



POLITECNICO DI TORINO
Repository ISTITUZIONALE

Compressed Sensing of Streams

Original

Compressed Sensing of Streams / Callegari, S.; Mangia, M.; Rovatti, R.; Setti, G.. - STAMPA. - (2019), pp. 835-838.
((Intervento presentato al convegno IEEE International Conference on Electronics, Circuits and Systems, ICECS 2019 tenutosi a ita nel 2019 [10.1109/ICECS46596.2019.8964730].

Availability:

This version is available at: 11583/2834906 since: 2020-06-12T10:00:39Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/ICECS46596.2019.8964730

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Compressed Sensing of $\Delta\Sigma$ Streams

Sergio Callegari*, Mauro Mangia[†] and Riccardo Rovatti[‡]

ARCES and DEI, University of Bologna, Italy

Email: *sergio.callegari@unibo.it,

[†]mauro.mangia2@unibo.it, [‡]riccardo.rovatti@unibo.it

Gianluca Setti

Polytechnic University of Turin, Italy

and ARCES, University of Bologna

Email: gianluca.setti@polito.it

Abstract—Compressed sensing (CS) is often applied at the *digital* level. We consider the case where CS follows a $\Delta\Sigma$ data converter and we show that CS can be practiced directly on the $\Delta\Sigma$ stream. In the proposed scheme, an appropriate sensing matrix incorporates the ability to get rid of the quantization noise from the $\Delta\Sigma$ modulator. We also show that a suitable sparsity basis enables the CS information recovery to be practiced directly at the Nyquist rate and that decimation, which is typically inherent in $\Delta\Sigma$ data acquisition, is not needed. Furthermore, the low depth of $\Delta\Sigma$ streams allows CS measures to be taken without multipliers, streamlining arithmetic blocks. A test case based on electrocardiograms is used to validate the approach.

Keywords—Compressed sensing, Delta-Sigma modulator, ECG

I. INTRODUCTION

Compressed sensing (CS) [1] consists in a set of mathematical methods enabling the recovery a signal from a set of measurements whose time density is much lower than the Nyquist rate. Such *undersampling* is permitted by an *a priori* knowledge of the signal structure, formalized through a *sparse signal model* [2]. A characterizing aspect is the quest for a strong asymmetry in the measurement-reconstruction process, where the resources needed by the former are minimized at the cost of supplemental computational/power requirements on the reconstruction side [3]. This is consistent with the asymmetry of centralized sensor networks where peripheral nodes send data to a collector counting on larger resources [4].

CS employs measurements taken on *signal segments* in the *discrete-time* domain by projecting them onto a set of *sensing vectors*. On the recovery side, reconstruction relies on numerically solving an *optimization* problem based on the sparsity model. CS is often presented as an *analog to information* technique [5] but many CS applications take the projections once the signal has been acquired in the digital domain. Digital CS is particularly frequent when reading biological data, as in electrocardiograms (ECGs) [6]. In this area, commercial solutions often employ $\Delta\Sigma$ modulators ($\Delta\Sigma$ M) for A/D conversion [7]. In digital CS, the arithmetic cost of projections can be reduced by adopting binary, ternary or otherwise *low resolution* sensing vectors [8], [9]. Among other CS developments, it has recently been realized that in addition to sparsity, more *a priori* knowledge on the input signals may exist, expressed by their average energy distribution over frequency. *Rakeness* based CS [10], uses such knowledge to design sensing vectors that collect more energy from the signal segments, further reducing the number of measurements needed to reach some target reconstruction accuracy. Complementary to rakeness, *denoising* CS uses sensing vectors [11]

The work of S. Callegari, M. Mangia and R. Rovatti was supported by the Project “Metodi di Ottimizzazione e ispirazione biologica nella Rappresentazione del SEgnale (MORSE)” in the framework of the ALMA IDEA 2017 funding initiative by the University of Bologna.

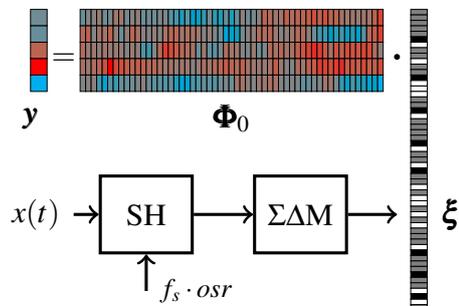


Figure 1. Measurement chain in a CS setup operating on a $\Delta\Sigma$ stream. The input $x(t)$ enters the sample and hold block SH operating a frequency osr times larger than the Nyquist frequency f_s before digitalization by the $\Delta\Sigma$ M and chunking into vectors ξ . The latter are multiplied by the sensing matrix Φ_0 to get the measurements y . In the illustration, the input vector for the CS machinery is ternary, and the sensing matrix is low pass.

that limit energy collection at frequency ranges where the signal is known to be corrupted by noise.

