Seismic imaging simulation

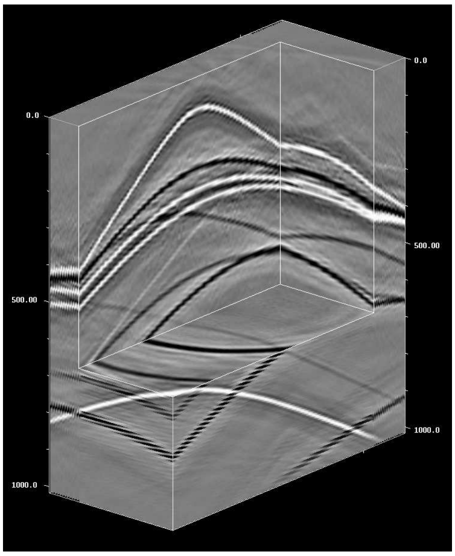

(a) Estimated
(16x faster, SNR=9.6 dB).

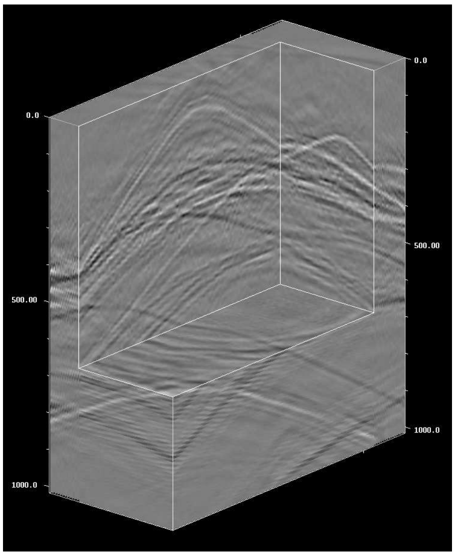

(b) Estimation error
(Figure 2b minus 5(a))

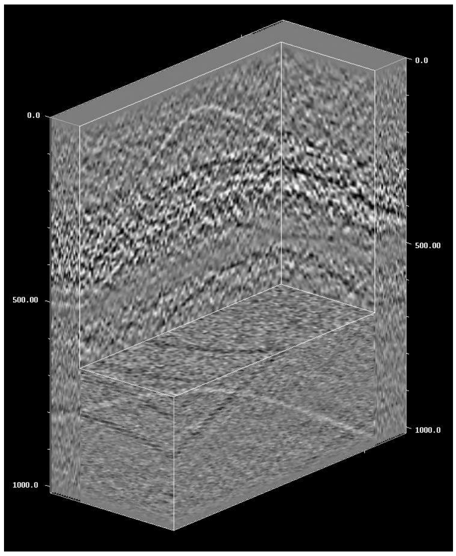

(c) Cross-correlation
estimate.

- Result produced with $16\times$ "compression" in the computations
- Can even take this example down to $32\times$

Randomly modulated integration

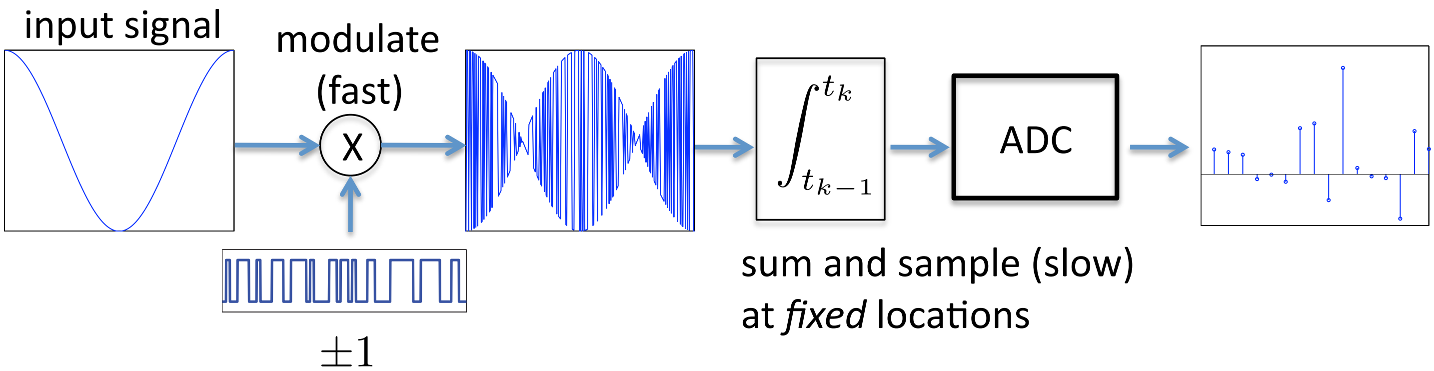

- Uses a standard "slow" ADC preceded by a "fast" binary mixing
- Mixing circuit much easier to build than a "fast" ADC
- In each sampling interval, the signal is summarized with a random sum
- Sample rate $\sim$ total *active* bandwidth

