Guest Editorial
Deep Packet Inspection: Algorithms, Hardware, and Applications

Ying-Dar Lin, *Fellow, IEEE*, Po-Ching Lin, Viktor K. Prasanna, *Fellow, IEEE*, H. Jonathan Chao, *Fellow, IEEE*, and John W. Lockwood, *Member, IEEE*

DEEP packet inspection (DPI) examines the content in packet payloads to search for signatures of network applications, signs of malicious activities, and leaks of sensitive information, rather than just examine packet headers for information such as IP addresses and port numbers. The inspection provides network devices with rich information of application protocol messages in packet payloads, and enables them to make intelligent decisions in packet processing based on the information. Therefore, the network devices equipped with the capability of DPI can provide numerous functions, such as network intrusion detection, traffic classification and content-aware policy control of network traffic, which will be otherwise much restricted if only packet headers are known.

DPI is inherently challenging due to the need to handle ever-increasing number of signatures and the diversity of application protocol messages. The signatures to be inspected must be also flexible and robust enough to resist possible evasion when facing the adversary of network attacks. Furthermore, the solutions usually should operate in real time in a high-speed network, while dealing with the above complexity. As a result, we believe that DPI still deserves careful study in depth, even though it has been studied for longer than a decade [1] and simultaneously searching a byte stream for thousands of patterns or even more at multi-giga bits per second is feasible in many state-of-the-art designs.

We received a total of 39 submissions, and selected 13 quality papers for publication after two rounds of reviews. The papers are organized into the following four sections: (1) Scalable Algorithms and Architectures for DPI, (2) Network Traffic Analysis with DPI, (3) Network Protocol Identification with DPI, and (4) Network Security Analysis with DPI.

It is essential that the algorithms for DPI should be scalable to accommodate a large number of signatures in limited memory space, and the architectures should be also scalable to operate in a high-speed network. The first section covers five papers to address the scalability issue.

Reading multiple bytes per memory access can potentially achieve high speed in DPI, but is also likely to suffer from the memory explosion problem. The paper "Kangaroo: Accelerating String Matching by Running Multiple Collaborative Finite State Machines" by Xiaofei Wang, Bin Liu, Junchen Jiang, Yang Xu, Yi Wang, and Xiaojun Wang presents a matching scheme that splits the original rule set into multiple sub-rule sets and partitions the input stream with a window of fixed size. The scheme scans multiple bytes in parallel using a compact state machine, while keeping the memory usage down at the same magnitude as a conventional single-byte scheme.

The size of deterministic finite automata (DFA) is usually exponential in the number of regular expressions (RegExes). The paper "Towards Fast and Optimal Grouping of Regular Expressions via DFA Size Estimation" by Tingwen Liu, Alex X. Liu, Jinqiao Shi, Yong Sun, and Li Guo presents an algorithm to estimate the DFA size for a given RegEx set without constructing the DFA, and a grouping algorithm which is much faster and memory efficient than conventional DFA construction.

The paper "TFA: A Tunable Finite Automaton for Pattern Matching in Network Intrusion Detection Systems" by Yang Xu, Junchen Jiang, Rihua Wei, Yang Song, and H. Jonathan Chao introduce an automaton representation of regular expressions to resolve the problems of state explosion and unpredictable performance in deterministic and non-deterministic finite automata (i.e., DFA and NFA) representation. This representation allows a limited number of concurrent active states, and the required memory space is significantly reduced.

The paper "Revisiting State Blow-up: Automatically Building Augmented-FA while Preserving Functional Equivalence" by Xiaodong Yu, Bill Lin and Michela Becchi is also intended to avoid state explosion in conventional DFA construction. The proposed method features several advantages, such as limited worst-case processing time and coverage of arbitrary regular expressions.

Deploying DPI systems in an ISP backbone network is challenging because of the high data rate. The paper "A Scalable Carrier-Grade DPI System Architecture using Synchronization of Flow Information" by NamUk Kim, GanHo Choi, and JaeHyeong Choi introduces a highly scalable DPI architecture based on the concepts of flow information synchronization and

Y.-D. Lin is with the Department of Computer Science, National Chiao Tung University, Hsinchu 30010, Taiwan (e-mail:ydlin@cs.nctu.edu.tw).

P.-C. Lin is with the Department of Computer Science and Information Engineering, National Chung Cheng University, Chiayi 62102, Taiwan (e-mail: pclin@cs.ccu.edu.tw).

V. K. Prasanna is with the Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: prasanna@usc.edu).

H. J. Chao is with theDepartment of Electrical and Computer Engineering, New York University, Brooklyn, NY 11201 USA (e-mail: chao@nyu.edu).