This manuscript provides a proof of principle of how digital CS can be *directly* practiced on the output of a $\Delta\Sigma$ M (Fig. 1). Recall that $\Delta\Sigma$ Ms are oversampled feedback systems capable of delivering a low depth output thanks to their ability to *spectrally shape quantization noise* until its power is pushed almost completely out of the signal band. Their deployment requires the embedding system to ultimately discard such noise [12]. In A/D conversion this is typically the task of a *decimation filter* [13], [14], that also gets rid of the oversampling. In this proposal, decimation is not needed since CS itself is used to deliver an equivalent result, thanks to denoising sensing vectors. Furthermore, it is shown that under suitable conditions CS can inherently undo the oversampling, practicing reconstruction directly at the Nyquist rate. Since oversampling does not change the information content on which CS focuses, the required number of measurements is substantially unaltered with respect to an equivalent Nyquist rate CS system. As a further advantage, because the proposed approach applies the projection operator to segments of the $\Delta\Sigma$ stream that have a low depth nature, the projection arithmetic can be simplified as in setups using binary/ternary sensing vectors, but in a “reversed” fashion requiring no such constraint. Finally, the approach can incorporate rakeness-based adaptation and further denoising strategies.

A test case based on the readout of synthetic ECG waveforms is presented for validation.

II. BACKGROUND

To define a notation, recall that in classical CS m measures are taken on an \mathbb{R}^n vector x by multiplying it by an $\mathbb{R}^{m \times n}$ sensing matrix Φ , with $m < n$ so obtaining an \mathbb{R}^m *measurement vector*

$$y = \Phi x . \quad (1)$$

In reconstruction, one relies on the assumption that a basis exists onto which vectors \mathbf{x} have a sparse representation, namely

$$\mathbf{x} = \Psi\boldsymbol{\alpha} \quad (2)$$

where $\Psi \in \mathbb{R}^{n \times n}$ is a matrix containing the basis vectors as its columns and $\boldsymbol{\alpha} \in \mathbb{R}^n$ has only a few nonzero entries. Reconstruction is practically achieved by solving a minimization problem

$$\arg \min_{\tilde{\boldsymbol{\alpha}} \in \mathbb{R}^n} \|\tilde{\boldsymbol{\alpha}}\|_{\ell_1} \quad \text{so that} \quad \|\mathbf{y} - \Phi\Psi\tilde{\boldsymbol{\alpha}}\|_{\ell_2} < \epsilon \quad (3)$$

where ϵ is a suitable small constant, so obtaining

$$\tilde{\mathbf{x}} = \Psi\tilde{\boldsymbol{\alpha}} \quad (4)$$

as an estimate of \mathbf{x} . All this works as long as Φ has features enabling it to collect sufficiently rich and diverse measures of \mathbf{x} . This requirement can be formalized in various ways, from bounds on the *incoherence* between Φ and Ψ to complex geometrical considerations [10, chapters 1 and 2]. Using *random* entries for the rows of Φ is generally an appropriate choice. The problem (3) is known as *basis pursuit with denoising* and ϵ is meant to tune robustness to noise. When dealing with *physical signals*, the framework above is generally applied assuming that \mathbf{x} is obtained from some signal $x(t)$ by sampling it at some frequency f_s and by chunking it into n sample windows. Frequency f_s is generally assumed to be the Nyquist one.

III. COMPRESSED SENSING OF OVERSAMPLED STREAMS

As a starting point for applying CS to $\Delta\Sigma$ stream, CS of oversampled streams is considered. Intuitively, if one substitutes $\mathbf{x}_o \in \mathbb{R}^{osr \cdot n}$ for \mathbf{x} , where \mathbf{x}_o is \mathbf{x} sampled at $osr \cdot f_s$ — i.e., at an oversampling ratio osr — the framework can still work. This requires the use of a new basis matrix Ψ_o , and a new sensing matrix Φ_o (ideally, with columns and rows in $\mathbb{R}^{osr \cdot n}$, respectively) so that the basis spans the space covering all the possible vectors \mathbf{x}_o . An interesting aspect to consider is whether a relationship can be established between Ψ_o and Ψ and what properties may derive from it.