Random modulated integration in time and frequency

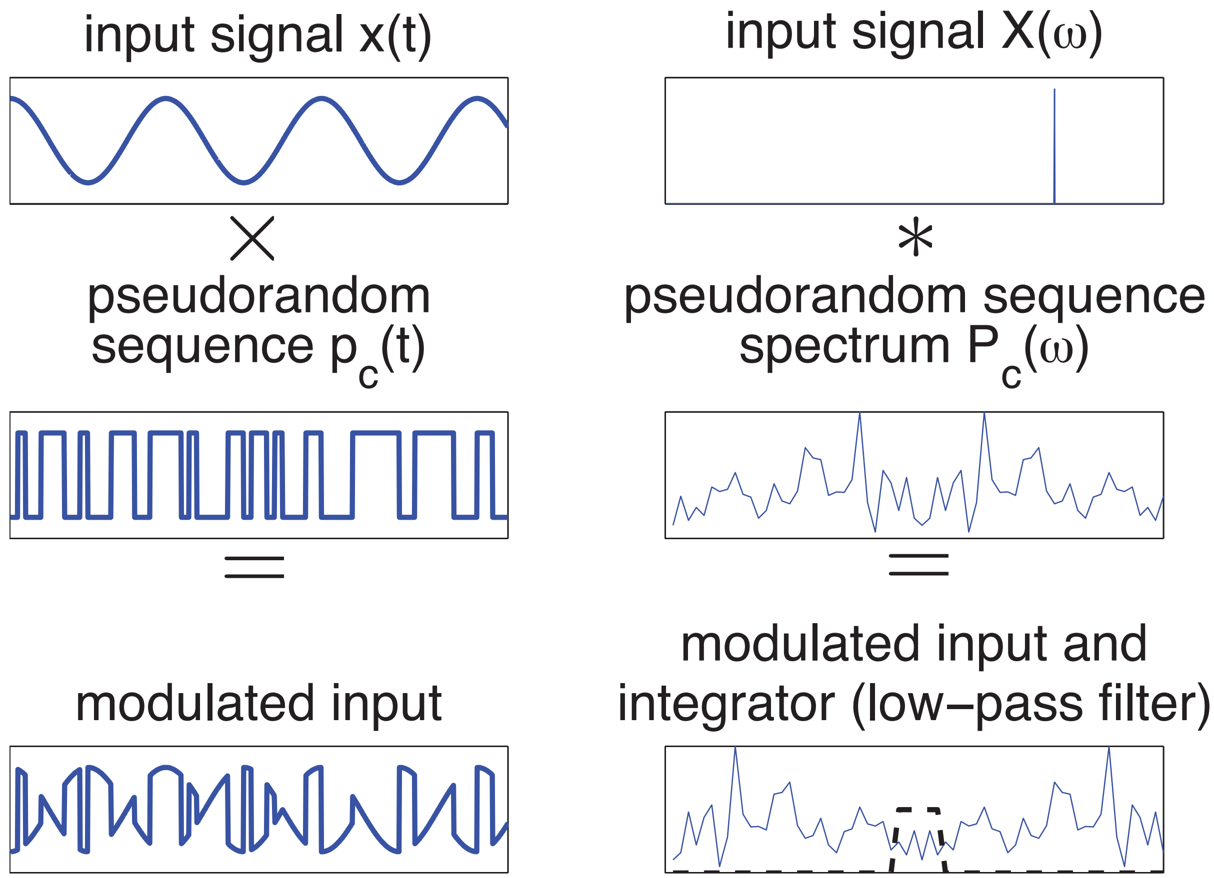

Multichannel modulated integration

