Compressive sampling of pulse trains: spread the spectrum !

Farid Movahedi Naini, Rémi Gribonval, Laurent Jacques, Pierre Vandergheynst

HAL Id: hal-00481954
https://hal.science/hal-00481954

Submitted on 7 May 2010

COMPRESSIVE SAMPLING OF PULSE TRAINS: SPREAD THE SPECTRUM !

Farid M. Naini[1], *Rémi Gribonval*[2], *Laurent Jacques*[1,3] *and Pierre Vandergheynst*[1]

[1] Signal Processing Laboratory, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
[2] Centre de Recherche INRIA Rennes, IRISA Campus de Beaulieu, F-35042 Rennes cedex, France
[3] Communications and Remote Sensing Laboratory, Université catholique de Louvain (UCL), B-1348 Louvain-la-Neuve, Belgium.

ABSTRACT

In this paper we consider the problem of sampling far below the Nyquist rate signals that are sparse linear superpositions of shifts of a known, potentially wide-band, pulse. This signal model is key for applications such as Ultra Wide Band (UWB) communications or neural signal processing. Following the recently proposed Compressed Sensing methodology, we study several acquisition strategies and show that the approximations recovered via ℓ_1 minimization are greatly enhanced if one uses Spread Spectrum modulation prior to applying random Fourier measurements. We complement our experiments with a discussion of possible hardware implementation of our technique.

Index Terms— Compressive sampling, compressed sensing, pulse trains, Ultra Wide Band, shift-invariant dictionary

1. INTRODUCTION

For many years, signal acquisition systems have been based on the Nyquist-Shannon sampling theorem that states that the number of samples needed to recover a signal without error is twice the bandwidth. Recently, the emerging field of *compressive sensing* (CS) has given a fresh look at data acquisition: the number of required measurements needed to reconstruct a signal without error depends on its *sparsity* and not on its bandwidth [1, 2, 3]. Hence, if the signal has a very sparse representation on some basis, or more generally on some dictionary [4], it is possible to sample it using very few, linear measurements.

Although the compressive sensing theory is flourishing, it is not still a completely mature field and important theoretical and practical questions remain to be solved. When the signal is a sparse train of Diracs, it was proved in the seminal paper of Candès and Tao [2] that it can be sampled efficiently using a small random subset of its Fourier coefficients. More generally, if the signal has a sparse representation on an orthonormal basis, it is known that a fairly large class of random matrices can be used for obtaining compressive samples. The interested reader is referred to the introductory review paper [5]

for more information. In the general case where the sparsifying system is a dictionary (a non necessarily tight frame) the results are themselves quite sparse, though, and it is not clear what sensing matrices should be used. Likewise, most results in CS are expressed in terms of the number of measurements needed to approximate the coefficients of the signal's sparse representation. But as we shall see below, there are many situations where a recovery algorithm can fail approximating the coefficients and still yield very precise approximations of the signal itself.

In this paper we concentrate on signals that are sparse on a dictionary made of all shifts of a fixed, known, waveform or pulse. This model is introduced and discussed in Section 2. We then study in Section 3 several measurement strategies used for compressively sampling these signals. Our main quality metric is the signal-to-noise ratio and we emphasize why this constraint has a strong influence on practical results. In particular, we design and discuss a novel sensing strategy based on pre-modulating the signal with a spread-spectrum sequence followed by random sampling in the Fourier domain. Section 4 contains ample experimental evidence that this strategy significantly outperforms other classical compressive sampling scenarios such as pure random Fourier sampling or gaussian measurements. In Section 5 we discuss how our scheme can be adapted to meet the requirements of potential hardware implementation and show that these modifications do not significantly degrade our results. Our approach bears several similarities with recent papers on Analog-to-Information conversion. In particular, [8] introduced random sampling of wide band multi-hop signals and [9] used the idea of random demodulation which is close in spirit to our spread spectrum pre-modulation.