J. W. Lockwood is with the Algo-Logic Systems, Santa Clara, CA 95050 USA (e-mail:jwlockwd@algo-logic.com).

Digital Object Identifier 10.1109/JSAC.2014.2371093

adaptive traffic control. The system is manufactured to operate at 40 Gbps with four 10 Gbps DPI modules.

Network analysis with DPI usually requires various underlying supports to isolate low-level functions from application developers. The second section covers two papers which present the frameworks for flow tracking, packet reassembly and application protocol parsing. Existing network traffic capture frameworks provide applications with raw packets and do not deal with complex operations such as flow tracking and TCP stream reassembly. The paper "Stream-Oriented Network Traffic Capture and Analysis for High-Speed Networks" by Antonis Papadogiannakis, Michalis Polychronakis and Evangelos P. Markatos presents the stream capture library, called scap, which can deliver flow-level statistics and reassembled streams to user applications. This design will facilitate obtaining application protocol messages for developers.

Application protocol parsing and field extraction is essential to understand the semantics of packet payloads, and is also an indispensable part in network traffic analysis. The paper "High-Speed Application Protocol Parsing and Extraction for Deep Flow Inspection" by Alex X. Liu, Chad Meiners, Eric Norige and Eric Torng presents an automated online method based on two models: counting regular grammars and counting automata. The models have the ability to facilitate parsing and extracting fields from context sensitive application protocols.

Traffic classification or network protocol identification is an essential part for content-aware network management. The third section covers three papers about protocol identification. Fine-grained traffic identification tells not only what network application generates a certain packet, but also what application function or user behavior generate the packet. The paper "MP-ROOM: Optimal Matching On Multiple PDUs for Fine-Grained Traffic Identification" by Hao Li and Chengchen Hu presents a method to split the identification rules into fields in a matching order, and construct a hierarchical layered matching tree, which reduces the space complexity and increases throughput for matching over multiple packet payloads.

The paper "Toward Unsupervised Protocol Feature Word Extraction" by Zhuo Zhang, Zhibin Zhang, Patrick P. C. Lee, Yunjie Liu, and Gaogang Xie designs an unsupervised method to systematically and efficiently extract protocol feature words. The feature words can be used to distinguish application protocols and are building blocks of payload analysis.

The paper "Efficient Methods for Early Protocol Identification" by Béla Hullár, Sándor Laki, and András Gyärgy presents a protocol identification method which inspects only the first few bytes of the first (or first few) packet(s) of each flow. The piece of information is analyzed by machine-learning-based methods with very low computational complexity, and high early classification accuracy is demonstrated on traffic traces from a diverse set of applications.

Network security analysis is traditionally an important application of DPI. The fourth section covers three papers in this respect. The paper "A Reconfigurable Platform and Programming Tools for High-Level Network Applications demonstrated as a Hardware Honeypot" by Sascha Mühlbach and Andreas Koch presents the NetStage platform, which allows rapid deployment of FPGA-accelerated attack-resilient interactive communication applications. This platform does not have software-programmable processors, and the FPGA configuration port is completely isolated from network traffic. The authors demonstrated the usage of the platform with a hardware honeypot.

It is important to detect infected hosts from the overwhelming alerts of intrusion detection systems. The paper "IDS Alert Correlation in the Wild with EDGe" by Elias Raftopoulos and Xenofontas Dimitropoulos presents an information theoretic measure to identify statistically significant temporal associations between the alerts, and detects infected hosts that exhibit recurring multi-stage behavior.

A fast-flux service network is widely adopted by attackers to hide hosts behind flux bots. The paper "Detect Fast-Flux Domains through Response Time Differences" by Fu-Hau Hsu, Chuan-Sheng Wang, Chi-Hsien Hsu, Chang-Kuo Tso, Li-Han Chen, and Song-Hui Lin presents a lightweight fast-flux domain detector based on fluctuation of the response time of a sequence of DNS requests with low false positive and false negative rates.

We would like to thank all the authors, who contribute their excellent work to this special issue, and the reviewers for their invaluable review comments. We also appreciate Prof. Muriel Medard, the Editor-in-Chief, Prof. Alberto Leon-Garcia, Laurel Greenridge, and Lauren Briede for their assistance during the entire process of the special issue, as well as the support from the other members in the JSAC Editorial Board. Finally, we thank the Editorial Staff at IEEE for the efforts in the production of the issue.

REFERENCES

[1] P. C. Lin, Y. D. Lin, Y. C. Lai, and T. H. Lee, "Using string matching for deep packet inspection," *IEEE Comput.*, vol. 41, no. 4, pp. 23–28, Apr. 2008.

