Network Coding

By YUNNAN WU

BAOCHUN LI

The internet delivers information by having routers store and forward packets, in a way similar to how the postal system delivers mail. Such a forwarding model has been deeply entrenched in our minds. There seemed hardly any doubt about it.

Around the year 2000, a group of information theorists—Rudolph Ahlswede, Ning Cai, Shuo-Yen Robert Li, and Raymond W. Yeung—challenged this model. An intermediate node in a network receives some packets and sends out some packets. Fundamentally, the node can output any functions of the received packets, not necessarily copies of the received packets. Ahlswede *et al.* demonstrated that it is possible to do better than forwarding, by mixing information (e.g., outputting a packet that is the XOR of two received packets). This is now called *network coding*. This groundbreaking work ignites an enormous amount of enthusiasm. It is time to revolutionize network communications!

Through this past decade, significant progress has been made toward a fundamental understanding and practical developments of network coding. Network coding has been successfully applied to a number of networking scenarios. These scenarios demonstrate many advantages of network coding over forwarding, including resource (e.g., bandwidth, energy, storage space) efficiency, computational efficiency, lower communication delay, and robustness to network dynamics. In the following we will walk through some of these scenarios, briefly explain the technology, and talk about the benefits network coding offers.

Ahlswede *et al.*'s work showed that for multicast (one-to-many) communication over a network, it is possible to achieve a higher throughput via network coding. Among the usage scenarios for network coding, multicast communication is best understood. To use network coding for multicast communication, the basic technology is to let each node generate output packets by taking random linear combinations of the received packets. Such a multicast strategy is drastically different from the conventional forwarding approach. In the forwarding approach, the delivery system needs to ensure that every receiver gets one copy of each source packet, like collecting enough distinct coupons to redeem a reward. Achieving that efficiently in a dynamic network is difficult. In contrast, in the network coding approach, random mixture packets flow in the network, each essentially providing an equation in terms of the source packets. A receiver can reconstruct the source packets upon collecting enough linearly independent packets. The network coding approach thus offers advantages in terms of delivery efficiency, ease of management, and resilience to network dynamics such as packet losses, node arrivals and departures.

Network coding, therefore, is potentially a good fit for applications with one-to-many or many-to-many communication patterns. There are many Internet applications with such communication patterns, such as live media broadcast, media on demand, stock quotes distribution, file sharing, gaming, conferencing, chat rooms, or

Digital Object Identifier: 10.1109/JPROC.2010.2044847

distance learning. Here are two concrete network coding–based systems for the Internet. A peer-to-peer file downloading protocol based on network coding, called Avalanche, has been developed at Microsoft Research, Cambridge. Live trials with several hundred nodes confirm that the protocol incurs little CPU and I/O overhead and results in smooth and fast downloads. Network coding is also used in a commercial large-scale peer-to-peer video on demand system of UUSee, an Internet streaming media company based in Beijing, China.

Network coding is especially useful for wireless networks, since its advantages match very well with the broadcast nature and other inherent characteristics of the wireless medium. For information multicast over multihop wireless networks, a network coding based solution has been shown to offer superior performance compared to state-of-the-art forwarding based solutions. Even for unicast (one-to-one) communications, network coding offers an attractive solution because its robustness can better deal with the dynamics in wireless networks. This is especially the case for highly dynamic and/or challenging network environments, such as mobile networks (e.g., vehicular networks, ad hoc networks formed in the battlefield), sensor networks where nodes sleep most of the time to conserve energy, and networks with sporadic connectivity.

In addition to efficiently supporting end-to-end communications in a wireless network, network coding has also been successfully applied in the link layer to efficiently perform relaying. Specifically, a wireless router opportunistically mixes packets passing by it, which may belong to different end-to-end communication sessions, to reduce the transmissions in the air. Such use of network coding is different from the random mixing used for end-to-end multicasting, where only packets belonging to the same end-to-end session can be mixed together. In this form of network coding, a wireless router examines the packets it needs to forward and identifies opportunities to mix the outgoing packets; the mixing is done carefully so that the mixture packets can be decoded at the immediate neighbors. Prof. Dina Katabi's research group at MIT reported throughput increases of up to fourfold using this technique in a wireless mesh test-bed.

The network coding schemes discussed so far mix the information packets: each packet is either lost or received free of error. Some recent physical layer communication techniques explore mixing at the signal level. By mixing at the modulator/channel coder or exploiting the inherent mixing of signals offered by the wireless medium, further performance gains can be achieved. For example, with a technique called analog network coding, two messages can be exchanged between two nodes with the help of a relay node in two timeslots. These techniques could potentially be useful for supporting the relaying feature for cellular technologies (e.g., the 3GPP LTE Advanced standard).

In addition to communications, storage is another important application area where network coding is expected to have significant impact. Various forms of error and/or erasure correction codes have long been used in storage systems, in memory systems, disks, RAID systems, etc., to protect against bit errors and component failures. Network coding, as a special form of coding, is a valuable addition to the portfolio of coding techniques for storage systems. For example, network coding is shown to be useful in significantly reducing the network traffic needed to repair an erasure-coded storage system from node failures. In this application, network coding is used to regenerate content when nodes fail, maximizing the overall protection strength under network traffic constraints. It is therefore useful in scenarios where the network bandwidth is a critical resource, such as distributed storage systems.

In the above, we have discussed several scenarios where network coding can provide benefits. In the past decade, the topic of network coding has attracted a substantial amount of research attention in both information theory and computer networking. It is based on a solid theoretical foundation, yet amenable to a very realistic potential towards practical applications, in wireless networking, the Internet, and storage systems, as we see from above. There exists a tremendous potential for the theory of network coding to affect the design of next-generation network protocols. Of course, network coding is not a panacea, and a successful application of it may require an in-depth understanding of the application properties and the strengths and weaknesses of network coding. The topic is still in active research and we look forward to seeing the full potential of network coding unleashed. Enjoy the fun and surprises of mixing! ■