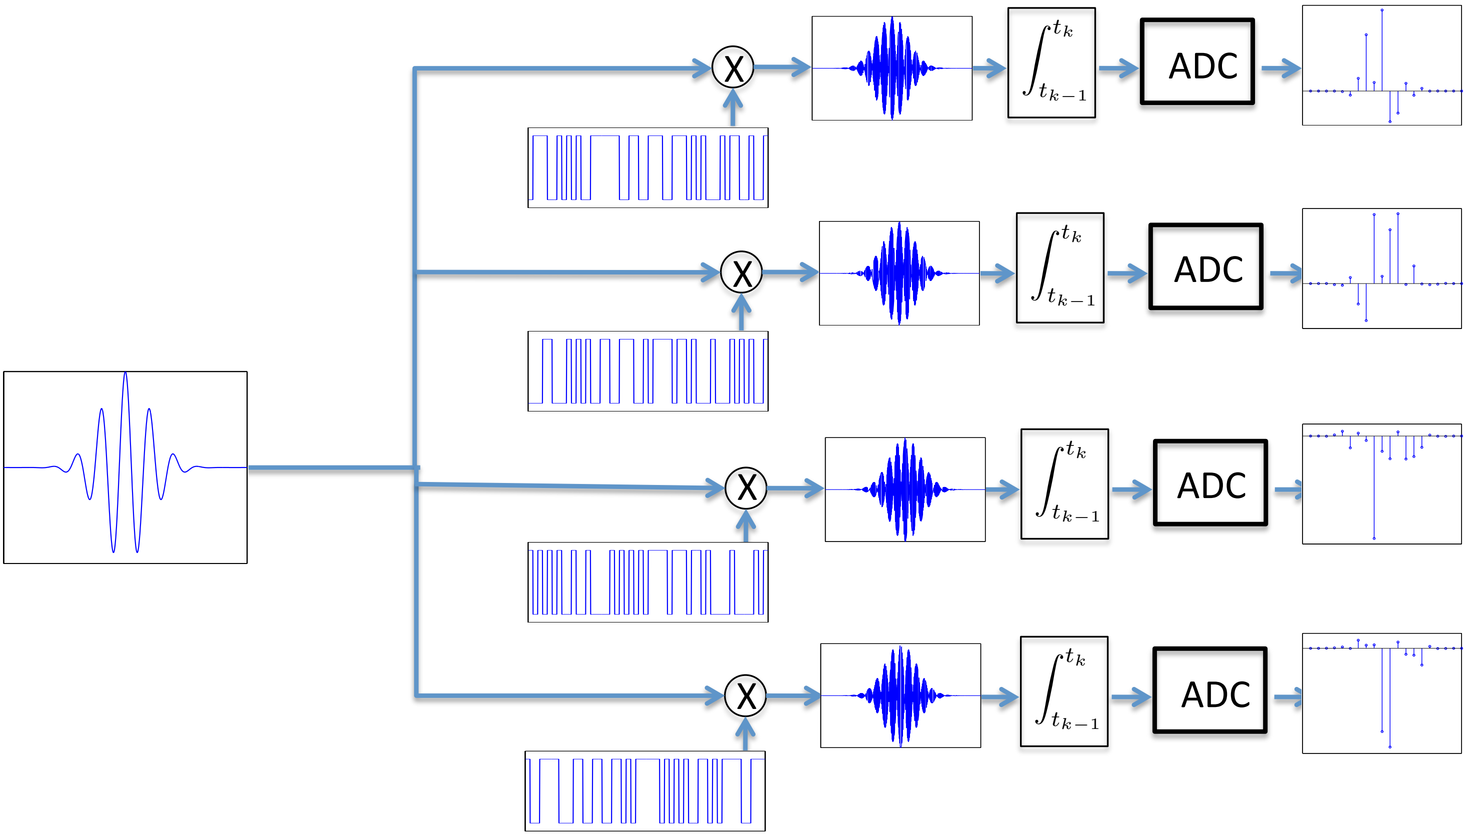

This architecture is being implemented as part of DARPA's Analog-to-Information program

