Computer Networks 56 (2012) 3007–3009

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Editorial

Special issue on "Challenges in high-performance switching and routing in the Future Internet"

High-performance, high-capacity packet switches and routers in the metro and backbone network are a basic infrastructure requirement to support the Future Internet. Trends in traffic evolution point to higher traffic volumes and increasingly unpredictable traffic variations caused by the mobility of large multimedia flows. These tendencies translate into challenges to switching and routing in the aggregation and core networks. This special issue is an effort to put together original papers on the design of high-performance switching and routing systems to address these challenges, including architectural, planning, traffic engineering and algorithmic aspects of electronic and optical high-speed networks.

In response to the call for papers for this special issue, we received 43 submissions for consideration for publication. After a rigorous two-step review process, we accepted 12 high quality papers reporting state of the art in high-performance switching and routing.

Seven out of the twelve papers included in this special issue [1–7], report novel and surveying proposals addressing the wire-speed electronic packet-by-packet processing at multi-Gbps. As traffic grows, nodes are required to aggregate thousands of user packet flows, while meeting their QoS requirements. This entails boards designed with dedicated architectures (ASIC or FPGA based) or specialized programmable devices (network processors). The packet processing cycle implemented in the line cards includes (i) packet classification to identify the packet flow, (ii) IP address lookup to obtain the packet forwarding information, (iii) packet filtering to implement security strategies in traffic forwarding, and (iv) a careful scheduling of the queues prior to packet transmission, to satisfy the QoS requirements. The use of active queue management schemes, that make selective and preventive packet drops to avoid link congestion, are also under debate.

Classification of packets in the line cards requires completing multidimensional searches at a very high speed. Hierarchical packet classification has been proposed in the literature to reduce the search time. Lim et al. [1] in their paper "A New Hierarchical Packet Classification Algorithm" propose an efficient packet classification algorithm using the hierarchical approach where a hierarchical binary search tree which is free of empty internal nodes is constructed. This makes the algorithm efficient due to the absence of back-tracking and empty internal nodes. They also propose two refinements of their algorithm for reducing and avoiding the rule copy. Their proposal improves search performance without increasing the memory requirement.

Ternary content-addressable memories (TCAMs) are also typically used for high-speed packet classification in routers by comparing packet header against rules in the database. To solve the issue of TCAM's unsuitability for rules containing range fields, Bremler-Barr et al. [2] in their paper "Layered Interval Codes for TCAM-based Classification" propose a technique for efficient representations of range rules by taking into account that disjoint ranges can be encoded much more efficiently than overlapping ranges. The proposed scheme has been shown to significantly reduce the number of redundant TCAM entries caused by range rules.

As the Internet moves to IPv6 due to exhaustion of IPv4 addresses, a large number of IP lookup algorithms have been proposed in the literature. Smiljanic and Cica [3] in their paper "Scalable Lookup Algorithms for IPv6" analyze the scalability of the IP lookup algorithms and their applicability to IPv6. The memory requirements of the algorithms in terms of lookup table sizes have been compared. Recommendation on using the optimal algorithm has been included as part of the study.

Although packet filtering is a well-established scheme for detecting intrusions and maintaining security in the network, computation and space requirement dramatically grows with increasing complexity of filtering rules and security strategies. A critical task in packet filtering is managing the large filtering tables based on prefix and range-based search. Neji and Bouhoula [4] in their paper "A prefix-based approach for managing hybrid specifications in complex packet filtering" proposed a solution based on signed prefixes to incorporate ranges in the data structures used for filtering. To show the advantages of their scheme, the authors have analyzed the worst case time and space complexity as a function of the size of the filtering table and packet header size.

1389-1286/$ - see front matter © 2012 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.comnet.2012.07.004

Quality of Service schedulers ensure fair service to users by allocating bandwidth to packets depending on the bandwidth negotiated to the corresponding flows. This can require hardware implementation of the algorithms at the routers, resulting in a complexity growth. Moreover, the hardware implementations are specific to the algorithm. Sanli et al. [5] in their paper "Hardware Design and Implementation of Packet Fair Queuing Algorithms for the Quality of Service Support in the High-speed Internet" present a generalized framework for Packet Fair Queuing Schedulers. The framework has been tested by implementing one of the Packet Fair Queuing algorithms on an FPGA and measuring the performance with real traffic flows.

Small buffer sizes in routers result in lower queueing delays and shorter round-trip times which aids better control of the TCP congestion windows size. The use of the bandwidth-delay product rule to determine buffer sizes mandates larger buffers with increasing link capacities. Larger buffer sizes are difficult to implement in on-chip buffers. Francini [6] in their paper "Periodic Early Detection for Improved TCP Performance and Energy Efficiency" proposes a new active queue management scheme, called Periodic Early Detection (PED), with a view to reducing the buffer size in routers, while retaining TCP throughput and fairness. The algorithm is suitable for on-chip buffer implementations for link rates up to 100 Gbps.

