P, green computing, SDN/NFV, and virtualization. He was selected into the National Program for Support of Top-Notch Young Professionals of the National Program for Special Support of Eminent Professionals. He is a Youth Scientist of the National 973 Basic Research Program Project of SDN-based Cloud Datacenter Networks. He was a StarTrack Visiting Faculty with Microsoft Research Asia in 2012–2013.

Dr. Liu has been the Editor-in-Chief of the *EAI Endorsed Transactions on Collaborative Computing*, a Guest Editor for *IEEE Network Magazine* and IEEE SYSTEMS JOURNAL, and served as TPC for ACM Multimedia 2014 and 2016, ACM e-Energy 2016–2017, IEEE INFOCOM 2013–2017, ICNP 2014 TPC and 2016 Poster/Demo Co-Chair, IWQoS 2016–2017, and ICDCS 2015–2017. He is the TPC Co-Chair of INFOCOM 2017 SmartCity workshop.

Ruppa K. Thulasiram (Tulsi) (M'00–SM'09) received the Ph.D. degree from the Indian Institute of Science, Bangalore, India.

He is a Professor and the Director of the Computational Finance Laboratory, Department of Computer Science, University of Manitoba, Winnipeg, MB, Canada. He spent years with Concordia University, Montreal, QC, Canada; Georgia Institute of Technology, Atlanta, GA, USA; and the University of Delaware, Newark, DE, USA, as a Postdoctoral Researcher, Research Staff, and Research Faculty before taking up a position with the University of Manitoba. He has graduated many students with M.Sc. and Ph.D. degrees. He has developed a curriculum for the cross-disciplinary computational finance course at the University of Manitoba for both graduate and senior undergraduate levels. He has authored or co-authored many papers in the areas of high-temperature physics, gas dynamics, combustion, computational finance, and grid/cloud computing. His current research interest includes grid/cloud computing, computational finance, cloud resources management, computational intelligence, ad hoc networks, and scientific computing. His research has been funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

Dr. Thulasiram has been an Associate Editor for the IEEE TRANSACTIONS ON CLOUD COMPUTING, and *International Journal of Aerospace Innovations* (MultiSceince Publishing). He is a member of the Editorial Board of many journals including the *International Journal of Computational Science and Engineering*. He has been a guest editor of many journals such as *Parallel Computing*, *Concurrency and Computation Practice and Experience*, the *International Journal of Parallel, Embedded and Distributed Systems*, and the *Journal of Supercomputing* for special issues. He was the recipient of many Best Paper Awards.

Erol Gelenbe (S'67–M'70–SM'79–F'86–LF'11) is a Professor with the Department of Electrical and Electronic Engineering, Imperial College London, London, U.K. He was an Associate Dean of Engineering and the Director of the School of Electrical Engineering and Computer Science, University of Central Florida (1998–2003), a Head of Electrical and Computer Engineering and Nello L. Teer Chair Professor, with secondary appointments as a Professor of computer science and experimental psychology with Duke University (1993–1998), a Professor with the University of Paris, and has also taught at the Ecole Polytechnique (1979–1993). His current research addresses four related areas: "autonomic" or self-aware and self-organizing software defined networks including the cognitive packet network, energy savings in the cloud and through energy harvesting for sensors in the Internet of Things (IoT), security in mobile networks, and reinforcement and deep learning with random neural networks. His work has also included analytical solutions to networked systems in engineering, physics and biology, including energy networks, neural networks, genetic networks, sensor networks, and economic agents in the Internet. Over the last 10 years, his publications have appeared in journals such as *Physical Review*, *IEEE Network*, IEEE TRANSACTIONS ON EMERGING TOPICS IN COMPUTING, IEEE JOURNAL OF SELECTED AREAS IN COMMUNICATIONS, *Energies*, *ACM Performance Evaluation Review*, IEEE ACCESS, IEEE TRANSACTIONS ON CLOUD COMPUTING, *Communications of the ACM*, *ACM Transactions on Sensor Networks*, *ACM Transactions on Internet Technology*, *Proceedings Royal Society*, *ACM Transactions on Adaptive and Autonomous Systems*, IEEE/ACM TRANSACTIONS ON BIOINFORMATICS AND COMPUTATIONAL BIOLOGY, IEEE TRANSACTIONS ON NANOBIOSCIENCE, *Neural Computation*, and *The Computer Journal*.

Dr. Gelenbe has served as Science and Technology Advisor to the French Minister of Universities (1984–1986). He serves on the Editorial Boards of several journals such as *Acta Informatica*, *Performance Evaluation*, *Computational Management Science*. Previous Editorial Boards he has served on include the *Proceedings of the Royal Society* (2006–2011), PROCEEDINGS OF THE IEEE (1994–1998), and *The Computer Journal* (2008–2012). He is currently the General Chair for MASCOTS 2016. He is a Fellow of the ACM, IET, the Royal Academy of Sciences, Arts and Letters of Belgium (2015), the Hungarian (2010), Polish (2013), Turkish (2012) Science Academies, the National Academy of Technologies of France (2008), and Academia Europaea (2005).