Analog-to-digital converter state-of-the-art

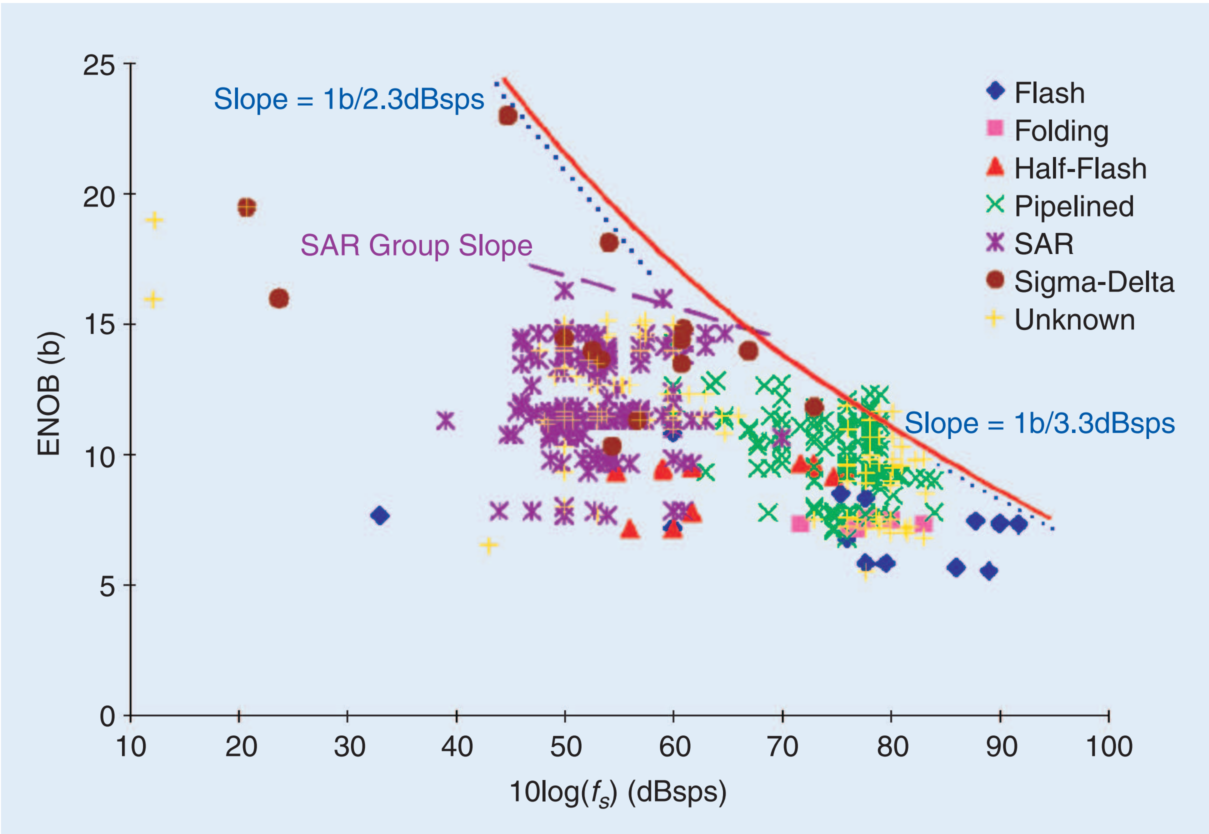

The bad news starts at 1 GHz (Le et al '05)

Analog-to-digital converter state-of-the-art

From 2008...

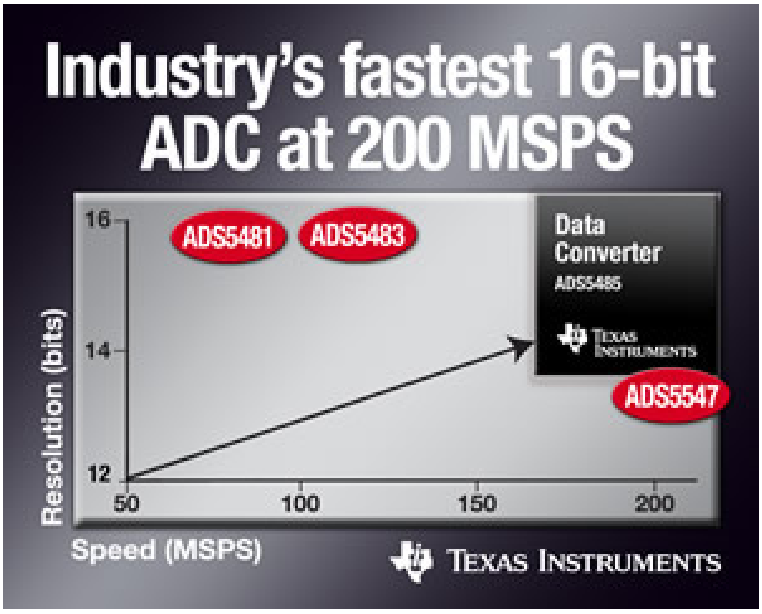

(Lots of RF signals have components in the 10s of gigahertz...)

