CMOS COMPRESSED IMAGING BY RANDOM CONVOLUTION

L. Jacques[1,3], P. Vandergheynst[1], A. Bibet[2], V. Majidzadeh[2], A. Schmid[2], Y. Leblebici[2]

[1]Signal Processing Laboratory. [2]Microelectronic Systems Laboratory (LSM).
Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland.
[3] Communications and Remote Sensing Laboratory (TELE), Université catholique de Louvain (UCL), Belgium.

ABSTRACT

We present a CMOS imager with built-in capability to perform Compressed Sensing. The adopted sensing strategy is the Random Convolution [6]. It is achieved by a shift register set in a pseudo-random configuration. It acts as a convolutive filter on the imager focal plane, the current issued from each CMOS pixel undergoing a pseudo-random redirection controlled by each component of the filter sequence. A pseudo-random triggering of the ADC reading is finally applied to complete the acquisition model. The feasibility of the imager and its robustness under noise and non-linearities have been confirmed by computer simulations, as well as the reconstruction tools supporting the Compressed Sensing theory.

Index Terms— Compressed Sensing, Imager, Analog Processing, Random Convolution, CMOS.

1. INTRODUCTION

The 20^{th} century has seen the development of a large variety of sensors capturing accurate representations of the physical world (e.g. optical sensors, radio receivers, seismic detector, ...). Since the purpose of these systems was to directly acquire a meaningful signal, a very fine sampling of this latter had to be performed. This was the context surrounding the Shannon-Nyquist condition stating that each continuous band-limited signal can be recovered from its discretization if its sampling rate is at least two times the bandwidth.

A recent theory named Compressed Sensing (or Compressive Sampling) [7] has shown that this lower bound on the sampling rate can be highly reduced, under the conditions that, first, the sampling is generalized to any linear measurement of the signal and second, specific a priori hypotheses on the signal are realized. In short, if the signal has only K non-zero (or important) coefficients in a given basis Ψ, then its generalized sampling can be achieved in only $M = O(K \log(M/K))$ *linear measurements*.

This straightforward statement is a real revolution for the physical design of many sensors. It means that a given signal does not need to be acquired in its initial space as previously, but can really be observed through a "distorting glass" (providing it is linear) with fewer measurements. The pair encoder (sensing) and decoder (reconstruction) is also asymmetric: the encoder is computationally light and linear, and also completely independent of the acquired signal (non-adaptive), while the decoder is non-linear and requires high computing power for real-time applications.

Interestingly, Compressed Sensing (described shortly in Section 2) reintroduces the concept of local analog processing of sensor signals, i.e. *in-situ*. Previous sampling schemes quickly lead to the digitalization of the recorded values, limiting as much as possible the analog path linking the real world to the digital output. However, the generalized sampling induced by CS increases the class of physical systems (intrinsically analog) leading to usable signal measurements. Last years have seen the development of such CS sensors: we may cite the one-pixel camera [1], CS Imager of Georgia Tech [2], Coded-Aperture Imaging [3], Ultra-wideband Frequency Hopping signals [4], and DNA microarrays [5].

In Section 3, we present a CMOS optical sensor array exploiting the key concepts of CS. This unit relies on a specific signal measurement named *Random Convolution* [6] implemented in the analog domain by the control of a shift register (1-bit memories sequence) acting as a convolutive filter in the focal plane. The imager array has been fully designed but has not yet been manufactured. Electrical simulations (analog and digital) have confirmed correct operation of the image sensor, while software simulations (Section 4) have been used to confirm the operation of the full system (encoder and decoder), and its stability under noise and non-linearities.

