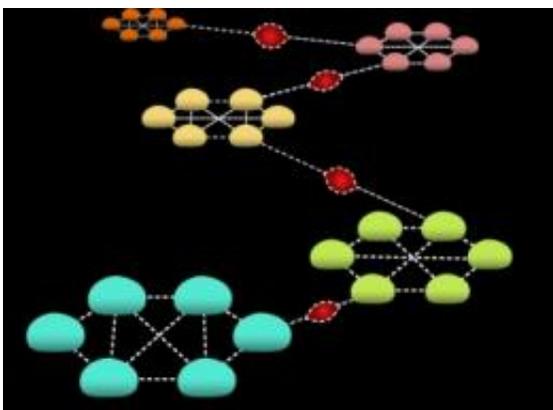


New algorithm enables much faster dissemination of information through self-organizing networks

January 11 2011, Larry Hardesty



A new algorithm spreads information (red) much more efficiently in networks characterized by sparse connections between densely interlinked clusters.
Graphic: Christine Daniloff

As sensors that do things like detect touch and motion in cell phones get smaller, cheaper and more reliable, computer manufacturers are beginning to take seriously the decade-old idea of "smart dust" -- networks of tiny wireless devices that permeate the environment, monitoring everything from the structural integrity of buildings and bridges to the activity of live volcanoes. In order for such networks to make collective decisions, however -- to, say, recognize that a volcano is getting restless -- they need to integrate information gathered by hundreds or thousands of devices.

But networks of cheap [sensors](#) scattered in punishing and protean environments are prone to “bottlenecks,” regions of sparse connectivity that all transmitted data must pass through in order to reach the whole [network](#). At the 2011 ACM-SIAM Symposium on Discrete Algorithms, which took place in New Orleans the weekend of Jan. 9, Keren Censor-Hillel, a postdoc at MIT’s Computer Science and Artificial Intelligence Laboratory, and Hadas Shachnai of Technion – Israel Institute of Technology presented a new algorithm that handles bottlenecks much more effectively than its predecessors.

The algorithm is designed to work in so-called ad hoc networks, in which no one device acts as superintendent, overseeing the network as a whole. In a network of cheap wireless sensors, for instance, any given device could fail: its battery could die; its signal could be obstructed; it could even be carried off by a foraging animal. The network has to be able to adjust to any device’s disappearance, which means that no one device can have too much responsibility.

Without a superintendent, the network has no idea where its bottlenecks are. But that doesn’t matter to Censor-Hillel and Shachnai’s algorithm. “It never gets to identify the bottlenecks,” Censor-Hillel says. “It just copes with their existence.”

Consistent inconsistency

The researchers’ analysis of their algorithm makes a few simplifying assumptions that are standard in the field. One is that communication between networked devices takes place in rounds. Each round, a device can initiate communication with only one other device, but it can exchange an unlimited amount of [information](#) with that device and with any devices that contact it. During each exchange, it passes along all the information it’s received from any other devices. If the devices are volcano sensors, that information could be, say, each device’s most

recent measurement of seismic activity in its area.

It turns out that if you're a sensor in a network with high connectivity — one in which any device can communicate directly with many of the others — simply selecting a neighboring device at random each round and sending it all the information you have makes it likely that every device's information will permeate the whole network. But take two such highly connected networks and connect them to each other with only one link — a bottleneck — and the random-neighbor algorithm no longer works well. On either side of the bottleneck, it could take a long time for information to work its way around to the one device that communicates with the other side, and then a long time for that device to bother to send its information across the bottleneck.

Censor-Hillel and Shachnai's algorithm works by alternating communication strategies from round to round. In the first round, you select a neighboring device at random and send it all the information you have — which, since it's the first round, is limited to the measurement that you yourself have performed. That same round, however, other devices may contact you and send you their information. In the second round, you don't just select a neighbor at random; you select a neighbor whose information you have not yet received. In the third round, you again select a neighbor at random. By the end of that round, since every device on the network forwards all the information it has, you've received not only the measurements performed by the devices you contacted, nor just the measurements performed by the devices that contacted you, but measurements performed by neighbors of your neighbors, and even neighbors of neighbors of neighbors. In the fourth round, you again select a device whose information you haven't received; in the fifth, you select a device at random; and so on.

“The idea is that the randomized steps I take allow me to spread the information fast within my well-connected subset,” says Censor-Hillel.

But in the alternate rounds, each device tracks down the devices it hasn't heard from, ensuring that information will quickly reach all the devices, including those that communicate across the bottleneck.

According to Alessandro Panconesi, a professor of computer science at Sapienza University of Rome and an expert on network analysis, the devices on ad hoc networks tend to have limited computational power and battery life, so the algorithms they execute must be very “lightweight.” The new algorithm is “an interesting contribution,” Panconesi says. “It’s a very simple, locally based algorithm. Essentially, a node in this network can wake up and start operating by using this algorithm, and if every node in the network does the same, then essentially you give communication capability to the entire network.” He points out that the current version of the algorithm, in which, every round, every device sends all the information it’s received, wouldn’t be practical: “The algorithm is very expensive in terms of the information that it needs to exchange.” But he believes that developing a less bandwidth-intensive version is “not unlikely.” “I’m optimistic,” he says.

Censor-Hillel agrees that “a major thing for future work would be to actually get practical bandwidth.” But for the time being, she’s collaborating with assistant professor of applied mathematics Jonathan Kelner and with CSAIL grad students Bernhard Haeupler and Petar Maymoukov to develop an [algorithm](#) that performs even better in the idealized case of unlimited bandwidth. “It’s a major improvement,” she says.

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

More information: Paper: “[Fast Information Spreading in Graphs](#)”

[with Large Weak Conductance”](#)

Provided by Massachusetts Institute of Technology

Citation: New algorithm enables much faster dissemination of information through self-organizing networks (2011, January 11) retrieved 26 April 2024 from <https://phys.org/news/2011-01-algorithm-enables-faster-dissemination-self-organizing.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.