Spectrally sparse RF signals

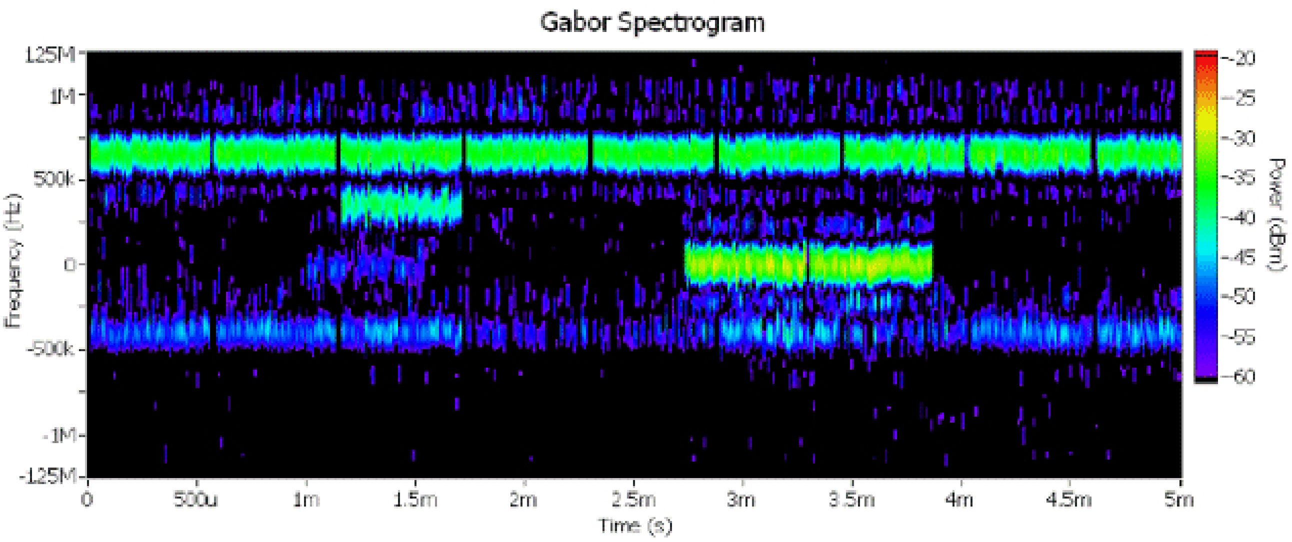

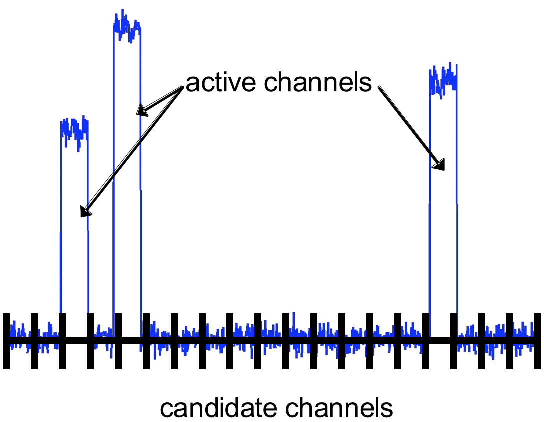

Randomly modulated integration receiver

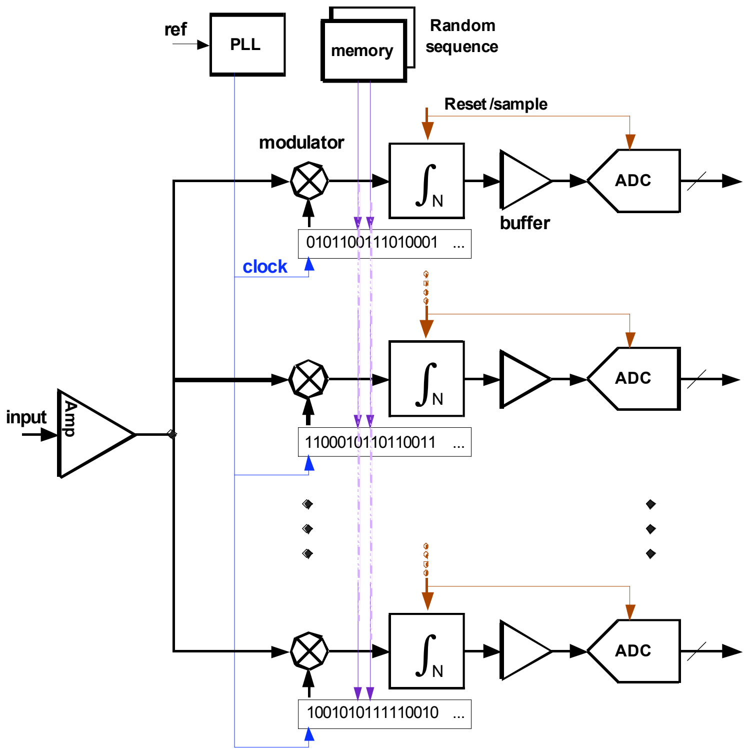

- Random demodulator being built at part of DARPA A2I program
  (Emami, Hoyos, Massoud)
- Multiple (8) channels, operating with different mixing sequences
- Effective BW/chan = 2.5 GHz
  Sample rate/chan = 50 MHz
- Applications: radar pulse detection, communications surveillance, geolocation

