

Toward practical compressed sensing

Massachusetts Institute of Technology

The last 10 years have seen a flurry of research on an emerging technology called compressed sensing. Compressed sensing does something that seems miraculous: It extracts more information from a signal than the signal would appear to contain. One of the most celebrated demonstrations of the technology came in 2006, when Rice University researchers produced images with a resolution of tens of thousands of pixels using a camera whose sensor had only one pixel.

Compressed sensing promises dramatic reductions in the cost and power consumption of a wide range of imaging and signal-processing applications. But it's been slow to catch on commercially, in part because of a general skepticism that sophisticated math ever works as well in practice as it does in theory. Researchers at MIT's Research Laboratory of Electronics (RLE) hope to change that, with a new mathematical framework for evaluating compressed-sensing schemes that factors in the real-world performance of hardware components.

"The people who are working on the theory side make some assumptions that circuits are ideal, when in reality, they are not," says Omid Abari, a doctoral student in the Department of Electrical Engineering and Computer Science (EECS) who led the new work. "On the other hand, it's very costly to build a circuit, in terms of time and also money. So this work is a bridge between these two worlds. Theory people could improve algorithms by considering circuit nonidealities, and the people who are building a chip could use this framework and methodology to evaluate the performance of those algorithms or systems. And if they see their potential, they can build a circuit."

Mixed reviews

In a series of recent papers, four members of associate professor Vladimir Stojanovic's Integrated Systems Group at RLE — Abari, Stojanovic, postdoc Fabian Lim and recent graduate Fred Chen — applied their methodology to two applications where compressed sensing appeared to promise significant power savings. The first was spectrum sensing, in which wireless devices would scan the airwaves to detect unused frequencies that they could use to increase their data rates. The second was the transmission of data from wireless sensors — such as electrocardiogram (EKG) leads — to wired base stations.

At last year's International Conference on Acoustics, Speech, and Signal Processing, the researchers showed that, alas, in spectrum detection, compressed sensing can provide only a relatively coarse-grained picture of spectrum allocation; even then, the power savings are fairly meager.

But in other work, they argue that encoding data from wireless sensors may be a more natural application of the technique. In a forthcoming paper in the journal *IEEE Transactions on Circuits and Systems*, they show that, indeed, in the case of EKG

Toward practical compressed sensing

Published on Electronic Component News (<http://www.ecnmag.com>)

monitoring, it can provide a 90 percent reduction in the power consumed by battery-powered wireless leads.

The reason the Rice camera could get away with a single-pixel sensor is that, before striking the sensor, incoming light — the optical signal — bounced off an array of micromirrors, some of which were tilted to reflect the signal and some of which weren't. The pattern of "on" and "off" mirrors was random and changed hundreds or even thousands of times, and the sensor measured the corresponding changes in total light intensity. Software could then use information about the sequence of patterns to reconstruct the original signal.

Ups and downs

The applications the RLE researchers investigated do something similar, but rather than using mirrors to modify a signal, they use another signal, one that alternates between two values — high and low — in a random pattern. In the case of spectrum sensing, the frequency of the input signal is so high that mixing it with the second signal eats up much of the power savings that compressed sensing affords.

Moreover, the time intervals during which the second signal is high or low should be of precisely equal duration, and the transition from high to low, or vice versa, should be instantaneous. In practice, neither is true, and the result is the steady accumulation of tiny errors that, in aggregate, diminish the precision with which occupied frequencies can be identified.

An EKG signal, however, is mostly silence, punctuated by spikes every second or so, when the heart contracts. As a consequence, the circuitry that mixes it with the second signal can operate at a much lower frequency, so it consumes less power.

Abari, however, says he hasn't given up on applying compressed sensing to spectrum sensing. A new algorithm called the sparse Fast Fourier Transform, [developed at MIT](#) [1], would modify the signal in the spectrum-sensing application in a way that offsets both the loss of resolution and the increase in power consumption. Abari is currently working with EECS professor Dina Katabi, one of the new algorithm's inventors, to build a chip that implements that algorithm and could be integrated into future compressed-sensing systems.

Source URL (retrieved on 12/21/2013 - 12:17am):

<http://www.ecnmag.com/news/2013/02/toward-practical-compressed-sensing>

Links:

[1] <http://web.mit.edu/newsoffice/2012/faster-fourier-transforms-0118.html>