Ying-Dar Lin (M'95–SM'06–F'13) received the Ph.D. degree in computer science from the University of California, Los Angeles, CA, USA, in 1993. He is a Distinguished Professor of Computer Science at National Chiao Tung University, Hsinchu, Taiwan. He served as the CEO of Telecom Technology Center during 2010–2011 and as a Visiting Scholar at Cisco Systems, San Jose, CA, during 2007–2008. Since 2002, he has been the Founder and Director of Network Benchmarking Lab (NBL, www.nbl.org.tw), which reviews network products with real traffic. NBL recently became an approved test laboratory of the Open Networking Foundation (ONF). He is a Research Associate of ONF. He also cofounded L7 Networks Inc. in 2002, which was later acquired by D-Link Corp. He published a textbook, *Computer Networks: An Open Source Approach* (www.mhhe.com/lin), with R.-H. Hwang and F. Baker (McGraw-Hill, 2011). It is the first text that interleaves open source implementation examples with protocol design descriptions to bridge the gap between design and implementation. His research interests include design, analysis, implementation, and benchmarking of network protocols and algorithms, quality of services, network security, deep packet inspection, wireless communications, embedded hardware/software co-design, and, recently, software-defined networking. His work on "multi-hop cellular" was the first along this line and has been cited over 600 times and standardized into IEEE 802.11s, IEEE 802.15.5, WiMAX IEEE 802.16j, and 3GPP LTE-Advanced. He is an IEEE Distinguished Lecturer (2014–2015). He is currently on the Editorial Boards of the IEEE TRANSACTIONS ON COMPUTERS, *IEEE Computer*, *IEEE Network*, *IEEE Communications Magazine—Network Testing Series*, IEEE WIRELESS COMMUNICATIONS, IEEE COMMUNICATIONS SURVEYS AND TUTORIALS, IEEE COMMUNICATIONS LETTERS, *Computer Communications*, *Computer*

Networks, *Journal of Network and Computer Applications*, and *IEICE Transactions on Information and Systems*. He has guest-edited several special issues in IEEE journals and magazines and co-chaired symposia at IEEE Globecom'13 and IEEE ICC'15.

Po-Ching Lin received the B.S. degree in computer and information education from National Taiwan Normal University, Taipei, Taiwan, in 1995 and the M.S. and Ph.D. degrees in computer science from National Chiao Tung University, Hsinchu, Taiwan, in 2001 and 2008, respectively. He was a Visiting Scholar with Prof. V. Paxson at the International Computer Science Institute, University of California, Berkeley, CA, USA, from 2007 to 2008 and became a Senior Engineer at the Networks and Multimedia Institute, Institute for Information Industry, from 2008 to 2009. In August 2009, he joined the faculty of the Department of Computer and Information Science, National Chung Cheng University, Chiayi, Taiwan, where he is currently an Assistant Professor. His research interests include network security, network traffic analysis, and performance evaluation of network systems. He was on the Program Committee of IEEE ICC'08 and IEEE Globecom'13.

Viktor K. Prasanna (M'84–SM'91–F'96) received the B.S. degree in electronics engineering from Bangalore University, Bangalore, India, the M.S. degree from the Indian Institute of Science, Bangalore, and the Ph.D. degree in computer science from the Pennsylvania State University, State College, PA, USA. He is a Charles Lee Powell Chair in Engineering and a Professor of electrical engineering and computer science at the University of Southern California (USC), Los Angeles, CA, USA, and serves as the Director of the Center for Energy Informatics. He is the Executive Director of the USC-Infosys Center for Advanced Software Technologies. He is an Associate Member of the Center for Applied Mathematical Sciences. He leads the integrated optimizations efforts at the USC-Chevron Center of Excellence for Research and Academic Training on Interactive Smart Oilfield Technologies (CiSoft) at USC and the demand response optimizations in the LA Smartgrid project. His research interests include high performance computing, parallel and distributed systems, reconfigurable computing, cloud computing, and embedded systems. He has published extensively and consulted for industries in the above areas. He is a Fellow of the Association for Computing Machinery and the American Association for Advancement of Science. He is the Steering Committee Co-Chair of the IEEE International Parallel and Distributed Processing Symposium (IPDPS) [merged IEEE International Parallel Processing Symposium and IEEE International Symposium on Parallel and Distributed Processing]. He is the Steering Committee Chair of the IEEE International Conference on High Performance Computing. In the past, he has served on the Editorial Boards of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS, IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, *Journal of Pervasive and Mobile Computing*, and the PROCEEDINGS OF THE IEEE. He serves on the Editorial Boards of the *Journal of Parallel and Distributed Computing* and the *ACM Transactions on Reconfigurable Technology and Systems*. During 2003–2006, he was the Editor-in-Chief of the IEEE TRANSACTIONS ON COMPUTERS. He was the Founding Chair of the IEEE Computer Society Technical Committee on Parallel Processing. He currently serves as the Editor-in-Chief of the *Journal of Parallel and Distributed Computing*. He was a recipient of the 2005 Okawa Foundation Grant and the Outstanding Engineering Alumnus Award from the Pennsylvania State University in 2009. He was a recipient of Best Paper Awards at several international forums, including ACM Computing Frontiers, IEEE IPDPS, IEEE International Conference on Parallel and Distributed Systems, International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), International Conference on Parallel and Distributed Computing and Systems, IEEE International Conference on High Performance Switches and Routers, among others.

