

Network Coding

By **BAOCHUN LI**, *Senior Member IEEE*
Guest Editor

YUNNAN WU
Guest Editor

As the Internet grows in scale since its inception, information bits have been packetized and then delivered from sending to receiving hosts by having routers and switches store and forward them, much like how the postal system delivers mail. After decades of technological advances in transport mechanisms, addressing and naming schemes, routing protocols, scheduling disciplines, and security in the Internet, the basic “store and forward” idea remains the same.

In 2000, a group of information theorists at the Chinese University of Hong Kong has published a seminal paper, titled “Network information flow,” to challenge this basic idea of storing and forwarding packets in routers in a communication network. The paper was the first to propose that an intermediate node in the network can forward any functions of incoming packets, not necessarily limited to copies of these packets. For example, an intermediate node can “mix” incoming packets using their linear combinations. The simplest way of mixing them is to forward an outgoing packet that is the XOR of incoming packets. The conventional wisdom in the Internet was that packets need to be delivered to receiving hosts intact and without any changes, as if they are commodities in a commodity flow, such as cars on highways. The new idea is that bits in information flows do not have to be delivered as commodities; they can be mixed however we wish, as long as the receiving hosts have received sufficient “evidence” or “clues” to reconstruct the original packets from the sending host.

This new idea is called *network coding*, and it is as exciting as it can be: it has been shown in early research that the use of network coding can help to maximize information flow rates in multicast communication sessions (from one sending host to multiple receiving hosts). Since transmitting evidence

about packets can be more useful than delivering the packets in their original form, can we substantially improve the efficiency of utilizing bandwidth in the current Internet by performing coding?

In the decade since 2000, there has been a large number of high-quality research papers—hundreds every year—in the literature on the topic of network coding, attracting the interests from a diverse group of researchers and practitioners in information theory, communications, and computer networking communities. With a substantial amount of research in network coding, work in the literature can largely be divided into two categories: earlier research is more focused on understanding fundamental theoretical properties and limits of network coding, and starting from 2003, more recent research is more focused on its practical challenges, implications, and implementations. It has become crystal clear that network coding is based on a solid theoretical foundation, yet amenable to a very realistic potential towards practical applications, in both wireless networking and the Internet, such as distributed storage and peer-assisted content distribution. There exists a tremendous potential for the theory of network coding to affect the design of next-generation network protocols.

The gradual shift from more theoretical investigations to more practical concerns has demonstrated that

Network coding means that bits in information flows do not have to be delivered as commodities; they can be mixed however we wish, as long as the receiving hosts have received sufficient “evidence” or “clues” to reconstruct the original packets from the sending host.

network coding research has reached a level of maturity, and it may be a good time to revisit and summarize some of the most prominent research results in network coding, from both theoretical and practical perspectives. It would also be a good time to introduce some of the technical advances on the topic of network coding to a wider engineering audience, so that the knowledge learned and expertise accumulated in the past decade may become useful to pave the way for its practical deployment. These reasons are precisely what motivated the organization of this Special Issue, and we believe that it is an excellent match to the broader interests of readers of the PROCEEDINGS OF THE IEEE.

With ten papers in this Special Issue covering a wide range of representative topics in both theoretical and practical aspects of network coding, some of the most prominent results in this area have been presented.

The Special Issue opens with a unique piece authored by Yeung, one of the coauthors of the first paper on network coding. In his interesting and concise piece, titled “Network coding: A historical perspective,” Yeung has offered a historical perspective of network coding as one of its coinventors, and elaborated on how the idea of network coding has come to fruition in the years leading to the publication of the first paper on this topic. It has been pointed out that research on network coding was motivated by a subtle change of problem formulation: instead of asking what communication system is needed for a given set of communication requests, the problem was formulated by asking what communication requests can be supported by a given communication system—in this case a network. The paper shows how network coding can outperform routing in a simple network topology, now known as the *Butterfly Network*.

The next four papers in this Special Issue cover theoretical advances in the area of network coding. A key milestone in the development of network

coding is the discovery that the maximum multicast communication capacity in a network can be achieved by using only linear network coding. Linear coding has many attractive properties. From a theoretical standpoint, linearity is a beautiful algebraic property supported by a rich and profound mathematical foundation. From an engineering standpoint, the simplicity of linear approaches made it possible to practically deploy network coding. Li *et al.*'s paper titled “Linear network coding: Theory and algorithms,” reviews the fundamentals of linear network coding and construction algorithms for optimal linear network codes. In addition, this paper also presented some deep connections between linear network coding and algebra.

Whereas the theory of network coding has been well understood for a single communication session, network coding for multiple concurrent communication sessions turned out to be a grand open theoretical challenge. In particular, linear network coding is no longer optimal for multi-session network communications. Dougherty *et al.*'s paper, titled “Network coding and matroid theory,” presents a number of intriguing examples that demonstrate the insufficiency of linear coding and reveal the inherent difficulties of multisession network coding. These examples were constructed elegantly by exploring the connection between network coding and matroid theory—a branch of mathematics that generalizes linear algebra and graph theory.

Errors may occur in a communication network with the use of network coding. Such errors may be random errors, erasures due to lost packets, or more seriously, errors caused by intentional attacks by malicious nodes in the network. Zhang's tutorial paper, titled “Theory and applications of network error correction coding,” presents an overview of the basic ideas proposed to design *network error correction codes*, which extend classical error correction coding in the time domain to new classes of codes in the space domain.

