From *SIAM News*, Volume 40, Number 8, October 2007

Compressed Sensing Creates a Buzz at ICIAM ’07

By Justin Romberg

Digital signal processing has become so pervasive that its impact would be hard to overstate. Converting data into digital format gives us almost unlimited flexibility in how we process it; as computation has gotten cheaper, DSP has become ever more widespread. Artifacts of the “DSP revolution” that has taken place over the last 30 years surround us: Any DVD player, iPod, or cell phone has an embedded DSP chip running algorithms that are the end-product of decades of research in applied mathematics.

The nascent field of compressed sensing, showcased at ICIAM 2007 in three special sessions and a plenary talk by Emmanuel Candès, gives us a fresh look at the fundamentals of DSP. Rather than looking for better ways to process the digital data, CS concentrates on the actual acquisition process, the translation of the signal into digital format. By looking at the acquisition process (“sampling” or “sensing”) from a mathematically abstract standpoint, research in CS has uncovered something surprising about the best way to sample signals.

The story of compressed sensing starts with classical methods for data compression. A common trick for compressing a set of digital data is to first “transform” the data as a superposition of a known set of basis signals. If the basis is chosen judiciously, the transform coefficients will be sparse—the salient information will be concentrated in a few large terms. Photograph-like images containing millions of pixels, for example, can be very sparse in the wavelet domain, where only a small percentage of the transform coefficients are significant. This sparsity structure allows us to throw away the vast majority of the transform coefficients, while suffering very little distortion. Roughly speaking, a transform coder operates by executing this change of basis, sifting through the coefficients for the important ones, and then carefully encoding their indices and values.

The traditional framework for capturing and then transmitting or storing a signal of interest is to sample it, possibly at very high rates, and then apply a digital transform coder to turn the large stream of sample values into a much smaller stream of coded transform coefficients. It seems wasteful, though, in terms of hardware cost, power, and complexity, to take a large number of samples only to produce a considerably smaller amount of compressed data. Com-pressed sensing avoids this waste by, as the name implies, integrating compression into the sampling process itself.

To do this, we need to broaden our notion of what it means to “sample” a signal. Whereas traditional sampling consists of simply recording the value of a signal at a discrete set of points, CS devices correlate the incoming signal against a set of known *test signals*, and record these correlations. (see Figure 1.) Now the question is: Which test signals should we use to minimize the number of measurements we have to take? It is tempting to match the test signals to the signal structure by using basis functions from our transform. We know, after all, that the signal is sparse in the transform domain, and so we should be able to capture it with a small number of measurements. But unless we know a priori *which* components will be active, and typically we do not, the fact that the signal is sparse does not help us at all.

The solution to this dilemma has a surprising twist: Correlating against a series of *random* waveforms allows us to simultaneously figure out which transform coefficients of the signal are important and what their actual values are. In doing this, a CS device is effectively mixing all of the important components of the signal together in a variety of ways. The measurement sequence will not look like the signal at all—in fact, it will look like noise. But subtly embedded in these measurements is all of the critical information about the signal.

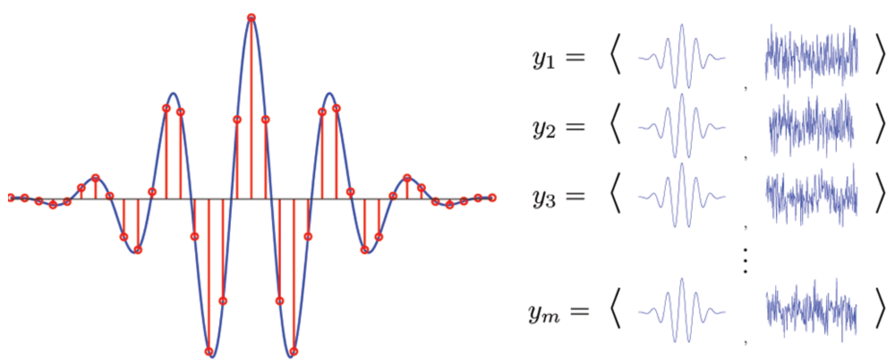

Figure 1. *Traditional sampling (left) turns a signal into a discrete list of numbers by simply evaluating it at a discrete set of points. Compressed sensing (right) turns a signal into a discrete list of numbers by correlating it with a series of random waveforms.*

The last ingredient for compressed sensing is an algorithm that will tease the signal out of the random measurements. This is where sparsity enters the picture. While many signals can explain the measurements we observed, only one signal with sparse structure can account for them. To find this signal, we solve an optimization program that searches for the sparsest signal (quantified using the ℓ_1 norm: the sum of the magnitudes of the transform coefficients) that could have produced the measurements in hand.

How many measurements do we need for this reconstruction procedure to work? It turns out that the number is roughly proportional (smaller, in practice, by a factor of about five) to the number of active components in the signal we are sensing. For a megapixel image that can be closely approximated with 20,000 wavelet coefficients, this means that we can recover the coefficients (and hence something close to the image) from 100,000 random measurements—a reduction by a factor of 10 compared with individual measurements of each of the 1 million pixels.

Compressed sensing casts a new light on the acquisition process. Instead of taking samples, we are *encoding* the signal: The compression and

the sensing are all part of the same process. The reconstruction algorithm (the optimization program) is thus in some sense trying to *decode* the measurements to recover the transform coefficients.

The development of CS is one of those rare instances in which a completely unintuitive suggestion from abstract mathematics turns out to work: An effective way to measure a structured signal is to correlate it with random noise. As we saw at ICIAM, this mathematical advance is directly influencing the design of next-generation sensors. The applications being explored include novel imaging devices that can capture high-resolution images with a single photodetector, high-resolution radar systems made from relatively inexpensive hardware, extremely low-powered cameras, and new analog-to-digital converters that can capture signals that contain extremely high frequencies.

Justin Romberg, organizer of three ICIAM '07 minisymposia on compressed sensing, is a professor of electrical and computer engineering at Georgia Institute of Technology.

Invited speaker Emmanuel Candès ("Compressive Sampling") discussed the novel sensing or sampling theory that allows the faithful recovery of signals or images from far fewer measurements than required with traditional methods.