


# Sharing the air

September 22 2009, by Larry Hardesty

---



In a wireless network, different transmission frequencies work better for different users. That's because the same transmission reaches each user along several different paths; at one frequency, the signals might reinforce each other, while at another frequency, they might cancel each other out. Photo - Christine Daniloff

(PhysOrg.com) -- In the old days, when a new wireless technology came along, it got its own swath of the electromagnetic spectrum: AM radio uses 535 to 1,605 kilohertz, so television got chunks between 54 and 806 megahertz. But the airwaves are getting so crowded that that approach won't work anymore. MIT researchers in the lab of Dina Katabi, an associate professor of electrical engineering and computer science, are teaching wireless technologies how to share what spectrum is left.

Giving each technology its own [frequency band](#) is intrinsically inefficient. In areas where a particular wireless service is underused, or

where use varies throughout the day, swaths of spectrum can sit idle for minutes or hours at a time. Historically, technical limitations meant that no alternative was practical. But improvements in computer processors, radio hardware, and signal-processing techniques have raised the possibility of devices that can look for unused spectrum and exploit it without stepping on each other's toes.

Regulatory bodies like the [Federal Communications Commission](#) are unlikely to grant additional technologies access to previously allocated spectrum anytime soon. But spectrum sharing could have immediate implications for the so-called white spaces — the frequency bands vacated when television moved from analog to digital. In the United States, the FCC has agreed to leave those bands unlicensed, at least for now, and a coalition of technology companies that includes Google, Microsoft, and Intel hopes to use them for high-speed data connections for portable devices — wherever they are. Technologies that want to use the white spaces, however, will have to show that they won't interfere with each other, or with devices already authorized to use the same spectrum. One advantage of the MIT researchers' work is that it takes such a general approach to the problem of spectrum sharing that it should work with most existing wireless data devices — and others yet unimagined.

According to Katabi, spectrum sharing poses two distinct problems. The first is figuring out which transmission channels in a given area are unoccupied. The second is deciding how to use the available channels efficiently.

At last year's Sigcomm, generally considered the major international conference in the field of networking, Katabi and her colleagues addressed the first question. Traditionally, says Katabi, wireless technologies trying to avoid each other would simply measure the power in a certain frequency band: high power meant that the band was in use,

low power meant that it wasn't. But "the fact that there is power in a particular frequency does not mean that you cannot use it, necessarily," says Katabi. Different transmitters might be able to use the same frequency, for instance, if their intended receivers are far enough apart. "The opposite is also not true," Katabi says. "The fact that certain frequencies do not have power does not mean that you can use them, because if you use them, you could potentially leak power to nearby frequencies." A radio transmitter uses filters to concentrate power into specific frequency bands, but the filters never work perfectly.

So Katabi and her colleagues propose that, instead of looking at the amount of power in a frequency band, wireless devices look at the changing power profiles of other devices sharing the same spectrum. Most wireless devices will cut their transmission rates if they encounter congestion. By tracking power over time, the MIT system determines whether a particular choice of frequency is forcing other devices to slow down.

## **Choosing a path**

On Wednesday, at the Mobicom mobile-computing and -networking conference in Beijing, Hariharan Rahul, a graduate student in Katabi's lab, presents a solution to the second problem. When a wireless technology has its own small allocation of spectrum, the frequencies it can use are close enough together that their performance will be roughly the same. But in a swath of spectrum shared by multiple technologies, the unoccupied frequencies may be far apart. As a consequence, they could have very different performance. That's because the same transmission reaches the user along several different paths: some signals travel directly, while others might first bounce off the ground or the walls of buildings. At one frequency, signals might reinforce each other; at another frequency, they might cancel each other out.

The MIT system provides an efficient way to determine which unoccupied frequencies work best for which users. Most emerging wireless technologies use a technique called orthogonal frequency division multiplexing (OFDM) to increase data transmission rates. OFDM requires senders to include in their transmissions known bit patterns, called pilots, to help a receiver to compensate for channel distortion. By analyzing how channel noise affect those patterns, the MIT system determines how well a given frequency will work for a given user and calculates the optimal transmission rate for each frequency. Since the OFDM devices are sending each other those bit patterns anyway, the new system imposes little additional burden on them.

Katabi and Rahul, who worked for Akamai Technologies before coming to MIT, implemented the new system in the lab, on a network that operates in the Wi-Fi spectrum. Data transmission rates on the network more than tripled.

Spectrum sharing in the white spaces is particularly amenable to Katabi and Rahul's approach. Because digital TV uses spectrum so efficiently, television stations broadcasting over the airwaves no longer need all of the bandwidth allotted them. The result is a host of unused [frequency](#) bands between television channels. Because the unlicensed bands are spaced so far apart, they're likely to exhibit the variable performance that Katabi and Rahul's system takes into account.

"It's very important; it's good stuff," says Anant Sahai, an assistant professor of electrical engineering and computer sciences at the University of California, Berkeley, who specializes in spectrum sharing. "I can see how this kind of thinking is going to be important in the white spaces." Sahai adds, however, that Katabi and Rahul's work is at the "protocol level" — the level of the transmission scheme — and that implementing it in the white spaces will require similar innovation in

hardware. Nonetheless, he says, "what's very encouraging about their work is that they've actually put together an implementation to test it out."

Provided by Massachusetts Institute of Technology ([news](#) : [web](#))

Citation: Sharing the air (2009, September 22) retrieved 27 April 2024 from <https://phys.org/news/2009-09-air.html>

<p>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.</p>
--