To solve the bias against small flows in the Internet, active queue management has been proposed by Divakaran [7] in their paper "A Spike-Detecting AQM to deal with Elephants". In contrast to previous schemes, the proposed mechanism does not need to track flow sizes to resolve the bias. The main idea is to analyze traffic signatures to detect large spikes in flows and drop packets from the flows during congestion. Results show that the proposed scheme performs better than the widely studied Random Early Detection queue management scheme.

Papers [8,9] in the special issue, give an insight in the evolution of the switching fabrics to meet the requirements of the Future Internet. In today's high-performance packet switches, the switching functionality is implemented by power-hungry multi-board electronic switch fabrics. The replacement of electronic switch fabrics with hybrid electro-optical boards is a hot research topic. While optical technology seems immature to accomplish packet header processing or lookup operations, it is emerging as a promising option for pure memoryless switching functionalities inside the routers. Among all the options, the paper [8] analyzes the combination of tunable lasers and the wavelength routing capabilities of Arrayed Waveguide Gratings (AWG) to build optical switch fabrics. The scalability of AWG-based backplanes is limited by in-band crosstalk when several lasers are tuned at the same wavelength. This can be overcome by scheduling algorithms that can reduce the probability of reusing wavelength in different ports of the AWG device. In the paper "AWG-based Optical Switches Performance using Crosstalk Limiting ", Hermida et al. [8] extend several previous scheduling algorithms to enhance their performance and scalability.

The Main Distribution Frame (MDF) in telephone systems does the permutation of the circuits connecting the subscribe telephone lines to Digital Subscriber Line Access Multiplexers. In recent years, the design of rearrangeable MDFs has been challenged by the need of completing very frequent rearrangements due to volatility in the users' ISP selection. MDFs traditionally used non-blocking rearrangeable multistage switches. In the paper "Design and Control of Next Generation Distribution Frames", Cuda et al. [9] proposed the design of Automated Main Distribution Frame based on Non-Interruptive Rearrangeable networks to solve the issues arising from rearrangeable multistage networks which are known to temporarily interrupt active circuits.

Paper [10] in the special issue: "GMPLS-enabled MPLS-TP/PWE3 node with integrated 10Gbps tunable DWDM transponders: design and experimental evaluation" by Vilalta et al., reports implementation of an IP/MPLS over DWDM node architecture, with a GMPLS control plane. Authors present this node as an alternative to SONET/SDH to reduce cost, simplify operations and offer scalability advantages. The proposed solution uses the efficient aggregation and statistical multiplexing of packet transport technologies (i.e., MPLS-TP) to support IP and Ethernet services. The node implementation uses commercial off-the-shelf components, and the forwarding engine is based on open source software.

Finally, papers [11,12] end this special issue, reporting recent advances in the planning of high-capacity networks. In the paper "The Cost Optimal Solution of the Multi-Constrained Multicast Routing Problem", Molnar et al. [11] define the cost optimal solution of the multi-constrained multicast routing problem with a view to finding a multicast structure which spans a source node and a set of destination nodes, taking into account a set of constraints while minimizing a cost function. The optimization would be useful for multicast trees requiring quality of service guarantees. The authors approach can help to identify the exact optimal solution among a number of possible solutions. In the paper "Trading Availability among Shared-Protected Dynamic Connections in WDM Networks", Lucerna et al. [12] address the planning of WDM networks in a bandwidth-on-demand scenario, where the network operator uses the knowledge of holding time and availability target of connections to effectively operate shared-path protection. This is done by monitoring the downtime experienced by the connection and network state change. The approach offers flexibility in managing the availability provided to active connections during their holding times. The quality of service provided by the network has been evaluated in terms of availability and probability of violation of the availability target specified in the service level agreement.

The guest editors would like to thank all authors for submission of their manuscript to this special issue, as well as all paper reviewers for their thorough work. Finally, we like also to specially thank the Editor in Chief, Harry Rudin, for all his support and guidance during the editorial process.