2. COMPRESSED SENSING: KEY CONCEPTS

Let us assume that an image $x \in \mathbb{R}^{N \times N}$ is "well described" in a certain orthonormal basis, e.g. the DCT or the Wavelet representation[1]. More precisely, by vectorizing x into an element of $\mathbb{R}^{\bar{N}}$ with $\bar{N} = N^2$, keeping the notation $x_{(p,q)}$ for the 2-D representation of the image, we assume that the decomposition $x = \Psi\alpha$ in a basis $\Psi \in \mathbb{R}^{\bar{N} \times \bar{N}}$, i.e. a set of $\bar{N}$ orthogonal elements $\psi_j \in \mathbb{R}^{\bar{N}}$ hosted in the columns of

LJ research is supported by the Belgian National Funds for Scientific Research (FRS-FNRS).

[1]This also holds for redundant basis such as the steerable wavelets or the curvelets.

Ψ, leads to a vector $\alpha \in \mathbb{R}^{\bar{N}}$ with few non-zero coefficients (strict sparsity) or with a power law decay in ordered amplitudes (compressible signal).

Classical sampling/compression strategies consist in observing all the x_i, i.e. at Nyquist rate, computing all the components of $\alpha = \Psi^T x$, and only keeping the K first coefficients of α to keep the essential information of x (or the exact one if x is K-sparse). This is a wasteful process however since in the best case $O(\bar{N})$ operations (e.g. for wavelet transform) are needed to compute α while only $K \ll N$ coefficients are kept.

Compressed Sensing (CS) theory [7] introduces a *sensing* or *measurement matrix* Φ made of $M \leqslant \bar{N}$ sensing vectors $\varphi_i \in \mathbb{R}^{\bar{N}}$ hosted in the rows of $\Phi \in \mathbb{R}^{M \times \bar{N}}$. CS shows that if $M \geqslant O(K \log N/K)$, then, for a matrix $\Phi = (\varphi_{ij}) \in \mathbb{R}^{M \times \bar{N}}$ generated randomly from a Gaussian distribution, i.e. $\varphi_{ij} \sim N(0, 1/M)$, x can be recovered from the *measurement vector* $y = \Phi x$ with an overwhelming probability. This is achieved by solving the Basis Pursuit (BP) problem

$$\min_u \|u\|_1 \quad \text{subject to} \quad y = \Phi\Psi u, \qquad \text{(BP)}$$

where $\|u\|_1 = \sum_i |u_i|$ (the ℓ_1 norm). The proof of this recovery relies on the Restricted Isometry Property (RIP) of such a random matrix Φ, i.e. the fact that there exists a constant $0 < \delta_k < 1$ such that $(1-\delta_k)\|x\|_2^2 \leqslant \|\Phi x\|_2^2 \leqslant (1+\delta_k)\|x\|_2^2$, for all K-sparse $x \in \mathbb{R}^N$.

Other random matrices such as the Bernoulli/Rademacher matrix, i.e. $\varphi_{ij} = \pm 1/\sqrt{M}$ with equal probability, are RIP. As described in the next Section, we use another particular sensing matrix almost as optimal as the Gaussian or the Rademacher matrices: the Random Convolution [6].

Outside of the ideal sensing, i.e. when the measurements are corrupted by some additive Gaussian noise n in the model $y = \Phi x + n$, a more stable reconstruction is provided by the *Basis Pursuit DeNoise* (BPDN) method, i.e.

$$\min_u \|u\|_1 \quad \text{subject to} \quad \|y - \Phi\Psi u\|_2 \leqslant \epsilon, \qquad \text{(BPDN)}$$

with ϵ set in function of the noise power and $\|w\|_2^2 = \sum_i |w_i|^2$. Both BP and BPDN can be solved efficiently using for instance Linear Programming techniques (LP) or Second Order Cone programming (SOC) respectively.