A. Idealized case: sampling rate change after signal chunking

To begin, assume that the oversampling does not *add information* to the signal segments, namely that an upsampling operator \mathbf{O} exists such that \mathbf{x}_o can be expressed as $\mathbf{O}\mathbf{x}$. In such situation, from $\mathbf{x} = \Psi\boldsymbol{\alpha}$ one gets $\mathbf{x}_o = \mathbf{O}\Psi\boldsymbol{\alpha} = \Psi_o\boldsymbol{\alpha}$, with $\Psi_o = \mathbf{O}\Psi$. In other words, if Ψ is a sparsity basis for \mathbf{x} , then Ψ_o is a sparsity basis for \mathbf{x}_o . Furthermore, with this specific choice for the basis, after a measurement-reconstruction process done on \mathbf{x}_o — leading to the estimation of some vector $\tilde{\boldsymbol{\alpha}}$ (and thus to some estimate $\tilde{\mathbf{x}}_o = \Psi_o\tilde{\boldsymbol{\alpha}}$ of \mathbf{x}_o) — one can readily get $\tilde{\mathbf{x}} = \Psi\tilde{\boldsymbol{\alpha}}$ as a Nyquist rate estimate of the signal segment.

For what concerns \mathbf{O} , the obvious expectation is that upsampling followed by downsampling is a no-op. Formally, \mathbf{O} should be a *right inverse matrix* of the uniform downsampling operator \mathbf{U} described by an $n \times osr \cdot n$ matrix with entries $u_{i,j}$ equal to 1 for $j = osr \cdot i$ and null otherwise (under the assumption of 0-based vector/matrix indexing), so that $\mathbf{U}\mathbf{O} = \mathbf{I}$, where \mathbf{I} is the identity operator. While there is not a unique \mathbf{O} satisfying such requirement, it is easy to find one by an appropriate domain restriction on \mathbf{U} . For instance, one may restrict to signals that do not have any energy at frequency components above $f_s/2$. This would lead to:

$$\mathbf{O} = osr \cdot \mathbf{F}^{-1} \mathbf{Z} \mathbf{F} \mathbf{U}^T \quad (5)$$

where $(\cdot)^T$ indicates transposition, \mathbf{F} is the $osr \cdot n \times osr \cdot n$ discrete Fourier transform (DFT) matrix, \mathbf{F}^{-1} is its inverse, and \mathbf{Z} is in

charge of nullifying all the Fourier components above the $f_s/2$, so that its entries $z_{i,j}$ are: 1 for $i = j \wedge (i < \frac{n}{2} \vee i > \frac{n}{2}(2osr - 1))$; $2^{-1/2}$ if $i = j \wedge (i = \frac{n}{2} \vee i = \frac{n}{2}(2osr - 1))$; and null otherwise. Other choices are obviously possible.

For what concerns the sensing matrix, if $\mathbf{U}\mathbf{O} = \mathbf{I}$, letting $\Phi_o = \Phi\mathbf{U}$ makes $\Phi_o\mathbf{x}_o = \Phi\mathbf{x}$. In other words, for this particular choice of Φ_o measurements taken by CS on the oversampled signal stay the same as those taken by the original CS system operating on signals sampled at f_s . The same holds for any $\Phi_o = \Phi\mathbf{M}$, as long as $\mathbf{M}\mathbf{x}_o = \mathbf{U}\mathbf{x}_o$ for any \mathbf{x}_o in the *image set* of \mathbf{O} . The \mathbf{M} matrices for which this happens are those satisfying $\mathbf{M}\mathbf{O}\mathbf{x} = \mathbf{U}\mathbf{O}\mathbf{x}$ for all \mathbf{x} , namely the *left inverse matrices* of \mathbf{O} .

This means that for the oversampled system, there is no need for the sensing matrix to be random. On the contrary, measurements sufficient for decoding can certainly be gathered via a Φ_o containing rows whose entries are strongly correlated, as long as they are obtained by upsampling the rows of a sensing matrix Φ known to work at the Nyquist rate via a suitable \mathbf{M} . Clearly, the sensing vectors can also be generated *ad hoc*. Yet, there is an obvious potential advantage in getting Φ_o by upsampling Φ as $\Phi\mathbf{M}$. In this case, thanks to the measurements staying the same as in the Nyquist rate CS system, it is not even necessary to introduce an oversampled sparsity basis, and the decoding can be directly practiced at the Nyquist rate by the minimization in (3), or its noise-robust variant.

