

Extended Reconstruction Approaches for Saturation Measurements Using Reserved Quantization Indices

Li, Peng; Arildsen, Thomas; Larsen, Torben

Published in:

Proceedings of the 12th IEEE International Symposium on Signal Processing and Information Technology

DOI (link to publication from Publisher):

[10.1109/ISSPIT.2012.6621286](https://doi.org/10.1109/ISSPIT.2012.6621286)

Publication date:

2012

Document Version

Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Li, P., Arildsen, T., & Larsen, T. (2012). Extended Reconstruction Approaches for Saturation Measurements Using Reserved Quantization Indices. In *Proceedings of the 12th IEEE International Symposium on Signal Processing and Information Technology* (pp. 198-202). IEEE Press.
<https://doi.org/10.1109/ISSPIT.2012.6621286>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

2012 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 collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Published at IEEE ISSPIT 2012.

Extended reconstruction approaches for saturation measurements using reserved quantization indices

Peng Li, Thomas Arildsen and Torben Larsen
Department of Electronic Systems, Aalborg University
Niels Jernes Vej 12, 9000 Aalborg, Denmark
E-mail: {pl|tha|tl}@es.aau.dk.

Abstract—This paper proposes a reserved quantization indices method for saturated measurements in compressed sensing. The existing approaches tailored for saturation effect do not provide a way to identify saturated measurements, which is mandatory in practical implementations. We introduce a method using reserved quantization indices to mark saturated measurements, which is applicable to current quantizer models. Two extended approaches based on the proposed method have been investigated compared to the existing approaches. The investigation shows that saturated measurements can be identified by reserved quantization indices without adding extra hardware resources while maintaining a comparable reconstruction quality to the existing approaches.

Index Terms—Compressed sensing, Quantization index, Saturated measurement, Signal reconstruction

I. INTRODUCTION

Compressed sensing (CS) is a recent signal acquisition and reconstruction technique enabling sampling and compression simultaneously [1], [2]. The compressive sampling paradigm may enable perfect reconstruction with a sub-Nyquist sampling frequency [2]–[5], provided that original signals are known to be sparse (or compressible) in some basis. The compressed sensing technique may relax the requirements of high-speed analog-to-digital converter (ADC) for signals with high frequencies. In practical systems, quantization error (or quantization noise) introduced by ADCs influences the quality of the signal reconstruction, which cannot be neglected in compressed sensing [6]. Therefore, several studies have been done to investigate the effects of quantization in compressed sensing systems [6]–[14].

Some researchers propose modified signal reconstruction algorithms based on the information of quantization, which could improve the quality of signal reconstruction [6]–[11]. Others optimized the design of quantizers according to features of their reconstruction algorithms, which could also benefit the signal reconstruction [12], [13]. Further, since saturation is hard to completely avoid in quantization, unbounded quantization error of saturated samples could significantly impact the reconstruction performance in compressed sensing. Two existing approaches, the rejection approach and the consistent approach, tailored for saturated measurements in compressed sensing were proposed in [14]. The saturated measurements are rejected or be enforced consistency in the two existing approaches, respectively. However, no method was provided to identify the positions of the saturated measurements, which is mandatory in practical implementations, see Fig. 1a.

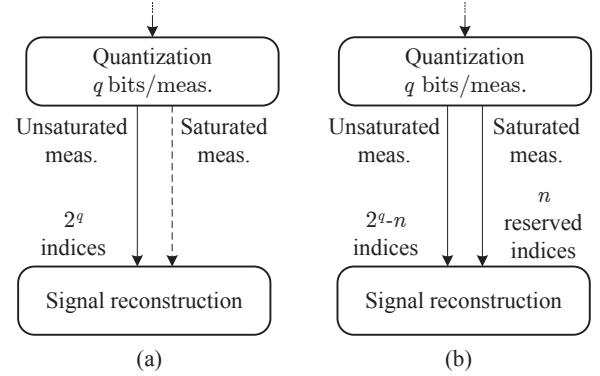


Fig. 1. (a) The existing approaches and (b) the extended approaches using reserved quantization indices.

In this paper, we propose a method reserving one or two of the set of possible quantization indices to mark saturated measurements, see Fig. 1b. It requires no extra hardware for the quantizers, and hence can be applied directly in the existing quantization models. We compare the reconstruction quality of the extended approaches based on the proposed method to the existing approaches. The simulation results show the feasibility of the proposed extended approaches.

II. METHODOLOGY

In this paper, a general CS-based structure is modeled with quantized compressed measurements. All input signals are presented as time-discrete vectors according to general CS theory [3].