It is often more efficient to impose that its *gradient* be sparse. Indeed, in that case the ℓ_1 norm in BP and BPDN can be replaced by the Total Variation (semi) norm $\|u\|_{\text{TV}} = \|\nabla u\|_1 = \sum_i \|(\nabla u)_i\|$, with the discrete *gradient operator* $(\nabla x) = (\nabla^1 x, \nabla^2 x) \in \mathbb{R}^{\bar{N}} \times \mathbb{R}^{\bar{N}}$, with $(\nabla^1 x)_{(p,q)} = x_{(p,q+1)} - x_{(p,q)}$ if $p < N$ and $(\nabla^1 x)_{(p,q)} = 0$ if $p = N$, and $\nabla^2 x = \Pi \nabla^1 \Pi\, x$ with $(\Pi\, x)_{(p,q)} = x_{(q,p)}$.

Experiments in fields such as Magnetic Resonance Imaging [8] show that this optimization scheme is however very efficient. We use this method for the reconstruction of images acquired by our CS imager. A simulation presented in Section 4 confirms the potential of the approach.

3. DESCRIPTION OF THE IMAGER

3.1. Framework and Sensing Strategy

In the sensing model $y = \Phi x$ of our imager, we could have taken for $\Phi \in \mathbb{R}^{\bar{N} \times \bar{N}}$ the Gaussian random matrix[2]. However, we have preferred the Random Convolution strategy explained in a recent work of J. Romberg [6]. In short, it dictates to pick M random values in the convolution of the image $x \in \mathbb{R}^{\bar{N}}$ with a random filter. The resulting sensing matrix is still optimal and requires a similar number of measurements[3], i.e. $M \geqslant O(K \log(\bar{N}/\delta))$ for a probability of successful recovery of $1 - \delta$, given $\delta \in [0, 1]$.

The random convolution of an image $x \in \mathbb{R}^{\bar{N}}$ by a random filter $a \in \mathbb{R}^{\bar{N}}$ is mathematically described by

$$y_i = (\Phi x)_i = \sum_i a_{r(i)-j}\, x_j \;=\; (x * a)_{r(i)}, \qquad (1)$$

where $r(i) \in \{1, \cdots, \bar{N}\}$ is selected uniformly at random. Initially, the filter a is defined in the Fourier domain by a vector of unit amplitudes and random phases. However, in this work, we use the less optimal choice of a filter a defined spatially as a Rademacher sequence of ± 1.

This sensing is interesting for two aspects. First, in the reconstruction stage that generally involves many matrix-vector computations with Φ and Φ^T, these operations are obviously simplified into the application of some FFTs of $O(\bar{N} \log_2 \bar{N})$ complexity.

Second, random convolution can be implemented very simply by the action of a shift-register (a chain of one-bit memories linked to each pixel) on the currents provided by the sensor array. We describe this striking aspect in the next section.

3.2. Microelectronic architecture

The system architecture of the imager array is depicted in (Fig. 1). A regular array of $N \times N$ standard CMOS Passive Pixel Sensors (PPS) with an active area of $30\mu\text{m} \times 30\mu\text{m}$ forms the core of the imager. Each pixel contains a photodiode delivering a maximal current of $200\mu\text{A}$. The PPS configuration has some drawbacks related to high consumption, average sensitivity, but enables a high design fill-factor.

A one-bit flip-flop memory is implemented in each PPS, in the close vicinity of the photodiode. This memory stores the information related to the random coefficient filter value a_i. Its input and output are connected to the memories of two neighboring pixels (according to the arrows in Fig. 1), thereby forming a $\bar{N}$-bits Shift Register (SR). If we push one 1-bit value into the input of the first pixel memory, the whole

[2] Or a pseudo-random alternative starting from a given *seed* to avoid the storage of this huge matrix.

[3] There is an additional constraint however imposing $M \geqslant O(\log^3(\bar{N}/\delta))$ independently of the sparsity level K.