Sampling correlated signals

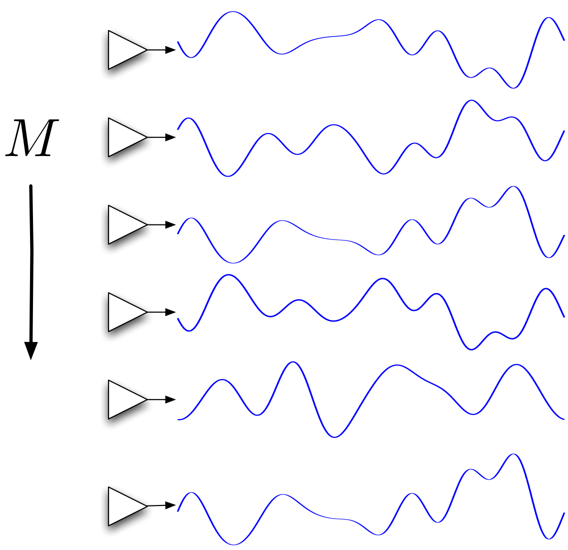

- Goal: acquire an *ensemble* of M signals
- Bandlimited to $W/2$
- "Correlated" $\rightarrow$ M signals are $\approx$ linear combinations of R signals

Sampling correlated signals

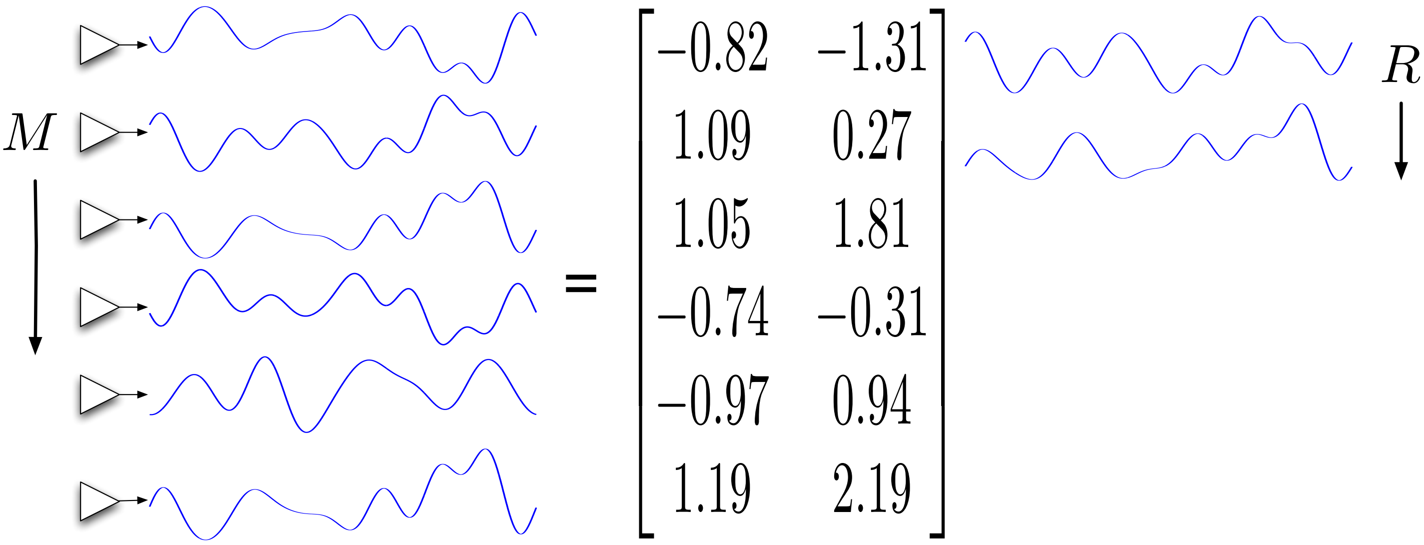

- Goal: acquire an *ensemble* of M signals
- Bandlimited to $W/2$
- "Correlated" $\rightarrow$ M signals are $\approx$ linear combinations of R signals