2. SIGNAL MODEL

Throughout this paper we will denote by x a real valued N-dimensional signal, i.e $x \in \mathbb{R}^N$. We wish to record few, say M with $M \ll N$, linear measurements of our signals and we thus model this sampling process by a measurement matrix Φ, with $\Phi \in \mathbb{R}^{M \times N}$. We collect these samples in the M-dimensional vector $y = \Phi x$. Many papers focus on signal models where x is a sparse vector or has a sparse rep-

resentation on an orthogonal basis. In this work however, we are motivated by signals which are linear superposition of K time-shifted pulses :

$$x[n] = \sum_{k=0}^{K-1} \alpha_k \psi[n - n_k],$$

where the mother pulse ψ has a limited support compared to N. This type of signal model is of particular interest in applications such as Ultra Wide Band communications [6] or bio-sensing [7]. In matrix form, our signals obey a sparse synthesis model, $x = \Psi\alpha$, where α is a sparse vector of coefficients, $\|\alpha\|_0 = K \ll N$, and the synthesis dictionary $\Psi \in \mathbb{R}^{N \times N}$. Moreover, the shift-invariant nature of our problem imposes that Ψ is a circulant matrix whose columns are shifted replicas of the mother pulse. It is very convenient to write $\Psi = F^*\Lambda F$, where F is the column normalized DFT matrix [1] and Λ is a diagonal matrix with $\text{diag}(\Lambda) = F\psi$. It has to be noted that our dictionary Ψ is not an orthogonal basis[2] and that we impose no specific restriction on the pulse ψ. With these notations, our compressive sampling problem can be formulated as follows. We would like to collect $M < N$ measurements of our signal $y = \Phi x = \Phi\Psi\alpha$ such that it is possible to reconstruct a good approximation of the original signal x. In this paper, the quality measure is the output Signal-to-Noise Ratio (SNR) :

$$\text{SNR}_{\text{out}}(x, \hat{x}) = -20 \log_{10} \|x - \hat{x}\|_2 / \|x\|_2. \quad (1)$$

This seemingly natural choice has a drastic influence on the way we present our results. Whereas most research papers in compressive sensing will discuss the recovery of the sparse vector α, exact recovery has in fact very little importance for us. Imagine indeed that for a given x we would recover all pulses *up to a little time shift of each pulse*. This would mean that we would not have recovered α at all. Yet, since our pulses are waveforms and not perfect spikes, the reconstructed signal would be very close to the original one ! In other words, we do not care about perfect identification of the support of the sparse vector α, instead what really matters for us is the SNR (1).

One of the difficulties with compressed sensing is that there are many possible algorithms to recover the signal from the measurements. One first restriction is that we need to select an algorithm that does signal *approximation* and that is resilient to noise added to the signal x or to the measurements y. We have tested several algorithms and decided to focus on the Basis Pursuit Denoising framework, i.e solving the following optimization problem :

$$\min \|\alpha\|_1 \text{ subject to } \|\Phi\Psi\alpha - y\|_2 \leq \sigma. \quad (\text{BPDN}_\sigma)$$

In all our experiments we have used the SPGL1[3] solver in

---

[1] The matrix F used throughout this paper has unit 2-norm columns, so the usual N-by-N DFT matrix is actually $\sqrt{N} \cdot F$.

[2] Even if $\Psi \in \mathbb{R}^{N \times N}$, the span of its columns is smaller than $\mathbb{R}^N$.

[3] http://www.cs.ubc.ca/labs/scl/spgl1/

combination with the SPARCO toolbox[4] for its handy syntax for designing sensing operators.

