

Compressed sensing by phase change memories: Coping with encoder non-linearities

Original

Compressed sensing by phase change memories: Coping with encoder non-linearities / Paolino, C.; Antolini, A.; Pareschi, F.; Mangia, M.; Rovatti, R.; Scarselli, E. F.; Gnudi, A.; Setti, G.; Canegallo, R.; Carissimi, M.; Pasotti, M.. - STAMPA. - 2021:(2021). (Intervento presentato al convegno 53rd IEEE International Symposium on Circuits and Systems, ISCAS 2021 tenutosi a Daegu, Corea del Sud [virtual] nel May 22-28, 2021) [10.1109/ISCAS51556.2021.9401176].

Availability:

This version is available at: 11583/2918014 since: 2021-08-17T17:47:13Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/ISCAS51556.2021.9401176

Terms of use:

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

©2021 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 by Phase Change Memories: coping with encoder non-linearities

Carmine Paolino*, Alessio Antolini^{†,‡}, Fabio Pareschi^{*,‡}, Mauro Mangia^{†,‡}
 Riccardo Rovatti^{†,‡}, Eleonora Franchi Scarselli^{†,‡}, Antonio Gnudi^{†,‡}, Gianluca Setti^{*,‡}
 Roberto Canegallo[§], Marcella Carissimi[§], Marco Pasotti[§]

* DET – Politecnico of Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy.

[†] DEI – University of Bologna, viale Risorgimento 2, 40136 Bologna, Italy.

[‡] ARCES – University of Bologna, via Toffano 2/2, 40125 Bologna, Italy.

[§] STMicroelectronics, Agrate Brianza, Italy.

email: carmine.paolino@polito.it, alessio.antolini2@unibo.it

Abstract—Several recent works have shown the advantages of using phase-change memory (PCM) in developing brain-inspired computing approaches. In particular, PCM cells have been applied to the direct computation of matrix-vector multiplications in the analog domain. However, the intrinsic nonlinearity of these cells with respect to the applied voltage is detrimental. In this paper we consider a PCM array as the encoder in a Compressed Sensing (CS) acquisition system, and investigate the effect of the non-linearity of the cells. We introduce a CS decoding strategy that is able to compensate for PCM nonlinearities by means of an iterative approach. At each step, the current signal estimate is used to approximate the average behaviour of the PCM cells used in the encoder. Monte Carlo simulations relying on a PCM model extracted from an STMicroelectronics 90 nm BCD chip validate the performance of the algorithm with various degrees of nonlinearities, showing up to 35 dB increase in median performance as compared to standard decoding procedures.

I. INTRODUCTION

Due to its successful integration in CMOS fabrication processes, its high throughput and read/write endurance, phase-change memory (PCM) can be considered as a valid technology for next-generation non-volatile memories (NVMs) [1]. PCM relies on the reversible transition of a chalcogenide material between its crystalline (or SET) and amorphous (or RESET) state. The electrical resistivity observed in the amorphous state decreases by several orders of magnitude in the crystalline one. Minimum impact on process complexity is guaranteed by the embedded PCM (ePCM) technology which includes also high-power and high-voltage integrated components [2], [3].

A widely investigated application of PCM technology is Analog In-memory Computing (AIMC), which is an enabling support for enhancing the next generation of application-oriented hardware. In the field of AIMC, analog elements are used to perform matrix operations simply exploiting Ohm's and Kirchhoff's Laws. With reference to the *ENCODER* in Fig. 1, given a cell with conductance g , a single multiplication is achieved by applying an input voltage v to the cell and reading the output current $i = gv$. If the outputs of several g_j cells, each with its own applied voltage v_j , are combined, a sum of products is obtained. It is then possible to interpret the whole memory as a conductance matrix G of size $m \times n$; applying a voltage vector to each row, a matrix-




Fig. 1. Architecture of a PCM-based system for CS applications with the proposed decoder. The encoder implements the ideal, binary sensing matrix A by PCM conductances. The decoder reconstructs the original signal by iteratively computing the estimate of the real sensing matrix \hat{G} .

vector multiplication is obtained [4]. All the computation is carried out inside the analog memory chip, avoiding digital data to be transferred between conventional memory chips and processors. Regrettably, PCM cells present several drawbacks when used as analog multipliers. One of these is the non-linearity of their $i(v)$ characteristic, which translates into a dependence of g on the applied voltage v .