H. Jonathan Chao (S'83–M'84–SM'95–F'01) received the B.S. and M.S. degrees in electrical engineering from National Chiao Tung University, Hsinchu, Taiwan, and the Ph.D. degree in electrical engineering from Ohio State University, Columbus, OH, USA.

He is the Department Head and Professor of electrical and computer engineering at Polytechnic Institute of New York University, Brooklyn, NY, USA, where he joined in January 1992. He has also served as a consultant for various companies, such as Huawei, Lucent, NEC, and Telcordia.

During 2000–2001, he was Co-Founder and CTO of Coree Networks, NJ, USA, where he led a team in the implementation of a multi-terabit Multi-Protocol Label Switching (MPLS) switch router with carrier-class reliability. From 1992 to 1999, he taught short courses three times a year in the subjects of SONET, ATM, IP, MPLS, and switch/router designs, to industry people through UC Berkeley and Oxford University's continuing education programs. From 1985 to 1992, he was a Member of Technical Staff at Telcordia, where he was involved in transport and switching system architecture designs and ASIC implementations, such as the world's first SONET-like framer chip, ATM layer chip, sequencer chip (the first chip handling packet scheduling), and ATM switch chip. From 1977 to 1981, he was a Senior Engineer at Telecommunication Labs of Taiwan performing circuit designs for a digital telephone switching system. He has coauthored three networking books, namely, *Broadband Packet Switching Technologies A Practical Guide to ATM Switches and IP Routers* (Wiley, 2001), *Quality of Service Control in High-Speed Networks* (Wiley, 2001), and *High-Performance Switches and Routers* (Wiley, 2007). He holds 44 patents with 10 pending and has published more than 200 journal and conference papers. He has been doing research in the areas of network designs in data centers, terabit switches/routers, network security, network on chip, and biomedical devices.

Prof. Chao is a Fellow of the IEEE for his contributions to the architecture and application of VLSI circuits in high-speed packet networks. He was a recipient of the Telcordia Excellence Award in 1987 and a co-recipient of the 2001 Best Paper Award from the IEEE TRANSACTION ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY.

John W. Lockwood (S'89–M'96) received the M.S., B.S., and Ph.D. degrees from the University of Illinois at Urbana-Champaign, IL, USA. He designs and implements networking algorithms in reconfigurable logic. He is the Founder and CEO at Algo-Logic Systems, Santa Clara, CA, USA. Algo-Logic's hardware-accelerated logic circuits enable networks to achieve ultra-low latency processing while carrying large volumes of data. From 2006 to 2009, he managed the NetFPGA program as a Consulting Associate Professor at Stanford University. At Stanford, he grew the worldwide deployment of NetFPGA hardware from 10 to 1021 units, managed the relationships with the corporate sponsors, presented the NetFPGA Developer Workshop and week-long training events at Stanford University and SIGCOMM, and hosted international NetFPGA events in the U.K., Europe, India, China, and Australia. Prior to joining Stanford in January 2007, he led the Reconfigurable Network Group, which was a part of the Applied Research Laboratory, Washington University in Saint Louis. At Washington University, he was a tenured Associate Professor at the Department of Computer Science and Engineering. He and his research group developed the Field programmable Port Extender (FPX) to enable rapid prototype of extensible network modules in Field Programmable Gate Array (FPGA) technology. He has served as the Principal Investigator on grants from the National Science Foundation, Xilinx, Altera, Agilent Technologies, Nortel Networks, Rockwell Collins, and Boeing. He has worked in industry for AT&T Bell Laboratories, IBM, Science Applications International Corporation, and the National Center for Supercomputing Applications. He served as a Co-Founder of Global Velocity, a networking start-up company focused on high-speed data security. He has published over 100 papers on the following topics: intrusion detection and prevention, Internet packet forwarding and classification, content-based routing and filtering, Bloom filters, TCP/IP flow processing, buffering and queuing, soft-core liquid architectures, thermal management of reconfigurable hardware, data classification and clustering, reconfigurable hardware platforms, and dynamic reconfiguration of FPGA platforms. He is a member of ACM, Tau Beta Pi, and Eta Kappa Nu.