3. MEASUREMENT STRATEGIES

Given our particular model, the signals we are interested in are convolutions of a spike train with the mother pulse ψ. Since random Fourier measurements are close to optimal for compressive sampling of spike trains, one may be tempted to reduce our problem to this ideal case. Indeed, let us define the following measurement matrix :

$$\Phi_I = RFF^*\Lambda^{-1}F, \quad (2)$$

where R is the sub-sampling operator : a M-by-N matrix filled up with zeroes except one element on each row that is equal to N/M and such that there is at most one non-zero element on each column. With this normalization, the matrix RF selects M rows of F and has unit norm columns. The effect of this measurement matrix is to first invert the convolution with the pulse and then try to recover the spike train α. Applying (2) to one of our signals indeed results in the following sampling strategy :

$$y = \Phi_I\Psi\alpha = (RFF^*\Lambda^{-1}F)(F^*\Lambda F)\alpha = RF\alpha, \quad (\text{INV})$$

which is simply standard Fourier sub-sampling of α. There are strong limits to this line of thought, though. For nothing guarantees that the inverse convolution is well-conditioned : Λ can even have zeroes. This technique is thus, unsurprisingly, highly sensitive to noise as illustrated on Figure 1: the amount of measurements M needed to achieve a 30dB SNR reconstruction increases rapidly as noise is added to the signal. The amount of noise is controlled by the input SNR, i.e. $\text{SNR}_{\text{in}} = -20 \log_{10} \|n\|_2 / \|x\|_2$, which, in this experiment, varied between $\text{SNR}_{in} = +\infty$ (no noise) and 40 dB.

In a sense, the previous measurement strategy fails because it wastes measurements : imagine, schematically, that Λ is zero on some intervals in the Fourier domain. The Fourier samples selected by the matrix R will have a non-negligible probability to fall in one of these intervals and will therefore bring no information about the signal. Clearly, what we need is a way to ensure that *every measurement counts*, i.e we need to make sure that the signal to which we apply the Fourier sub-sampling matrix RF occupies the full spectrum. We thus propose as alternative strategy to pre-modulate the input signal with a spread spectrum sequence $p[n]$. There are many good such sequences, some (pseudo)-random, some deterministic and it is behind the scope of this paper to describe the full theory of spread spectrum techniques. It should be noted, though, that we tested several sequences and observed that the results do not vary significantly. Our proposed sensing strategy is thus the following :

$$y = \Phi_S\Psi\alpha = RFP\Psi\alpha, \quad (\text{SPREAD})$$

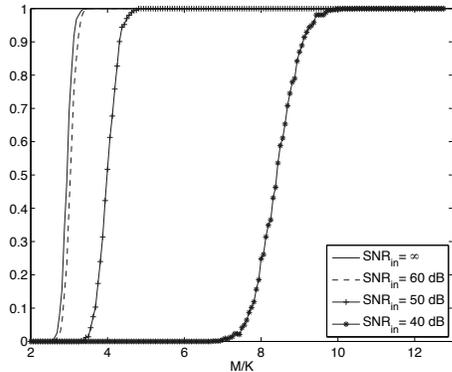
[4] http://www.cs.ubc.ca/labs/scl/sparco/

Fig. 1. Probability of reaching an output SNR greater than 30 dB as a function of the over-sampling ratio M/K. The results were obtained by drawing at random 1000 signals of length $N = 1024$ and sparsity $K = 80$. The number of measurements was varied from $2K$ to $N - 1$.

where RF is the Fourier sub-sampling matrix described above and P is a diagonal matrix with $\text{diag}(P) = p$.

4. EXPERIMENTAL RESULTS

In order to evaluate the performance of the SPREAD sampling strategy, it is important to compare it to other measurement matrices. In our experiments we systematically compared with the INV strategy described above and with two classical measurement matrices :

$$y = \Phi_F \Psi \alpha = RF \Psi \alpha, \qquad \text{(FOURIER)}$$

which is the traditional Fourier sub-sampling, and

$$y = \Phi_G \Psi \alpha = G \Psi \alpha \qquad \text{(GAUSS)}$$

where G is a M-by-N matrix of independent centered gaussian variables [5]. The mother pulse $\psi[n]$ defining Ψ has been set to a truncated first derivative of Gaussian with 17 non-zero elements.