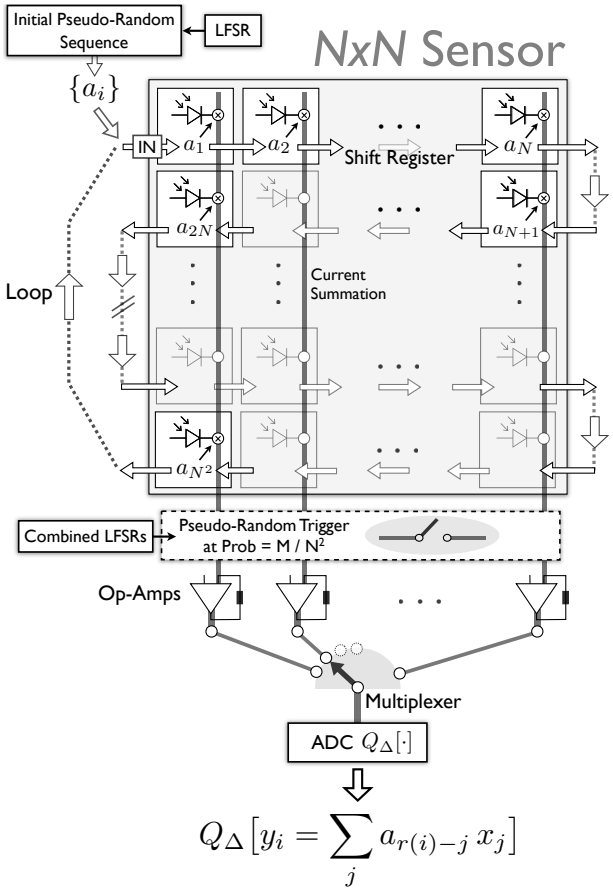

Fig. 1. Scheme of the CMOS Compressed Imager.

sequence a is moved by one element in the SR, which is the exact behavior required to implement a convolution.

The image acquisition process is achieved according to the following steps. First, as an initialization stage, a pseudo-random Rademacher sequence a is generated by a Linear Feedback Shift Register (LFSR) with a cyclic period larger[4] than $\bar{N}$. This pseudo-random sequence can be regenerated on the decoder knowing the *seed* of the LFSR. As it is generated, this sequence is pushed into the $\bar{N}$-bits SR using the memory input of the first pixel. The system is ready to perform the first image acquisition after $\bar{N}$ clock cycles. In each PPS, the output current is proportional to the light intensity x_i. The sign of this current is adjusted by the 1-bit value stored in the memory. This current is collected according to Kirchoff's law (added or subtracted) on a wire connecting all the pixels of the same grid column. Each column is connected to the input of one Operational Amplifier (Op-Amp). The voltage output of the Op-Amp is subsequently converted to a digital value by an Analog to Digital converter (ADC), using time-domain multiplexing of the input. The ADC output is in turn accumulated to form the final compressed image value. Scanning the columns is intended to limit the current provided to the Op-Amp, but requires a higher processing frequency, and N digital summations (accumulations) of the ADC outputs to obtain one CS measurement. This scenario was adopted to limit the physical width of the column lines, as a benefit of a smaller dynamic range of the current to be handled.

Second, for the realisation of the next measurements, the content of the SR has to be adapted. According to the aforementioned developments, by pushing the last 1-bit value of the grid, i.e. a_{N^2}, into the first pixel memory[5], the system is potentially ready to acquire a second measurement. However, to fit the random convolution model (1), a random triggering of the measurement reading, i.e. of the Op-Amps/ADC blocks activation, must be applied. This triggering is obtained by logically combining several LFSRs[6] so that it is activated with a certain rational probability p. If the triggering is off, a new SR shift is performed without any reading. If it is on, a measurement is acquired and quantized by the Op-Amps/ADC layer according to the scheme described for the first acquisition.

After $\bar{N}$ shifts of the SR, which correspond to its cycling period, $M \simeq p\bar{N}$ triggerings/measurements are produced, i.e. an average of $\bar{N}/M = p^{-1}$ clock cycles per measurement. In our project, $M = \bar{N}/3$ measurements may be provided in 400ms, taking into account the bandwidth of the custom Op-Amps (214kHz) and an initial setup of $N = 64$. In a near future, we plan to improve these technological characteristics to reach 25 frames per second (fps), i.e. 40ms per frame, with $\bar{N} = 256^2$ pixels, and use an Active Pixel Sensor (APS) configuration.