Compressed Sensing (CS) is an ideal candidate for AIMC. CS is a signal acquisition technique capable of reconstructing a signal from a number of scalars, called *measurements*, that is smaller than in a traditional Nyquist-rate system. The encoding step producing those measurements is indeed a matrix-vector product and can be implemented by a PCM array. Signal decoding, however, requires the exact a priori knowledge of G , which is not possible if it is signal dependent.

In this work, we propose an iterative technique allowing the recovery of the acquired signal even when measurements are produced by strongly varying conductances. At its core, the proposed method iteratively executes known decoding algorithms. At each iteration, the input signal is estimated, then a better representation of the signal-dependent conductance matrix is obtained and the updated matrix exploited in the subsequent iteration until convergence. The technique shows promising results, enabling the use of highly voltage-dependent SET conductance states which would have unsatisfactory performance with a standard decoder.

The adopted technology and its characterization are described in Section II, while an overview of Compressed

Sensing theory is outlined in Section III. The proposed technique is defined in Section IV. Section V presents the results obtained in system-level, Monte Carlo simulations. Finally, the conclusion is drawn.

II. PCM TECHNOLOGY AND CHARACTERIZATION

We consider an ePCM test chip designed and manufactured by STMicroelectronics [1] in a 90-nm BJT-CMOS-DMOS (BCD) technology featuring a specifically optimized Ge-rich Ge-Sb-Te (GST) alloy. The chip is intended for digital storage in automotive applications and contains 8 independent 256-kB macrocells, together with voltage and current regulators. The ePCM elementary cell is based on an NMOS selector [5] and occupies $0.19\mu\text{m}^2$ of silicon area. A custom PCM evaluation board (designed ad-hoc for testing purposes) has been employed.

A total of 5120 cells have been programmed with SET-pulses of different intensities, and immediately characterized in the i/v domain. Typical behaviors have been obtained by averaging all the i/v curves, associated to different physical cells, falling within a $\pm 5\%$ of an arbitrarily selected reference current i_{ref} , evaluated at a reference voltage v_{ref} . Fig. 2a shows three such average behaviors, each corresponding to a different reference current and normalized with respect to the maximum value of the state #1 curve. The arbitrary reference points have been identified so as to highlight significantly different conductance behaviours. The saturation of current profiles clearly visible at $v > 0.8$ is due to the reading circuit transistors entering the saturation region. In the following, we will consider the three curves depicted in the figure as the nominal $i(v)$ behavior of typical cells in three different programming states, denoted as #1, #2 and #3. The corresponding large-signal, normalized conductances defined as $g = i/v$ are shown in Fig. 2b in which non-constant trends reach a maximum around $v = 0.5$

Although we expect conductances in state #1 to perform better in terms of reconstruction quality, they lead to the least power-efficient solution, having larger absolute conductance values. Hence the motivation to devising a strategy enabling the use of lower, more voltage-sensitive conductance states.

III. COMPRESSED SENSING

Compressed Sensing [6]–[8] enables simultaneous signal acquisition and compression by performing a low-complexity linear encoding phase which can be considered as a trivial matrix-vector product. With respect to a standard compression algorithm, complexity is moved to the decoder side, where a minimization problem needs to be solved for signal recovery. CS can be applied if the n -dimensional input signal $x \in \mathbb{R}^n$ (i.e., a sequence of n consecutive samples x_j , $j = 1, \dots, n$) is *sparse* when represented in a suitable vector space. Mathematically, being $D \in \mathbb{R}^{n \times n}$ the *sparsity* basis, and $x = D\xi$, x is sparse if only a small fraction of the coefficients in ξ are significantly non-null. The acquisition of x requires the computation of m measurements y_1, \dots, y_m , with $m \ll n$, each computed as: $y_k = \sum_{j=1}^n a_{kj}x_j$, ($k = 1, \dots, m$) where




