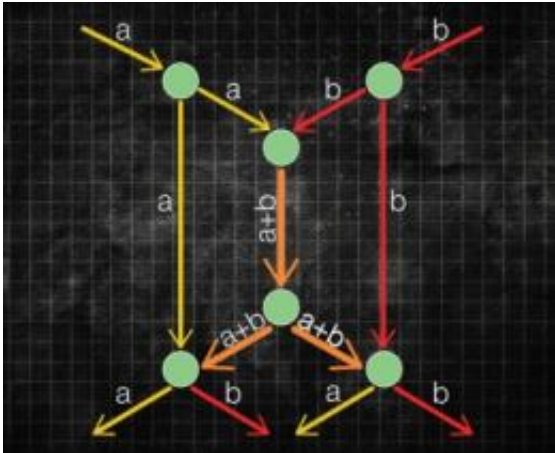


# The elusive capacity of data networks

May 15 2012, by Larry Hardesty

---



The principle behind network coding is often explained by reference to a so-called butterfly network. When messages A and B reach the same node, they're scrambled together, and their combination ( $A+B$ ) is passed to the next node. Further downstream, one node uses A to recover B from  $A+B$ , while another uses B to recover A from  $A+B$ . Graphic: Christine Daniloff

In its early years, information theory — which grew out of a [landmark 1948 paper](#) by MIT alumnus and future professor Claude Shannon — was dominated by research on error-correcting codes: How do you encode information so as to guarantee its faithful transmission, even in the presence of the corrupting influences engineers call "noise"?

Recently, one of the most intriguing developments in information theory has been a different kind of coding, called [network](#) coding, in which the question is how to encode information in order to maximize the capacity

of a network as a whole. For information theorists, it was natural to ask how these two types of coding might be combined: If you want to both minimize error and maximize capacity, which kind of coding do you apply where, and when do you do the decoding?

What makes that question particularly hard to answer is that no one knows how to calculate the data capacity of a network as a whole — or even whether it can be calculated. Nonetheless, in the first half of a [two-part paper](#), which was published recently in *IEEE Transactions on Information Theory*, MIT's Muriel Médard, California Institute of Technology's Michelle Effros and the late Ralf Koetter of the University of Technology in Munich show that in a wired network, network coding and error-correcting coding can be handled separately, without reduction in the network's capacity. In the forthcoming second half of the paper, the same researchers demonstrate some bounds on the capacities of wireless networks, which could help guide future research in both industry and academia.

A typical data network consists of an array of nodes — which could be routers on the Internet, wireless base stations or even processing units on a single chip — each of which can directly communicate with a handful of its neighbors. When a packet of data arrives at a node, the node inspects its addressing information and decides which of several pathways to send it along.

## Calculated confusion

With network coding, on the other hand, a node scrambles together the packets it receives and sends the hybrid packets down multiple paths; at each subsequent node they're scrambled again in different ways. Counterintuitively, this can significantly increase the capacity of the network as a whole: Hybrid packets arrive at their destination along multiple paths. If one of those paths is congested, or if one of its links

fails outright, the packets arriving via the other paths will probably contain enough information that the recipient can piece together the original message.

But each link between nodes could be noisy, so the information in the packets also needs to be encoded to correct for errors. "Suppose that I'm a node in a network, and I see a communication coming in, and it is corrupted by noise," says Médard, a professor of electrical engineering and computer science. "I could try to remove the noise, but by doing that, I'm in effect making a decision right now that maybe would have been better taken by someone downstream from me who might have had more observations of the same source."

On the other hand, Médard says, if a node simply forwards the data it receives without performing any error correction, it could end up squandering bandwidth. "If the node takes all the signal it has and does not whittle down his representation, then it might be using a lot of energy to transmit noise," she says. "The question is, how much of the noise do I remove, and how much do I leave in?"

In their first paper, Médard and her colleagues analyze the case in which the noise in a given link is unrelated to the signals traveling over other links, as is true of most wired networks. In that case, the researchers show, the problems of error correction and network coding can be separated without limiting the capacity of the network as a whole.

## **Noisy neighbors**

In the second paper, the researchers tackle the case in which the noise on a given link is related to the signals on other links, as is true of most wireless networks, since the transmissions of neighboring base stations can interfere with each other. This complicates things enormously: Indeed, Médard points out, information theorists still don't know how to

quantify the capacity of a simple three-node wireless network, in which two nodes relay messages to each other via a third node.

Nonetheless, Médard and her colleagues show how to calculate upper and lower bounds on the capacity of a given wireless network. While the gap between the bounds can be very large in practice, knowing the bounds could still help network operators evaluate the benefits of further research on network coding. If the observed bit rate on a real-world network is below the lower bound, the operator knows the minimum improvement that the ideal code would provide; if the observed rate is above the lower bound but below the upper, then the operator knows the maximum improvement that the ideal code might provide. If even the maximum improvement would afford only a small savings in operational expenses, the operator may decide that further research on improved coding isn't worth the money.

"The separation theorem they proved is of fundamental interest," says Raymond Yeung, a professor of information engineering and co-director of the Institute of [Network Coding](#) at the Chinese University of Hong Kong. "While the result itself is not surprising, it is somewhat unexpected that they were able to prove the result in such a general setting."

Yeung cautions, however, that while the researchers have "decomposed a very difficult problem into two," one of those problems "remains very difficult. ... The bound is in terms of the solution to another problem which is difficult to solve," he says. "It is not clear how tight this bound is; that needs further research."

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT research, innovation and teaching.*

Provided by Massachusetts Institute of Technology

Citation: The elusive capacity of data networks (2012, May 15) retrieved 27 April 2024 from <https://phys.org/news/2012-05-elusive-capacity-networks.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.