A. Signal acquisition

To implement an existing CS technique in signal processing systems, it is necessary to find sparsity in the original signal [2], which means that the original signal should be sparse or compressible:

$$\mathbf{z} = \Psi \mathbf{x}, \quad (1)$$

where $\mathbf{z} \in \mathbb{R}^{N \times 1}$ is the original signal vector, $\Psi \in \mathbb{C}^{N \times N}$ is the dictionary and $\mathbf{x} \in \mathbb{C}^{N \times 1}$ is the sparse vector representation of \mathbf{z} in Ψ . In this model, \mathbf{x} is a vector containing only a few non-zero elements. \mathbf{x} is sparse if all other elements are zeros, thus $1 \leq \|\mathbf{x}\|_0 \ll N$, where $\|\mathbf{x}\|_0$ is the number of non-zero elements in \mathbf{x} , or \mathbf{x} is compressible if all other

elements are small enough to be accurately approximated by zeros.

According to general CS theory [2], a measurement matrix Φ is used to sample the original signal, which is described by:

$$\mathbf{y} = \Phi \mathbf{z} = \Phi \Psi \mathbf{x}, \quad (2)$$

where $\mathbf{y} \in \mathbb{C}^{M \times 1}$ is the compressed measurement vector, $\Phi \in \mathbb{C}^{M \times N}$ is the measurement matrix and M is the number of compressed measurements. If the product of the measurement matrix Φ and the dictionary matrix Ψ obey the Restricted Isometry Property (RIP) [2], the original signal can be recovered from M measurements, where $\|\mathbf{x}\|_0 < M \ll N$. In this paper, Φ is a random Gaussian matrix, which means all elements in Φ are randomly chosen from a collection with Gaussian distribution, to obey the requirement of RIP [2].

B. Quantization and Saturation

According to existing CS theory [4], [5], an original signal can be reconstructed from a compressed measurement vector \mathbf{y} . However, in practise, measurements need to be converted from analog to digital before further processing in a DSP. Since quantization is a necessary part of current ADCs [15], quantization error is an important factor influencing the quality of signal reconstruction. For a general view, uniform quantizers are used in this paper:

$$\mathbf{y}_Q = \mathbf{y} + \mathbf{e}, \quad (3)$$

where \mathbf{y} is the input measurement vector, \mathbf{y}_Q is the quantized measurement vector, and \mathbf{e} is an additive quantization error/noise vector. If we choose a one-dimensional scalar uniform quantizer with a resolution of q bits/meas., the entire quantization range is divided into W equal quantization partitions, where $W = 2^q$ [15]. In this work we consider scalar quantization of each of the elements of the vector \mathbf{y} .

Usually, the quantization range of a quantizer is chosen based on the type or class of input signals applied. If μ is the expectation of input data and a quantization range is bounded by $[\mu - G, \mu + G]$ ($G > 0$), G is called the saturation level and any input data exceeding this range is saturated. In this paper, input data is quantized to the mid-point values of the quantization partitions in uniform quantizers. Since the measurement matrices are random Gaussian matrices in this paper, the elements of the measurement vector \mathbf{y} are assumed to have a Gaussian distribution too. Therefore, saturation level and saturation rate are correlated [16] as:

$$G = \sigma \sqrt{2} \cdot \text{erf}^{-1}(1 - r), \quad (4)$$

where erf^{-1} denotes the inverse error function, σ is the standard deviation of \mathbf{y} and r ($0 \leq r \leq 1$) is the saturation rate, i.e., the ratio of the number of saturated measurements to the number of total measurements.

According to (4), quantizers with a small saturation level G lead to increased saturation rate r and vice versa. Then, if saturated measurements are allowed to occur in the models of the existing rejection and consistent approaches [14], the

average quantization error \mathbf{e} of non-saturated values will become smaller when using the same number of quantization indices W (due to smaller step-size) at the expense of more saturated values. The quality of the reconstructed signal may be better due to smaller quantization errors [14].

III. ALGORITHM

A. Reserved quantization indices method

To tailor reconstruction for saturation effects, both the existing rejection and consistent approach need to identify the indices of saturated measurements for further processing [14]. However, the two existing approaches do not provide a method to identify saturated measurements. It is not feasible for the reconstruction algorithms to know this information without any change in the quantizers. Therefore, the reserved quantization indices method is proposed to extend the existing approaches, see Fig. 1. It is achieved by reserving some of the available quantization indices to represent saturated measurement values and then allocating the remaining quantization indices to represent unsaturated measurement values. The number of reserved quantization indices is $n \in \{1, 2\}$ depending on the specific reconstruction algorithm. Based on the proposed method, we modify the two existing approaches in [14]. The extended approaches provide the feasible ways to tailor for the saturated measurements in CS in practical systems.