As a particular case, consider upsampling operators \mathbf{O} that satisfy $\mathbf{O}^T\mathbf{O} = \mathbf{I}$, namely where \mathbf{O} is *column orthonormal*. For such operators \mathbf{O}^T belongs to the previously introduced family. Column orthonormal upsampling operators clearly include $\mathbf{O} = \mathbf{U}^T$ and most standard interpolations belong to the class (at least with some approximation or adjustment). For instance, the \mathbf{O} matrix in Eqn. (5) has this property, apart from a scaling factor $1/\sqrt{osr}$ that can be corrected on the measurement vectors. The adoption of $\Phi_o = \Phi\mathbf{O}^T$ with the \mathbf{O} in Eqn. (5) or given by other interpolations makes the measurement vectors inherently low-pass. This choice provides automatic denoising in case a practical implementation causes \mathbf{x}_o to include disturbances above $f_s/2$, because it makes negligible or null the energy collectible in that frequency range [11]. Such property will be of utter importance in Sect. IV, where CS is practiced on $\Delta\Sigma$ streams, to automatically treat quantization noise.

B. Real case: different sampling rates before chunking

Unfortunately, the discussion so far is challenged when one tries to apply it to real world problems where \mathbf{x} and \mathbf{x}_o *both* derive from a single physical signal $x(t)$ that extends arbitrarily long in time. To consider such case, assume that $x(t)$ is bandlimited to $f_s/2$, so that it can be sampled at any frequency equal or larger than f_s without aliasing. Also assume that two different samplings of it exist, $\hat{x}[j] = x(j/f_s)$ and $\hat{x}_o[j] = x(j/(osr \cdot f_s))$. Finally, define \mathbf{x} and \mathbf{x}_o as two *windows* of $\hat{x}[j]$ and $\hat{x}_o[j]$. Even if \mathbf{x} and \mathbf{x}_o are fully aligned (namely, if the entries in \mathbf{x} are $\hat{x}[i + \bar{j}]$ for $i = 0, \dots, n - 1$ and the entries of \mathbf{x}_o are $\hat{x}_o[i + \bar{j} \cdot osr]$ for $i = 0, \dots, osr \cdot n - 1$), the relationship from \mathbf{x}_o to \mathbf{x} is not injective. In a nutshell, there can be different vectors \mathbf{x}_o for the same \mathbf{x} and it is impossible to define an operator \mathbf{O} such that $\mathbf{x}_o = \mathbf{O}\mathbf{x}$, because the information in \mathbf{x} alone is insufficient to determine \mathbf{x}_o . It is so because the different sampling rates exist already *before* the signal chunking. Furthermore, one does not have any control on them because in the context of the current application the faster one is the sampling practiced by the $\Delta\Sigma$ data converter and the slower one is the reference sampling that one would get out of the converter decimation filter.

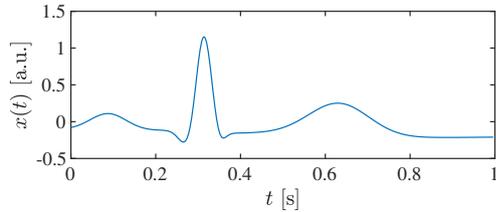


Figure 2. Synthetic ECG used for validation (one segment).

Incidentally, a related issue can be observed in the frequency domain noticing that the DFTs of x and x_o cannot overlap up to frequency $f_s/2$, with the DFT of x_o being null elsewhere. This is because the multiplication of $x(t)$ by the rectangular chunking window used to get x and x_o corresponds to a convolution operator in the frequency domain causing *leakage* [15]. Not only the two DFTs differ in band: at the higher rate leakage pushes energy above $f_s/2$.

The lack of an unequivocally defined oversampling operator \mathcal{O} calls into question the applicability of the idealized results in Sect. III-A. In fact it now appears not fully rigorous to attempt CS on x_o using sensing vectors with no energy above frequency $f_s/2$. Furthermore, if there is not an unambiguous way to get the CS matrices Φ_o and Ψ_o used at the higher sampling rate from the matrices Φ and Ψ used in the Nyquist rate system, the possibility to use the CS measurements taken at the higher rate to directly get an estimate of the signal at the Nyquist rate via $\tilde{x} = \Psi\tilde{\alpha}$ as in Sect. III-A becomes dubious.