Fig. 2. (a) Average, normalized i/v characteristics of PCM in three different states. The curves have all been normalized with respect to the maximum current in state #1, I_0 ($i = I/I_0$). The applied voltages have been normalized as well with respect to the maximum applied voltage V_0 ($v = V/V_0$). (b) Large-signal, normalized conductance $g = i/v$.

the a_{kj} 's are suitable weights. In matrix form, the previous expression can be rewritten as:

$$y = Ax, \quad (1)$$

where $A = [a_{kj}]$ is called the *sensing* or *acquisition* matrix.

The cost of such a simple acquisition phase is a complex signal decoding, that is typically performed on a device with high computational capabilities. Under assumptions that are usually verified [6], [9], the recovered input signal \hat{x} can be found by looking at the sparsest among all signals generating the same measurement vector y . Formally, this corresponds to the solution of the following minimization:

$$\hat{x} = D \cdot \operatorname{argmin}_{\xi} |\xi|_0 \quad \text{s.t. } AD\xi = y. \quad (2)$$

where $|\cdot|_0$ is the ℓ_0 pseudonorm, which counts the number of nonzero elements in its argument, and that can be conveniently replaced [9] by the easier-to-compute ℓ_1 standard norm.

The general theory of CS states that the solution of (2) is the original input signal for Gaussian, i.i.d. weights a_{kj} . Although, it has been shown that x can be retrieved without significant performance degradation even by generating the a_{kj} with a sub-gaussian distribution limited to a few discrete levels only [10], allowing a reduction in complexity of the encoder hardware implementation. We will exploit this and consider a binary encoding, i.e., we set the sensing matrix entries either to 0 or to 1 [11]–[13], thus minimizing hardware complexity.

Note however that correct decoding requires the exact knowledge of A , which is impossible in our analog implementation of the matrix-vector product because of the voltage dependence of the PCM cells conductance that replaces (1) with

$$y = A(v) \cdot x.$$

Even assuming the exact knowledge of the voltage dependence of a PCM cell, the decoder is unable to know the effective

sensing matrix $A(x)$ since x is unknown. Some preliminary studies on signal decoding after a non-linear sensing can be found [14], [15], though they cannot be used in situations with strong variations such as those of Fig. 2b. In the following, we will describe a technique that copes with the signal dependence of A , allowing a significant performance increase with respect to the straightforward application of known decoding algorithms.

In the following, as shown also in Fig. 1, we will refer to $A \in \{0, 1\}^{m \times n}$ as the ideal, binary sensing matrix, while $G = [g_{kj}]$ will denote its physical realization. The zeroes are implemented through PCM cells in the RESET state (i.e., high-impedance), for which it is reasonable to assume $g_{kj} = 0$. Conversely, for the non-zero values we use PCMs in the SET state, such that $g_{kj} = g(x_j)$, according to the conductance curves in Fig. 2b.

IV. PROPOSED DECODING TECHNIQUE

In principle, our technique can use any of the known CS decoding algorithms [16]–[18]. In the following, we will refer to the inner algorithm as the *minimizer*, while we will refer to the entire reconstruction procedure as the *decoder*. The purpose of the minimizer is to solve the minimization problem described in (2) by assuming a constant, known sensing matrix.

The proposed decoding scheme is formally described in Algorithm 1 and depicted in Fig. 1. Given the ideal sensing matrix A , the observed measurements y , a reference scalar value x_{ref} and a model describing the voltage dependence of the conductance elements $g(x) : \mathbb{R} \rightarrow \mathbb{R}$, the initial estimate of the conductance matrix $\hat{G}_{|0}$ is evaluated as $A \cdot g(x_{\text{ref}})$. The matrix $\hat{G}_{|0}$ is used by the minimizer to compute $\hat{x}_{|1}$, the initial signal estimate. The conductance matrix is then updated by computing $g((\hat{x}_{|1})_j) \cdot a_{kj}$ for each of its elements, and the new signal estimate $\hat{x}_{|2}$ is evaluated to be used in a further iteration. If the signal estimates $\hat{x}_{|p}$ converge, the conductance matrices $\hat{G}_{|p}$ converge, as well.