Irrelevantly from the final number of frames per second, it is important to understand that our scheme assumes that the observed scene be still over the time elapsed between two consecutive frames, i.e. between two full acquisitions of M measurements.

4. SIMULATION

The output of the imager is simulated as the measurement vector $\tilde{y}$ obtained from the quantization of a noisy random convolution of an 256×256 image (Fig. 2) with a Rademacher pseudo-random pattern. In other words, $\tilde{y} = Q_\Delta\big[\Phi x + n_0\big] \simeq \Phi x + n$, where n_0 is a white noise on the measurements (e.g. thermal noise), i.e. $(n_0)_i \sim N(0, \sigma_0^2)$, and Q_Δ is the quantization operator of step size Δ. This value is set so that $\tilde{y}$ can be coded in 11-bits, i.e. $\Delta = 2\,\|\Phi x\|_\infty / 2^{11}$. Thus, the final noise n combines the quantization noise, the measurement noise n_0, and possible non-linearities in the system, e.g. due to the Op-Amp current-to-voltage conversion or to the ADC. We assume here $n_i \sim N(0, \sigma^2)$ with $\sigma^2 = \sigma_0^2 + \sigma_{ADC}^2 + \frac{\Delta^2}{12}$ and

[4]The number of registers in the LFSR has to be larger than $\log_2 \bar{N}$.

[5]Equivalently, if the LFSR period is equal to $\bar{N}$, the desired loop occurs naturally in the pseudo-random sequence without any physical connection between the first and the last pixel memories.

[6]For instance, with a global AND operation on n LFSR outputs, the triggering occurs with probability $p = 1/2^n$ at each clock signal

Fig. 2. (left) Original Image. (right) Reconstructed image with $M = \lfloor \bar{N}/3 \rfloor$, from noisy measurements quantized on 11-bits. PSNR 27.3 dB.

$\sqrt{\sigma_0^2 + \sigma_{ADC}^2} = \|\Phi x\|_\infty / 100$.

From this imager simulation, we have run the reconstruction stage (i.e. the decoder) using a regularized BPDN solver named TwIST [9] defined with the TV norm (see Section 2). The regularizing parameter has been tuned iteratively so that the fidelity term $\|\Phi u - \tilde{y}\|_2 \simeq \epsilon$, where $\epsilon \simeq \sigma\sqrt{M}$ is the noise power. Notice that in case where ϵ cannot be easily estimated a *Cross-Validation* technique could be used to avoid noise overfitting in the reconstructed image [10].

The result of the reconstruction is presented in Fig. 2 for a number of measurement $M = \lfloor \bar{N}/3 \rfloor$. The reconstruction reaches a PSNR of 27.3 dB.

5. PREVIOUS WORKS

Our system exhibits similarities with the CMOS Analog Imager (CAI) of R. Robucci et al. [2]. Nevertheless, our system is optimized for Compressed Imaging while the CAI is a more general architecture aims at realizing alternate analog signal processing (e.g. DCT or wavelet transform). This generality reflects into a larger electronic system, e.g. due to the storage of the (random or structured) sensing matrix Φ out off the array. Our system is more performant in terms of memory, where random convolution needs only a $\bar{N}$ bit storage on the focal plane, and faster since the SR configuration can be adapted for the next measurement within few clock cycles (i.e. $\bar{N}/M$).