Still, if the signal segments are long enough and the oversampling is not too large, one may expect the results in Sect. III-A to remain applicable with good approximation when a standard kernel \mathcal{O} is used for the upsampling of the basis matrix Ψ into Ψ_o and as \mathcal{O}^T for the oversampling of the sensing matrix Φ into Φ_o . Certainly, one shall expect performance to become somehow sensitive to the specific oversampling operator being used. For instance, the one in Eqn (5) may not be the best one (because of the *complete* erasure of high frequency components). Similarly, the choice of an appropriate denoising parameter ϵ in basis pursuit may become more cogent.

C. Validation of performance expectations in the real case

Since the purpose of this paper is mainly to offer a proof of principle, the exploration of the best choices for \mathcal{O} and ϵ will be left to further works. Conversely, the focus will be on a mere validation of the previous expectations by a practical test case. The latter involves synthetic ECG signals with rates from 50 bpm to 75 bpm (fig. 2) and a 50 dB SNR for the Nyquist rate system. For the oversampled system, the input SNR is set to be $10 \log_{10} osr$ dB lower with noise spreading uniformly on the whole available bandwidth, in order to test the denoising ability of CS at frequencies above the signal bandwidth. For the reference Nyquist rate system, signal segments are 256-samples long and $osr = 64$ and a *symmlet* wavelet sparsity basis is employed. The setup includes two types of tests. The first one relies on random sensing vectors, while the second one uses vectors designed following the *rakeness*-based CS paradigm [10]. In either case, for the system working on oversampled streams the sensing matrix and the basis are obtained from the corresponding matrices of the Nyquist rate system by standard interpolations.

Fig. 3 compares the performance of the oversampled CS setup to that of Nyquist rate CS in terms of reconstructed SNR (estimated over 64 runs) for growing numbers of measures m . As evident

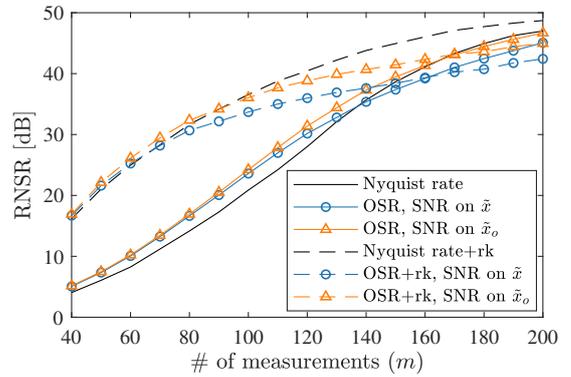


Figure 3. Comparison of CS practiced on Nyquist rate and oversampled signals in terms of reconstruction SNR (RNSR). Black solid curve is the baseline for CS practiced at the Nyquist rate with random sensing vectors. Solid curves always correspond to using random vectors in the sensing matrix while dashed curves correspond to using sensing matrices optimized for *rakeness*. Curves with the circle marker measure the SNR on Nyquist rate reconstructed signals \tilde{x} , while those with the triangular marker measure SNR on the oversampled reconstructed signals \tilde{x}_o .

from the plots, the performance of the system operating on the oversampled signal is consistent with that of the Nyquist rate system and a performance loss is only seen at large SNR levels (for the particular case above 35 dB, which can be set as a target for being often considered the minimum value for medical diagnosis of ECG waveforms [16]). Most important, the system operating on the oversampled streams is capable to approximately get to the same performance of the non-oversampled one notwithstanding the input SNR being worse by about 20 dB. This follows from the ability to automatically discard the out of band noise. Surprisingly, the higher rate system does even better than the Nyquist rate one when the number of measurements is particularly low. In the *rakeness* optimized setup one always gets better performance than in the corresponding non optimized test, showing that the advantages of the *rakeness* approach are preserved. In this case, the performance loss of the higher rate system is already evident at a lower number of measurements, though, due to the higher SNRs. Finally, the possibility to directly decode at the Nyquist rate after CS is practiced on the oversampled stream is also validated, showing a slightly larger performance loss. All this is quite consistent with the expectations from Sect. III-B.