Algorithm 1: Iterative CS decoder for signal dependent sensing matrices

Data: $A, y, g(x), x_{\text{ref}}$

Result: \hat{x}

$\hat{G} \leftarrow g(x_{\text{ref}}) \cdot A$

repeat

$\hat{x} \leftarrow \text{minimizer}(y, \hat{G} \cdot D)$

$\hat{G}_{k,j} \leftarrow g(\hat{x}_j) A_{k,j} \quad \forall k = 1, \dots, m; j = 1, \dots, n$

Compute convergence metric

until *convergence* or *timeout*

Note that each iteration of the decoding algorithm requires the solution of the minimization problem in (2), a costly operation in itself. Hence, a possible stopping criterion would have to balance the resources spent against the achievable reconstruction quality.

V. NUMERICAL RESULTS

The validation of the proposed technique has been conducted by performing system-level simulations with synthetically generated signals. Given one of the three conductance models depicted in Fig. 2b, each signal instance is encoded by means of a PCM-based array whose cells in the SET state (implementing a 1) are all described by the selected model. From the corresponding measurements vector y , the decoder then tries to recover the original signal, knowing the conductance model and iteratively using it to estimate the effective, signal-dependent conductances in the encoder.

The iterative reconstruction technique has been tested with 1000 signal instances, sparse in the Discrete Cosine Transform basis, with $n = 256$, a sparsity level of 26 non-null coefficients, and a high-pass spectrum profile [10, Section 2.3]. The sensing matrices are binary, of size 128×256 , with ideal zeroes, and ones being implemented by one of the conductance models in Fig. 2b. The minimization procedure being applied is the Orthogonal Matching Pursuit (OMP) [16].

Before the encoding step is performed, white gaussian noise is added to the signal to reach a predetermined value of Signal to Noise ratio (SNR), denoted as ISNR. Up to some denoising effects, which Compressed Sensing is known to possess [10], this defines a target for correct recovery of the original signal and makes the experimental setup more realistic. The reconstruction quality obtained at iteration p is measured by the Reconstruction Signal-to-Noise Ratio (RSNR), defined as:

$$\text{RSNR} = \text{SNR}(x, \hat{x}_{|p}) = 20 \log \frac{\|x\|_2}{\|x - \hat{x}_{|p}\|_2}$$

Two sets of results are shown. The first, in Fig. 3a, observes the reconstruction quality as the number of iterations is increased, the second, in Fig. 3b, studies the effect of the dynamic range of the applied voltages, emphasizing the performance gain between the first and the last iterations.

The results in Fig. 3(a) have been obtained for the dynamic range $[0, 0.8]$, thus avoiding the region where the current saturates, which we expect to be problematic to work with. The empirical Cumulative Distribution Function (CDF) of all the RSNR datapoints being obtained at a fixed number of iterations clearly show how the reconstruction quality grows as more runs are performed, with a saturation of the obtainable quality in a number of iterations that depends on the behavior of the conductances.

For conductances of type #1, the initial quality has a median of 33 dB and in less than 5 iterations saturates around a median of 49 dB. Conductances of type #3, result in unacceptable RSNR values, which do not improve significantly by increasing the number of runs due to the higher voltage dependence. Conductances of type #2, show the most interesting behavior, starting with a median RSNR of 12 dB, they rise to 30 dB in 5 iterations and 48 dB in 20 runs. The performance increase makes this programming state, whose performance at iteration 1 is insufficient for virtually all applications, worth considering, with the advantage of a reduction in energy consumption. Indeed, the expected power consumption (accounting for the probability distribution of the input signal samples) of the




Fig. 3. (a) Empirical CDFs of RSNR values for the three conductance models, with input signals in the [0:0.8] dynamic range at a selected number of iterations for the decoding algorithm. (b) Empirical CDFs of RSNR values for the three conductance models, 30 iterations, at different dynamic ranges for the input signals. Dotted curves represent the results at the first iteration. Dash-dotted vertical lines, both in (a) and (b), define the ISNR level, i.e. the target reconstruction quality. (c) Mapping of RSNR values against ASNR, at a fixed dynamic range [0:0.8]. For each of the 1000 signal instances, 30 data points are represented, one for each iteration of the decoding procedure. The horizontal line represents the ISNR value. The oblique line is the plane bisector.