Sensor arrays

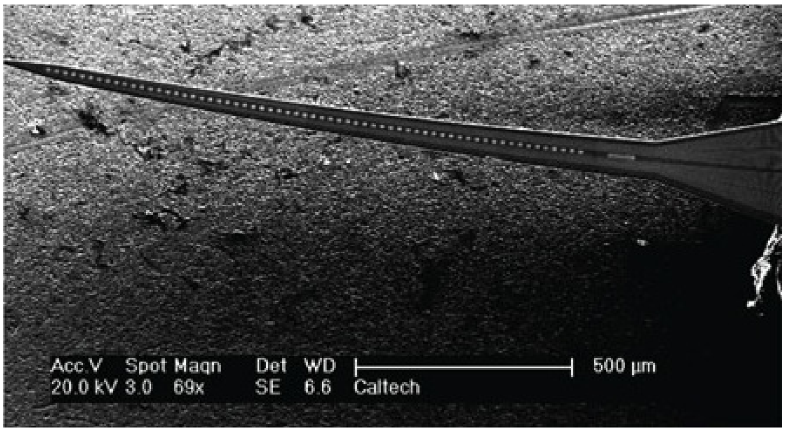

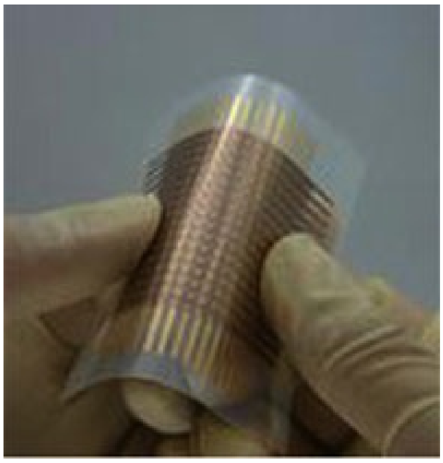

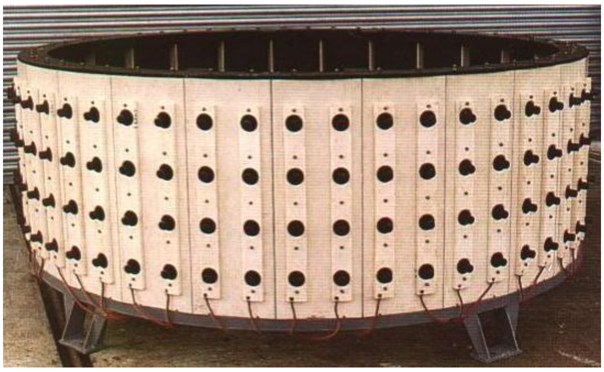

Low-rank matrix recovery

- Given P *linear samples* of a matrix,

$$y = \mathcal{A}(\mathbf{X}_0), \quad y \in \mathbb{R}^P, \quad \mathbf{X}_0 \in \mathbb{R}^{M \times W}$$

we solve

$$\min_{\mathbf{X}} \ \|\mathbf{X}\|_* \quad \text{subject to} \quad \mathcal{A}(\mathbf{X}) = y$$

where $\|\mathbf{X}\|_*$ is the nuclear norm: the sum of the singular values of $\mathbf{X}$.

Low-rank matrix recovery

- Given P *linear samples* of a matrix,

$$y = \mathcal{A}(\mathbf{X}_0), \quad y \in \mathbb{R}^P, \quad \mathbf{X}_0 \in \mathbb{R}^{M\times W}$$

  we solve

$$\min_{\mathbf{X}} \ \|\mathbf{X}\|_* \quad \text{subject to} \quad \mathcal{A}(\mathbf{X}) = y$$

  where $\|\mathbf{X}\|_*$ is the nuclear norm: the sum of the singular values of $\mathbf{X}$.

- If $\mathbf{X}_0$ is rank-R and $\mathcal{A}$ obeys the mRIP:

$$(1-\delta)\|\mathbf{X}\|_F^2 \ \le \ \|\mathcal{A}(\mathbf{X})\|_2^2 \ \le \ (1+\delta)\|\mathbf{X}\|_F^2 \quad \forall \text{ rank-}2R\ \mathbf{X},$$

  then we can stably recover $\mathbf{X}_0$ from y. (Recht et. al '07)

Low-rank matrix recovery

- Given P *linear samples* of a matrix,

$$y = \mathcal{A}(\mathbf{X}_0), \quad y \in \mathbb{R}^P, \quad \mathbf{X}_0 \in \mathbb{R}^{M \times W}$$

  we solve

$$\min_{\mathbf{X}} \ \|\mathbf{X}\|_* \quad \text{subject to} \quad \mathcal{A}(\mathbf{X}) = y$$

  where $\|\mathbf{X}\|_*$ is the nuclear norm: the sum of the singular values of $\mathbf{X}$.