Our first experiment consists in evaluating the stability of the sampling strategy with respect to the amount of noise added to the signal, i.e the measurements record $y = \Phi(x + n)$, where n is centered gaussian noise and Φ is set to one of the sensing strategies described above. The input SNR varied between -10 dB to 60 dB. The input signal had length $N = 1024$ and sparsity $K = 15$. We gathered $M = 7K \simeq 100$ measurements and measured the output SNR. The same experiment was repeated for all sampling strategies (INV, FOURIER, GAUSS and SPREAD) and the results were averaged over 1000 randomly simulated input signals[5]. The outcome, displayed in Figure 2, clearly shows the attractiveness of the SPREAD technique. First, the results

---

[5] K locations were selected randomly in α and filled with random Gaussian values.

are always better than the other three measurement matrices. At low input SNR, the INV methodology is clearly the worst one as explained before, while GAUSS and FOURIER behave similarly. As the input SNR get higher, INV catches up with FOURIER and GAUSS: as we already know, FOURIER wastes measurements and once the effect of noise disappears the reconstruction quality stalls evidently.
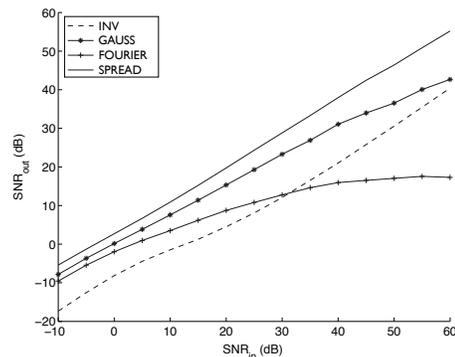

Fig. 2. Output SNR as a function of input SNR averaged over 1000 signals of length $N = 1024$ drawn at random. The sparsity was fixed at $K = 15$ and the number of measurements at $M = 7K$.

Our second experiment is aimed at testing how many measurements we need to achieve a given output SNR. More precisely, we randomly simulated signals of length $N = 1024$ and varied the sparsity K between 10 and 100. The input SNR was fixed at 30 dB. For every K, and for all sampling strategies, we then recorded how many measurements were needed, in average over all signals, to achieve an output SNR greater than 25 dB. Figure 3 shows the results in terms of the over-sampling factor M/N versus the relative sparsity K/N. One notices that over-sampling diminishes rather quickly and, for a quite realistic 10% sparsity, it is roughly of the order of 4K.

5. HARDWARE PERSPECTIVE

The previous set of experiments clearly shows the interest of the SPREAD strategy for compressive sampling of pulse trains signals. Our experiments and findings were carried out in a purely digital setting, though, while clearly this technique is aimed at sampling *analog* signals. In other words, for a mother wideband pulse $\psi(t)$, the ideal analog model is in fact $x(t) = \sum_{k=0}^{K-1} \alpha_k \psi(t - t_k)$, on a given time interval of length Δt. In this case, one would like to keep the whole signal processing chain as long as possible within the analog domain, and only manipulate sub-Nyquist digital samples. Looking towards this perspective, let us streamline how the SPREAD strategy can be modified to manipulate analog signals in an idealized analog architecture (Fig. 4).

First, the spread spectrum sequence will have to be an analog signal $p(t)$. The pre-modulation $x(t) \times p(t)$ can be
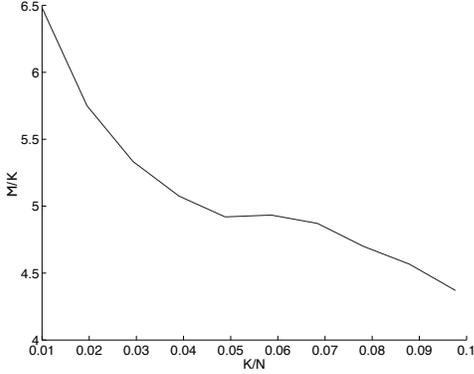
Fig. 3. Over-sampling ratio M/K as a function of the relative sparsity M/N for the SPREAD strategy. In this experiment, M is the average number of measurements needed to reach an output SNR greater than 25 dB. In all runs, the input SNR was fixed at 30 dB.
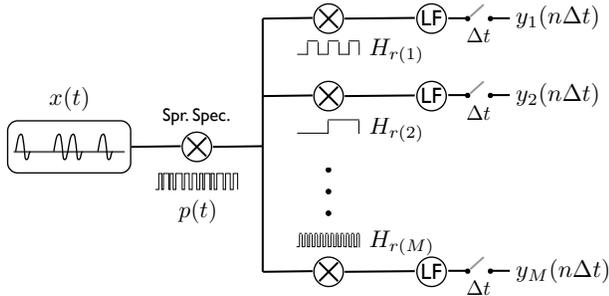