B. Rejection approach

1) *Existing formulation:* The existing rejection approach was proposed by Laska et al. in [14], which is used to accommodate the saturation effect in CS. We define \mathcal{S} as the set of indices of the unsaturated measurements. The vector of unsaturated quantized measurements $\tilde{\mathbf{y}}_Q$ of length \tilde{M} ($\tilde{M} \leq M$) and the measurement matrix $\tilde{\Phi}$ consisting of the rows corresponding to the unsaturated measurements are defined as:

$$\tilde{\mathbf{y}}_Q = \mathbf{y}_Q^{\mathcal{S}}, \quad \tilde{\Phi} = \Phi^{\mathcal{S}}. \quad (5)$$

The existing rejection approach is then defined as [14]:

$$\begin{aligned} \min. \quad & \|\tilde{\mathbf{x}}\|_1 \\ \text{s.t.} \quad & \|\tilde{\Phi} \Psi \tilde{\mathbf{x}} - \tilde{\mathbf{y}}_Q\|_2 < \epsilon, \end{aligned} \quad (6)$$

$$\epsilon = \sqrt{\tilde{M} + 2\sqrt{2\tilde{M}} \cdot \sigma_e}, \quad (7)$$

where the threshold ϵ is calculated according to [17] that requires estimation of the standard deviation σ_e of the quantization noise \mathbf{e} in (3). In the existing rejection approach, the saturated CS measurements and the corresponding rows in the measurement matrix are discarded. Then, the unsaturated measurements and the corresponding rows in the measurement matrix are used exactly as in the conventional CS reconstruction approach [18]. If the resolution of the quantizer is q bits/meas., the number of quantization indices for unsaturated measurement is $W = 2^q$ in this case that how to indicate which values are saturated is not considered in the existing rejection approach.

2) *Extended formulation*: The existing rejection algorithm needs to know the locations of saturated measurements in the measurement vectors to generate $\tilde{\mathbf{y}}_Q$ and $\tilde{\Phi}$, and this information is necessary for the signal reconstruction. The existing rejection algorithm in [14] does not take into account how to index saturated measurements in practice. Therefore, in the proposed extended rejection approach, one reserved quantization index ($n = 1$) is used to mark all saturated measurements. If the resolution of the quantizer is q bits/meas., and $\tilde{\mathbf{y}}_{QR}$ is the vector of the quantized unsaturated measurements, there are $W - 1$ ($W = 2^q$) quantization indices for quantizing each element in $\tilde{\mathbf{y}}_{QR}$ in this case. The entire quantization range is divided into $W - 1$ ($W = 2^q$) quantization partitions to quantize the unsaturated measurements. The extended rejection approach is defined as:

$$\tilde{\mathbf{y}}_{QR} = \mathbf{y}_{QR}^S, \quad (8)$$

$$\begin{aligned} \min. \quad & \|\tilde{\mathbf{x}}\|_1 \\ \text{s.t.} \quad & \|\tilde{\Phi}\Psi\tilde{\mathbf{x}} - \tilde{\mathbf{y}}_{QR}\|_2 < \epsilon. \end{aligned} \quad (9)$$

The extended rejection approach is a simple method tailored for identification of saturated measurements which merely requires a change to the quantizer and not the reconstruction method itself.

C. Consistent approach

1) *Existing formulation*: Since the information in saturated measurements is wasted in the rejection approach, an advanced approach, called the consistent approach, was proposed in [14]. The existing consistent approach makes use of saturated measurements while they are treated differently by enforcing consistency, which means the saturated measurements should be consistent when sampling the original signal and the reconstructed signal. Provided that \mathcal{S}^+ and \mathcal{S}^- represent the sets of indices of the positive saturated measurements and negative saturated measurements, respectively, the measurement matrix Φ is divided into two sub-matrices: an unsaturated matrix $\tilde{\Phi} \in \mathbb{C}^{\tilde{M} \times N}$; and a saturated matrix $\hat{\Phi} \in \mathbb{C}^{(M-\tilde{M}) \times N}$, such that:

$$\hat{\Phi} = \begin{bmatrix} \Phi^{\mathcal{S}^+} \\ -\Phi^{\mathcal{S}^-} \end{bmatrix}. \quad (10)$$

Then, the existing consistent approach is defined as [14]:

$$\begin{aligned} \min. \quad & \|\tilde{\mathbf{x}}\|_1 \\ \text{s.t.} \quad & \|\tilde{\Phi}\Psi\tilde{\mathbf{x}} - \tilde{\mathbf{y}}_Q\|_2 < \epsilon, \\ \text{and} \quad & \hat{\Phi}\Psi\tilde{\mathbf{x}} \geq G \cdot \mathbf{1}, \end{aligned} \quad (11)$$

where $\mathbf{1} \in \mathbb{C}^{(M-\tilde{M}) \times 1}$ is a vector of ones. The existing consistent approach still uses the unsaturated measurements like conventional l_1 norm algorithms [18]. Like the existing rejection approach, if the resolution of the quantizer is q bits/meas., the number of quantization indices for unsaturated measurement is $W = 2^q$ in this case that how to indicate which values are saturated is not considered in the existing consistent approach. The existing consistent approach generally provides a better signal reconstruction quality than

the existing rejection approach [14]. However, it should be noted that the total computational complexity also increases in the existing consistent approach [14].