- If $\mathbf{X}_0$ is rank-R and $\mathcal{A}$ obeys the mRIP:

$$(1-\delta)\|\mathbf{X}\|_F^2 \ \leq \ \|\mathcal{A}(\mathbf{X})\|_2^2 \ \leq \ (1+\delta)\|\mathbf{X}\|_F^2 \quad \forall \text{ rank-}2R \ \mathbf{X},$$

  then we can stably recover $\mathbf{X}_0$ from y. (Recht et. al '07)

- An 'generic' (iid random) sampler $\mathcal{A}$ (stably) recovers $\mathbf{X}_0$ from y when

$$\begin{aligned} \#\text{samples} &\gtrsim \ R \cdot \max(M, W) \\ &\gtrsim RW \quad \text{(in our case)} \end{aligned}$$

CS for correlated signals: modulated multiplexing

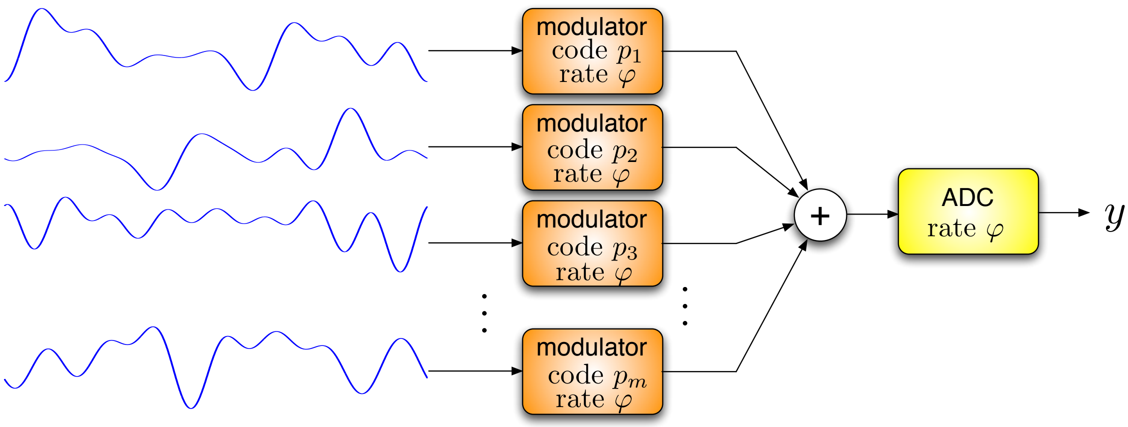

- If the signals are spread out uniformly in time, then the ADC and modulators can run at rate

$$\varphi \gtrsim RW \log^{3/2}(MW)$$

- Requires signals to be (mildly) spread out in time

Summary

- Main message of CS:

  We can recover an S-sparse signal in $\mathbb{R}^N$ from
  $\sim S \cdot \log N$ measurements

  We can recover a rank-R matrix in $\mathbb{R}^{M \times W}$ from
  $\sim R \cdot \max(M, W)$ measurements

- Random matrices (iid entries)
  - easy to analyze, optimal bounds
  - universal
  - hard to implement and compute with
- Structured random matrices (random sampling, random convolution)
  - structured, and so computationally efficient
  - physical
  - much harder to analyze, bound with extra log-factors

Compressive sensing tells us ...

Sensing...

- ... we can sample *smarter* not faster
- ... we can replace front-end acquisition complexity with back-end computing
- ... injecting randomness allows us to *super-resolve* high-frequency signals (or high-resolution images) from low-frequency (low-resolution) measurements
- ... the acquisition process can be *independent* of the types of signals we are interested in

Compressive sensing tells us ...

Sensing...

- ... we can sample *smarter* not faster
- ... we can replace front-end acquisition complexity with back-end computing
- ... injecting randomness allows us to *super-resolve* high-frequency signals (or high-resolution images) from low-frequency (low-resolution) measurements
- ... the acquisition process can be *independent* of the types of signals we are interested in

Mathematics...

- ... there are unique *sparse* solutions to underdetermined systems of equations
- ... random projections keep sparse signals separated
- ... a seemlingly impossible optimization program (subset selection) can be solved using a tractable amount of computation