PCM array decreases by 28% if conductances of type #2 are chosen in place of type #1.

The effect of different input voltage dynamic ranges (DR) depicted in Fig. 3(b) is studied at constant noise power, i.e. the ISNR decreases at smaller dynamic ranges, as expressed by:

$$\text{ISNR}|_{\text{DR}} = \text{ISNR}|_{\text{DR}_{\text{ref}}} - 20 \log \frac{\text{DR}_{\text{ref}}}{\text{DR}}.$$

Being the effective ISNR the target reconstruction quality, maximizing it by making use of the largest possible dynamic range is desirable. However the higher the allowed voltage excursions, the higher the variation of the conductances, the lower the chances that the iterative technique will be able to recover the original signal. This trade-off suggests the presence of an optimal dynamic range, as indeed shown in the figure. The top two rows, corresponding to conductances of type #1 and #2, confirm the expectation that the dynamic range has to be maximized, only avoiding the region where the access transistors saturate, resulting in a low differential conductance. The results for the [0:1] dynamic range show indeed a long tail towards RSNR values lower than those of the [0:0.8] range, affecting around 25% of the tested signal instances. Conversely, for conductances of type #3, the ISNR value is reached only for the minimum dynamic range, [0:0.2]. Limiting the allowed dynamics to constrain the observed variability is indeed the typical design choice in this context, and under such a severe variation the proposed technique is not able to extend the range further than that.

Up to this point, no convergence metric has been defined, and the proposed decoding technique has been characterized against a predefined number of iterations. Since we have observed how the obtainable reconstruction quality saturates in a number of iterations that depends heavily on the behavior of the conductance, being lower for less variable ones, the

definition of an efficient halting criterion is highly desirable to minimize the number of runs.

A good candidate for a convergence metric is $\|\hat{x}_{|p} - \hat{x}_{|p-1}\|_2$, since this also implies convergence of the effective sensing matrices $\hat{G}_{|p}$. It is convenient to observe the SNR value computed from this difference, named Algorithmic SNR (ASNR), defined as

$$\text{ASNR} = \text{SNR}(\hat{x}_{|p} - \hat{x}_{|p-1}) = 20 \log \frac{\|\hat{x}_{|p}\|_2}{\|\hat{x}_{|p} - \hat{x}_{|p-1}\|_2},$$

and depicted in Fig. 3c against the corresponding values of RSNR, for all the instances used in Fig. 3a. Empirical evidence supports that RSNR and ASNR are strongly linked (at least until the ISNR is reached) though the computation of the latter does not need the knowledge of the original signal x .

Observing that the clouds corresponding to different programming states reach different maximum levels of ASNR, defining convergence by setting a threshold on the observed ASNR is not efficient, since the threshold itself would have to be adapted to the conductance type. Moreover, it has been empirically observed that several signal instances are not able to achieve the threshold and run out of iterations nonetheless. A promising approach yet to be explored is to monitor the ASNR for the current signal instance until a local maximum is reached.

VI. CONCLUSION

An iterative decoding procedure for Compressed Sensing has been presented. It effectively copes with the variability of signal-dependent sensing matrices implemented by Phase Change Memory arrays and allows a substantial extension of the dynamic voltage range. Numerical simulations prove the validity of the technique, enabling the use of conductance states with a stronger voltage dependence and leading to a significant performance increase while reducing the overall power consumption of the PCM array.