Security is another important subject that deserves significant attention from both theoretical and practical angles. In Cai and Chan's paper, titled “Theory of secure network coding,” the basic theory of linear secure network coding is presented with a rigorous, information-theoretic treatment. The paper explores the fundamental limit for confidential communication in networks in the presence of malicious eavesdroppers.

Five papers in the second half of this Special Issue cover important research topics towards the application of network coding with more practical perspectives. In wireless networks, Nazar and Gastpar's paper, titled “Reliable physical layer network coding,” provides extensive coverage of using the idea of network coding in the physical layer, where interference from signals can be treated and taken advantage of as a network code, in order to improve throughput in a variety of wireless network models, starting from a simple scenario involving a two-way relay channel.

Fragouli's paper, titled “Network coding: Beyond throughput benefits,” presents an interesting perspective that bridges network coding theory with its practical applications, especially when compared with routing without coding. It has shown that network coding is beneficial in highly dynamic networks that changes rapidly over time, and has illustrated a number of examples on how network coding can be practically used in a simple and decentralized fashion, and what solutions exist to mitigate disadvantages of network coding from a practical standpoint.

In Dimakis *et al.*'s paper, titled “A survey on network codes for distributed storage,” a new application scenario is proposed in which the application of network coding is beneficial: distributed storage systems where data reliability is maintained using erasure coding. In such distributed storage systems, if one of the storage nodes fails, its data need to be repaired by accessing data

from healthy nodes. Depending on the type of repairs to be made, different network codes may be designed to best meet their specific requirements.

Next, Sundararajan *et al.*'s paper, titled "Network coding meets TCP: Theory and implementation," seeks to integrate the application of network coding into the most commonly adopted transport layer protocol in the Internet, the transmission control protocol (TCP). TCP uses the lack or reordering of acknowledgment packets to indicate errors, and to induce retransmission from the sending host. If network coding is applied, a new mechanism needs to be designed: rather than acknowledging the receipt of packets in their original form, the receiving host can communicate the degree of freedom in the batch of coded packets it has received so far to the sending host. The work presented makes it feasible to apply network coding in an incremental fashion in the Internet, without any changes to TCP.

Finally, Li and Niu's paper, titled "Random network coding in peer-to-

peer networks: From theory to practice," introduces important milestones in the past decade of network coding research towards bringing theoretical benefits of network coding to practical systems. In the context of peer-to-peer networks, peers are end hosts at the edge of the Internet, and can implement network coding without any changes to the Internet infrastructure. It has been pointed out that applications that involve peer-to-peer networks, such as file sharing and video streaming, may become the most promising scenario for network coding to be deployed in real-world systems. It shows protocols used to apply network coding in a large-scale on-demand video streaming system that has already been deployed in a production setting, and discusses the corresponding practical design and implementation challenges.

Looking forward, we believe that research in network coding, from both theoretical and practical perspectives, will see sustained or even accelerating interests from both research communities of information theory and communication networks.

The Chinese University of Hong Kong has recently established the Institute of Network Coding, which envisions and is working towards new technological advances of network coding in the Internet. We hope that papers in this Special Issue may help convince our readers that network coding may find practical applications in a wide variety of communication networks and systems, from cloud computing systems to mobile applications. With this issue, we wish to draw attention to highlights in the past decade of network coding research, and to pave the way towards future theoretical and practical developments on this topic.

In closing, we would like to express our sincere appreciation to J. Calder, Managing Editor, J. Sun and M. Meyer, our Publication Editors of the PROCEEDINGS OF THE IEEE, for their assistance and support throughout the preparation of this Special Issue. We would like to thank the authors and reviewers for their participation and contributions.

Enjoy the fun and benefits of mixing packets along the way! ■

ABOUT THE GUEST EDITORS

Baochun Li (Senior Member, IEEE) received the B.Eng. degree in computer science from the Department of Computer Science and Technology, Tsinghua University, China, in 1995 and the M.S. and Ph.D. degrees in computer science from the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, in 1997 and 2000, respectively.

Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Professor. He holds the Nortel Networks Junior Chair in Network Architecture and Services from October 2003 to June 2005, and the Bell University Laboratories Endowed Chair in Computer Engineering since August 2005. His research interests include large-scale multimedia systems, cloud computing, peer-to-peer networks, applications of network coding, and wireless networks.

Dr. Li was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems, in 2000. In 2009, he was a recipient of the Multimedia Communications Best Paper Award from the IEEE Communications Society, and a recipient of the University of Toronto McLean Award. He is a member of the Association for Computing Machinery (ACM).



Yunnan Wu received the Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ, in January 2006.

He is currently a Software Engineer at Facebook, Inc., Palo Alto, CA. From August 2005 to March 2010, he was a Researcher at Microsoft Corporation, Redmond, WA. His research interests include multimedia, networking, storage systems, and information theory. He has been doing research on network coding since 2003.

Dr. Wu has presented tutorials on network coding at the 2007 IEEE Conference on Computer Communications and ACM SIGMETRICS conferences. He has co-organized special sessions on network coding at the 2006 Conference on Information Sciences and Systems and the 2009 Asilomar Conference. He has served in the Technical Program Committee of the annual network coding workshop since 2008. He has served as Program Co-Chair of the Second Workshop on Wireless Network Coding, in conjunction with the 2009 IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (IEEE SECON 2009).