Fig. 4. Possible Analog Architecture.

achieved by flipping (at Nyquist rate) the polarity of the input signal using (pseudo)-random or deterministic strategies which can be implemented very fast in analog electronics [9]. Second, the Fourier sub-sampling operator, which would require a large amount of very finely tuned analog band-pass filters, is replaced by the more analog friendly Hadamard transform: the pre-modulated signal would be multiplexed M times analogically and each replica multiplied by a randomly selected Hadamard sequence[6] $H_{r(i)}(t)$, where $r(i)$ models the ith random selection. The output of each branch is then accumulated analogically by low-pass filtering (LF) to form, by time sampling every Δt seconds, the corresponding scalar product $y_i(n\Delta t) = \int_{n\Delta t}^{(n+1)\Delta t} H_{r(i)}(t)\, p(t)\, x(t)\, \mathrm{d}t$. In this way, M measurements are computed in parallel over a time frame Δt of the original signal. One of the issues to understand with this modified scheme is just how sensitive the SPREAD strategy would be if one trades the Fourier transform for the Hadamard transform. We ran simulations, using the settings that generated Figure 2 but replaced the DFT matrix F by the corresponding Hadamard matrix. Let us call this new strategy H-SPREAD. The outcome is shown on Figure 5

---

[6]Using the same polarity flipping trick for the analog multiplication, the Hadamard sequences being composed also of ± 1.

where we compare SPREAD, H-SPREAD and GAUSS. Although a small degradation is observed, H-SPREAD behaves almost exactly like the GAUSS strategy, although the latter would not be easily implemented in hardware. The modified sampling technique thus shows really promising potential for a compressive sampling architecture that would operate as much as possible in the analog domain.
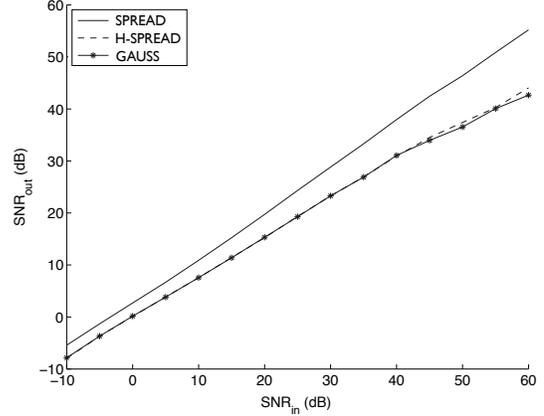

Fig. 5. Comparison of the SPREAD, H-SPREAD and GAUSS sampling strategies. The settings are the same as those used in Figure 2.

6. REFERENCES

[1] E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, pp. 969–986, 2007.

[2] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, pp. 489–509, 2006.

[3] D. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.

[4] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *IEEE Trans. Inform. Theory*, vol. 54, no. 5, pp. 2210–2219, 2008.

[5] E. Candes and M. Wakin, "An introduction to compressive sampling [a sensing/sampling paradigm that goes against the common knowledge in data acquisition]," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 21 – 30, Mar 2008.

[6] J. Paredes, G. Arce, and Z. Wang;, "Ultra-wideband compressed sensing: Channel estimation," *IEEE Journ. Sel. Top. Sig. Proc.*, vol. 1, no. 3, pp. 383 – 395, Oct 2007.

[7] 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.

[8] 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," *IEEE Dallas Circuits and Systems Workshop (DCAS)*, Dallas, Texas, 2006.

[9] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud " Theory and implementation of an analog-to-information converter using random demodulation," *IEEE Int. Symp. on Circuits and Systems (ISCAS)*, New Orleans, Louisiana, 2007.