IV. COMPRESSED SENSING OF $\Delta\Sigma$ STREAMS

The framework for CS of oversampled signals may have some interest by itself but not a significant practical value. Its opportunity lays in setting the basis for the CS of $\Delta\Sigma$ streams. Recall that $\Delta\Sigma$ is often adopted in data acquisition because it enables the use of quantizers with a quite limited number of levels (often just 2 or 3), so relaxing linearity related issues [13]. Low resolution is paid by the introduction of a large amount of quantization noise. However, $\Delta\Sigma$ smartly use oversampling to push such noise *above* the signal bandwidth, so that it can be removed later on in the processing chain.

Given the proof of principle nature of this work, as the sole example Fig. 4 shows a 2nd order single quantizer $\Delta\Sigma$, with a 3 level output. Fig. 5, illustrates the kind of spectra that can be obtained from its application to ECG signals at $osr = 64$. While much more sophisticated forms of noise shaping can in principle be pursued [12], [13], profiles such as those in the plot are common as a reasonable trade-off with the cost of the modulator. Because of this compromise,

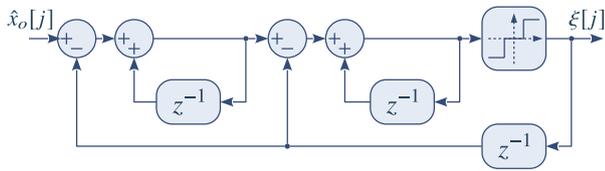


Figure 4. Second order, single quantizer, 3-level $\Delta\Sigma$ modulator used for validation.

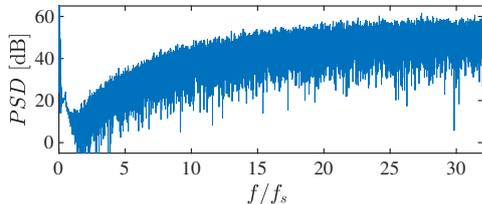


Figure 5. Power spectral density (PSD) plot showing Quantization noise shaping by the $\Delta\Sigma$ in Fig. 4 operating on the test ECG signals.

residual quantization noise in the signal band can significantly cap SNR. For ECG signals (that sit most of the time around zero to only occasionally peak, so representing an unfavourable case), the $\Delta\Sigma$ in Fig. 4 can provide an in-band theoretical SNR just over 40 dB at $osr = 64$, i.e., slightly above the minimum required by the application.

Following the theory in Sect. III, one can treat the $\Delta\Sigma$ stream as a case of an oversampled stream that is particularly noisy above the Nyquist frequency. Because the ability of CS to remove high frequency noise has already been validated, one can think of feeding the $\Delta\Sigma$ stream directly into the CS system. With this, the decimation filter becomes unneeded. More interestingly, the projection arithmetic gets simplified from a multiply-and-accumulate process into a conditional-sum process by the ternary nature of the $\Delta\Sigma$ output. This is similar to [8], [9], but without imposing constraints on the sensing matrix. Furthermore, the downsampling is also saved, since — as already shown — it can be done implicitly during signal reconstruction.

Fig. 6 validates these concepts by providing SNR curves directly comparable to those in Fig. 3. The plots are totally consistent, with the sole exception that the curves in Fig. 6 reach their plateau earlier, at about 40 dB, this being the intrinsic SNR cap imposed by the $\Delta\Sigma$ used for the tests. Reconstruction SNRs at about 35 dB remain achievable at a reasonable compression factor.

V. CONCLUSIONS

The possibility of practicing CS directly on $\Delta\Sigma$ stream has been validated both by experiments and through the development of a framework for CS of oversampled signals. The presented examples showed that measurement matrices suitable for $\Delta\Sigma$ streams can be readily obtained from others known to work at the Nyquist rate, including matrices for rakesness-based CS. In future work, the ad hoc crafting of sensing matrices for the task will likely be addressed, as part of an ongoing effort in evaluating the suitability and potential of pulse-based signals in processing tasks.