Another analog implementation is the Single-Pixel Camera [1] of the Rice group. In this system, the analog sensing is obtained optically by focusing the reflection of an image on a Digital Micromirror Device (DMD) in a random sensing configuration on an unique photodiode (pixel). As for any micro-mechanical system used to perform analog processing, the pair DMD-photodiode is subject to various nonlinearities (e.g. nonuniform reflectance of the mirrors through the focusing lens, nonuniform mirror positions, light to current photodiode conversion). We are convinced that our imager suffers less from these imperfections since it relies on an homogeneous analog processing in the electric domain, and uses a mature CMOS fabrication technology. Errors and nonlinearities induced by all the micro-electronic modules (e.g. PPS, Op-Amp, ADC) can be reduced, modeled and on-chip calibration applied to counter their effects.

6. CONCLUSION

We have presented a new microelectronic system for compressed imaging operating mainly in the analog domain, where ADC quantization is applied on the final measurements. A moderate size grid of 64^2 pixels has been selected, and a first prototype has been developed in a 0.35μm CMOS technology. Fabrication, testing and calibration of the device is expected to provide insights into noise perfomance and actual non-linearities. The simulation model will include them, enabling extracting the correct noise power ϵ.

In a near future, we plan to adapt the same technology to 2-D grid of biosensors for analysing the electrical activity of a group of connected neural cells [11]. The biosignal produced is indeed sparse both in the spatial and in time domain, confirming the applicability of CS.

7. REFERENCES

[1] M.F. Duarte, M.A. Davenport, D. Takbar, J.N. Laska, T. Sun, K.F. Kelly, and R.G. Baraniuk, "Single-Pixel Imaging via Compressive Sampling," *Sig. Proc. Mag., IEEE*, vol. 25, no. 2, pp. 83–91, 2008.

[2] R. Robucci, L.K. Chiu, J. Gray, J. Romberg, P. Hasler, and D. Anderson, "Compressive sensing on a CMOS separable transform image sensor," *Int. Conf. on Acoustics, Speech and Sig. Proc., ICASSP 2008*, pp. 5125–5128, 2008.

[3] R.F. Marcia and R.M. Willett, "Compressive coded aperture superresolution image reconstruction," *Int. Conf. on Acoustics, Speech and Sig. Proc., ICASSP 2008*, pp. 833–836, 2008.

[4] J. Laska, S. Kirolos, Y. Massoud, R. Baraniuk, A. Gilbert, M. Iwen, and M. Strauss, "Random sampling for analog-to-information conversion of wideband signals," *Proc. IEEE Dallas Circuits and Systems Workshop (DCAS)*, 2006.

[5] F. Parvaresh, H. Vikalo, S. Misra, and B. Hassibi, "Recovering Sparse Signals Using Sparse Measurement Matrices in Compressed DNA Microarrays," *IEEE Journal of Sel. Topics in Sig. Proc.*, vol. 2, no. 3, pp. 275–285, 2008.

[6] Justin Romberg, "Sensing by Random Convolution," *IEEE Int. Work. on Comp. Adv. Multi-Sensor Adaptive Proc., CAMPSAP 2007*, pp. 137–140, 2007.

[7] E.J. Candes and J. Romberg, "Quantitative Robust Uncertainty Principles and Optimally Sparse Decompositions," *Found. Comp. Math.*, vol. 6, no. 2, pp. 227–254, 2006.

[8] M. Lustig, D. Donoho, and J.M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182, 2007.

[9] JM Bioucas-Dias and MAT Figueiredo, "A New TwIST: Two-Step Iterative Shrinkage/Thresholding Algorithms for Image Restoration," *IEEE Trans. Image Proc.*, vol. 16, no. 12, pp. 2992–3004, 2007.

[10] P. Boufounos, M.F. Duarte, and R.G. Baraniuk, "Sparse Signal Reconstruction from Noisy Compressive Measurements using Cross Validation," *IEEE/SP 14th Work. Stat. Sig. Proc., 2007. SSP'07.* , pp. 299–303, 2007.

[11] M Jenkner, M Tartagni, A Hierlemann, and R Thewes, "Cell-based cmos sensor and actuator arrays," *IEEE Journ. Solid-State Circuits*, vol. 39, no. 12, pp. 2431 – 2437, Dec 2004.