References

[1] Hyesook Lim, Soohyun Lee, Earl E. Swartzlander, A new hierarchical packet classification algorithm, Computer Networks, 2012.
[2] Anat Bremler-Barr, David Hay, Danny Hendler, Layered interval codes for TCAM-based classification", Computer Networks, 2012.
[3] Aleksandra Smiljanic, Zoran Cica, Scalable lookup algorithms for IPv6, Computer Networks, 2012.
[4] Nizar Ben Neji, Adel Bouhoula, A prefix-based approach for managing hybrid specifications in complex packet filtering, Computer Networks, 2012.
[5] Mustafa Sanli, Ece Schmidt, Hasan Guran, Hardware design and implementation of packet fair queuing algorithms for the quality of service support in the high-speed Internet, Computer Networks, 2012.
[6] Andrea Francini, Periodic early detection for improved TCP performance and energy efficiency, Computer Networks, 2012.
[7] Dinil Mon Divakaran, A spike-detecting AQM to deal with elephants, Computer Networks, 2012.
[8] D. Ferńandez Hermida, M. Rodelgo Lacruz, A. Bianco, D. Cuda, G. Gavilanes Castillo, C. Ĺopez Bravo, F.J. Gonźalez Casta~no, AWG-based optical switches performance using crosstalk limiting, schedulers, Computer Networks, 2012.
[9] Davide Cuda, Paolo Giaccone, Massimo Montalto, Design and control of next generation distribution frames, Computer Networks, 2012.
[10] Ricard Vilalta, Raul Munoz, Ramon Casellas, Ricardo Martinez, Javier Vilchez, GMPLS-enabled MPLS-TP/PWE3 node with integrated 10 Gbps tunable DWDM transponders: design and experimental evaluation, Computer Networks, 2012.
[11] Mikĺos Molńar, Alia Bellabas, Samer Lahoud, The cost optimal solution of the multi-constrained multicast routing problem, Computer Networks, 2012.
[12] Diego Lucerna, Massimo Tornatore, Biswanath Mukherjee, Achille Pattavina, Trading availability among shared-protected dynamic connections in WDM networks, Computer Networks, 2012.

Pablo Pavon-Marino
Mohammed Atiquzzaman
Joan Garcia-Haro

Pablo Pavon-Marino received the M.Sc. degree from the Universidad de Vigo (Spain), in 1999, and the Ph.D. degree in telecommunications from the Universidad Politécnica de Cartagena (Spain), in 2004. He also received the M.Sc. degree in mathematics from the Universidad Nacional de Educación a Distancia (Spain), in 2010. He joined the Universidad Politécnica de Cartagena in 2000, and is currently an Associate Professor with the Department of Information Technologies and Communications.

He is author or co-author of more than 30 journal papers in the field of optimization and planning of electronic and optical communication networks. He is Technical Editor of the Optical Switching and Networking journal, and serves or has served in the technical program committee in several international conferences. He has served as General Chair of the IEEE High Performance Switching and Routing 2011 (IEEE HPSR 2011, Cartagena 2011) conference, and is General Chair of the International Conference on Transparent Optical Networks (ICTON 2013, Cartagena 2013). He is coordinating the FIERRO project (Future Internet: Efficiency in High-Speed Networks), which groups 21 institutions in Spain.

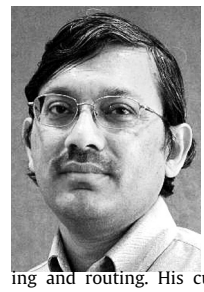

Mohammed Atiquzzaman received the B.Sc. degree in Electrical and Electronic Engineering in 1982, M.Sc. in Electrical Engineering in 1984 and Ph.D. in Electrical Engineering in 1987 from the University of Manchester, UK. He holds the Gaylord Presidential professorship of Computer Science at The University of Oklahoma where he teaches courses in Data Networks and Computer Architecture. His research interests and publications are in next generation computer networks, wireless and mobile networks, satellite networks, switching and routing. His current research activities are supported by the National Science Foundation (NSF), National Aeronautics and Space Administration (NASA), U.S. Air Force, Cisco, Honeywell and the Oklahoma Department of Transportation through grants totaling over $4 M.

He is the Editor-in-Chief of Journal of Network and Computer Applications and past co-editor-in-chief of Computer Communications. He is serving/served as editor of IEEE Communications Magazine, Journal of Wireless and Optical Communications, International Journal of Communication Systems, International Journal of Sensor Networks, International Journal of Communication Networks and Distributed Systems, and Journal of Real-Time Image Processing. He has Guest edited 12 special issues in various journals including IEEE Communications Magazine, and serves/served on the conference advisory board of over 10 international conferences, and in the technical program committee in over 60 conferences.

He is a senior member of IEEE and the recipient of the 2006 IEEE Communications Society Fred W. Ellersick Prize and the 2002 IET Premium paper Award. He is also the recipient of the 2007 Regents' Award for Superior Accomplishment in Research and Creative Activity.

Joan Garcia-Haro received the M.Sc. and Ph.D. degrees in telecommunication engineering from the Universitat Politecnica de Catalunya, Barcelona, Spain, in 1989 and 1995, respectively. He is currently a Professor with the Universidad Politecnica de Cartagena (UPCT). He is author or co-author of more than 70 journal papers mainly in the fields of switching, wireless networking and performance evaluation. From April 2002 to December 2004 he served as Editor-in-Chief of the IEEE Global Communications Newsletter, included in the IEEE Communications Magazine. He was Technical Editor of the same magazine from March 2001 to December 2011. He also holds an Honorable Mention for the IEEE Communications Society Best Tutorial paper Award (1995). He has been a visiting scholar at Queen's University at Kingston, Canada (1991–1992) and at Cornell University, Ithaca, USA (2010–2011).