REFERENCES

- [1] D. L. Donoho, “Compressed sensing”, *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006. DOI: 10.1109/TIT.2006.871582
- [2] E. J. Candes and T. Tao, “Decoding by linear programming”, *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005. DOI: 10.1109/TIT.2005.858979

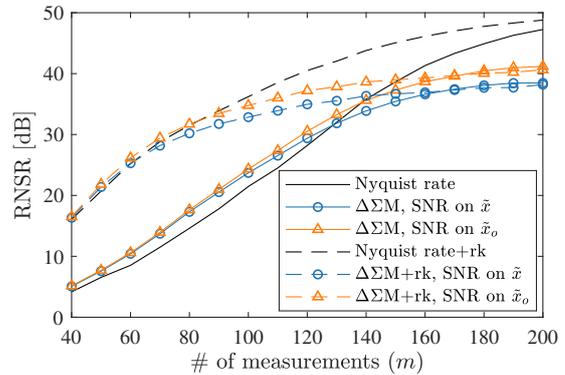


Figure 6. Comparison of conventional CS and CS practiced on $\Delta\Sigma$ streams. The plots can be read as those in Fig. 3.

- [3] E. J. Candes and M. B. Wakin, “An introduction to compressive sampling”, *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008. DOI: 10.1109/MSP.2007.914731
- [4] J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak, “Compressed sensing for networked data”, *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 92–101, Mar. 2008. DOI: 10.1109/MSP.2007.914732
- [5] J. N. Laska, S. Kirolos, M. F. Duarte, T. S. Ragheb, R. G. Baraniuk, and Y. Massoud, “Theory and implementation of an analog-to-information converter using random demodulation”, in *Proc. of ISCAS’07*, May 2007, pp. 1959–1962. DOI: 10.1109/ISCAS.2007.378360
- [6] F. Pareschi, P. Albertini, G. Frattini, M. Mangia, R. Rovatti, and G. Setti, “Hardware-algorithms co-design and implementation of an analog-to-information converter for biosignals based on compressed sensing”, *IEEE Trans. Biomed. Circuits Syst.*, vol. 10, no. 1, pp. 149–162, Feb. 2016. DOI: 10.1109/TBCAS.2015.2444276
- [7] K. Soundarapandian and M. Berarducci, “Analog front-end design for ECG systems using Delta-Sigma ADCs”, Texas Instruments, Application Report SBAA160A, Apr. 2010. [Online]. Available: www.ti.com/lit/an/sbaa160a/sbaa160a.pdf
- [8] M. Mangia, F. Pareschi, V. Cambarelli, R. Rovatti, and G. Setti, “Rakesness-based design of low-complexity compressed sensing”, *IEEE Trans. Circuits Syst. I*, vol. 64, no. 5, pp. 1201–1213, May 2017. DOI: 10.1109/TCSI.2017.2649572
- [9] J. Zhang, Z. Gu, Z. L. Yu, and Y. Li, “Energy-efficient ECG compression on wireless biosensors via minimal coherence sensing and weighted ℓ_1 minimization reconstruction”, *IEEE J. Biomed. Health Inform.*, vol. 19, no. 2, pp. 520–528, Mar. 2015. DOI: 10.1109/JBHI.2014.2312374
- [10] M. Mangia, F. Pareschi, V. Cambarelli, R. Rovatti, and G. Setti, *Adapted Compressed Sensing for Effective Hardware Implementations*. Springer, 2018. DOI: 10.1007/978-3-319-61373-4
- [11] M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, “Implicit notch filtering in compressed sensing by spectral shaping of sensing matrix”, in *Proc. of ISCAS’16*, May 2016, pp. 257–260. DOI: 10.1109/ISCAS.2016.7527219
- [12] S. Callegari and F. Bizzarri, “Optimal design of the noise transfer function of $\Delta\Sigma$ modulators: IIR strategies, FIR strategies, FIR strategies with preassigned poles”, *Signal Processing*, vol. 114, pp. 117–130, Sep. 2015. DOI: 10.1016/j.sigpro.2015.02.001
- [13] S. Pavan, R. Schreier, and G. C. Temes, *Understanding Delta-Sigma data converters*, 2nd ed. Wiley-IEEE Press, 2017.
- [14] M. Laddomada, “Comb-based decimation filters for $\Sigma\Delta$ A/D converters: Novel schemes and comparisons”, *IEEE Trans. Signal Process.*, vol. 55, no. 5, pp. 1769–1779, May 2007. DOI: 10.1109/TSP.2006.890822
- [15] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Prentice-Hall, 1989.
- [16] Y. Zigel, A. Cohen, and A. Katz, “The weighted diagnostic distortion (WDD) measure for ECG signal compression”, *IEEE Trans. Biomed. Eng.*, vol. 47, no. 11, pp. 1422–1430, Nov. 2000. DOI: 10.1109/TBME.2000.880093