REFERENCES

- [1] G. W. Burr, B. N. Kurdi, J. C. Scott, C. H. Lam, K. Gopalakrishnan, and R. S. Shenoy, "Overview of candidate device technologies for storage-class memory," *IBM Journal of Research and Development*, vol. 52, no. 4/5, pp. 449–464, Jul. 2008. doi: 10.1147/rd.524.0449
- [2] G. D. Sandre *et al.*, "A 90nm 4Mb embedded phase-change memory with 1.2V 12ns read access time and 1MB/s write throughput," in *2010 IEEE International Solid-State Circuits Conference - (ISSCC)*, Feb. 2010, pp. 268–269. doi: 10.1109/ISSCC.2010.5433911
- [3] M. Pasotti *et al.*, "A 32-KB ePCM for Real-Time Data Processing in Automotive and Smart Power Applications," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 7, pp. 2114–2125, Jul. 2018. doi: 10.1109/JSSC.2018.2828805
- [4] Z. Sun, G. Pedretti, E. Ambrosi, A. Bricalli, W. Wang, and D. Ielmini, "Solving matrix equations in one step with cross-point resistive arrays," *Proceedings of the National Academy of Sciences*, vol. 116, no. 10, pp. 4123–4128, Mar. 2019. doi: 10.1073/pnas.1815682116
- [5] M. Carissimi *et al.*, "2-Mb Embedded Phase Change Memory With 16-ns Read Access Time and 5-Mb/s Write Throughput in 90-nm BCD Technology for Automotive Applications," *IEEE Solid-State Circuits Letters*, vol. 2, no. 9, pp. 135–138, Sep. 2019. doi: 10.1109/LSSC.2019.2935874
- [6] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. on Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006. doi: 10.1109/TIT.2005.862083
- [7] D. L. Donoho, "Compressed Sensing," *IEEE Trans. on Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006. doi: 10.1109/TIT.2006.871582
- [8] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, Mar. 2008. doi: 10.1109/MSP.2007.914728
- [9] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. on Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005. doi: 10.1109/TIT.2005.858979
- [10] M. Mangia, F. Pareschi, V. Cambareri, R. Rovatti, and G. Setti, *Adapted compressed sensing for effective hardware implementations: A design flow for signal-level optimization of compressed sensing stages*. Springer International Publishing, 2018.
- [11] H. Mamaghani, N. Khaled, D. Atienza, and P. Vandergheynst, "Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes," *IEEE Trans. on Biom. Eng.*, vol. 58, no. 9, pp. 2456–2466, Sep. 2011. doi: 10.1109/TBME.2011.2156795
- [12] M. Shoaran, M. H. Kamal, C. Pollo, P. Vandergheynst, and A. Schmid, "Compact low-power cortical recording architecture for compressive multichannel data acquisition," *IEEE Trans. Biomed. Circuits Syst.*, vol. 8, no. 6, pp. 857–870, Dec. 2014. doi: 10.1109/TB-CAS.2014.2304582
- [13] M. Mangia, D. Bortolotti, F. Pareschi, A. Bartolini, L. Benini, R. Rovatti, and G. Setti, "Zeroing for HW-efficient compressed sensing architectures targeting data compression in wireless sensor networks," *Microprocessors and Microsystems*, vol. 48, pp. 69–79, Feb. 2017. doi: 10.1016/j.micpro.2016.09.007
- [14] T. Blumensath, "Compressed sensing with nonlinear observations and related nonlinear optimization problems," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3466–3474, Jun. 2013. doi: 10.1109/TIT.2013.2245716
- [15] H. Chen, H. T. Kung, and M. Comiter, "Nonlinear compressive sensing for distorted measurements and application to improving efficiency of power amplifiers," in *2017 IEEE international conference on communications (ICC)*, May 2017, pp. 1–7. doi: 10.1109/ICC.2017.7996386
- [16] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. on Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007. doi: 10.1109/TIT.2007.909108
- [17] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, May 2009. doi: 10.1016/j.acha.2008.07.002
- [18] E. Van den Berg and M. P. Friedlander, "Sparse optimization with least-squares constraints," *SIAM J. on Optimization*, vol. 21, no. 4, pp. 1201–1229, 2011. doi: 10.1137/100785028