2) *Extended formulation*: The information of indices of positive saturated measurements and negative saturated measurements is necessary in the above signal reconstruction method. However, the acquisition and recording of this information is not included in the existing consistent approach. Therefore, in the proposed extended consistent approach, two reserved quantization indices ($n = 2$) are used to represent positive and negative saturated measurements, respectively. If the quantizer resolution is q bits/meas. for the entire system, and $\tilde{\mathbf{y}}_{QC}$ is the quantized unsaturated measurements vector, the number of quantization indices for $\tilde{\mathbf{y}}_{QC}$ is then $W - 2$ ($W = 2^q$). The extended consistent approach is defined as:

$$\tilde{\mathbf{y}}_{QC} = \mathbf{y}_{QC}^S, \quad (12)$$

$$\begin{aligned} \min. \quad & \|\tilde{\mathbf{x}}\|_1 \\ \text{s.t.} \quad & \|\tilde{\Phi}\Psi\tilde{\mathbf{x}} - \tilde{\mathbf{y}}_{QC}\|_2 < \epsilon, \\ \text{and} \quad & \hat{\Phi}\Psi\tilde{\mathbf{x}} \geq G \cdot \mathbf{1}. \end{aligned} \quad (13)$$

Reserving two quantization indices to mark positive and negative saturated measurements in the proposed approach, respectively, the saturated measurements can be identified for use in the reconstruction algorithm. When the quantizer has a high resolution q , lacking two quantization indices for unsaturated values may represents only a minor decrease in the available number of quantization levels and should cause only minor increase of quantization error.

IV. SIMULATION AND EMPIRICAL RESULTS

In this section, numerical results are presented to evaluate two proposed extended approaches. We compare the extended approaches in (III-B2, III-C2) to the existing approaches in (III-B1, III-C1) by the reconstruction quality defined in terms of Normalized Mean Square Error (NMSE) ρ :

$$\rho = \frac{\|\tilde{\mathbf{z}} - \mathbf{z}\|_2^2}{\|\mathbf{z}\|_2^2} = \frac{\|\Psi\tilde{\mathbf{x}} - \Psi\mathbf{x}\|_2^2}{\|\Psi\mathbf{x}\|_2^2}, \quad (14)$$

where \mathbf{z} is the original signal, \mathbf{x} is its sparse representation, $\tilde{\mathbf{z}}$ is the reconstructed signal and $\tilde{\mathbf{x}}$ is the reconstructed sparse representation. If ρ_1 and ρ_2 are two values of NMSE based on the same specification in the existing and extended approaches (rejection or consistent), respectively, the difference of reconstruction quality for the two approaches is evaluated by:

$$\left| \frac{\rho_2 - \rho_1}{\rho_1} \right| \times 100\%. \quad (15)$$

For a general view, we test above four approaches in the scenarios of different saturation rates r , quantizer resolutions q and numbers of measurements M .

In all simulations, multi-tone signals are used as original signals, which are sparse in the frequency domain. We stress the fact that the multi-tone signal is just an example of all possible kinds of original signals for the approaches in our simulations, which helps us to focus the analysis on the effects

of the proposed methods. The tones are randomly located in a range of $[0, 500]$ Hz in the frequency domain. The minimum guard space is 10Hz between each tone. According to (1) and (2), we apply the four approaches to multi-tone signals randomly generated according to the following specifications: size of sparse vector $N = 1000$; number of tones $K = 10$; number of compressed measurements $M \in \{80, 160\}$. The expectation of compressed measurements, μ , is zero in all simulations. The dictionary Ψ is an inverse discrete Fourier transform (DFT) matrix in our simulations according to (1). The measurement matrix Φ has i.i.d. zero-mean Gaussian entries $\sim \mathcal{N}(0, 1/M)$. We test uniform quantizers with resolution $q \in \{2, 4, 6, 8\}$ bits/meas. In each simulation, we use 1000 randomly generated multi-tone signals as training signals to estimate the saturation level G for the uniform quantizer according to different saturation rate $r \in \{0.5\%, 1\%, 2\%, 4\%, 8\%, 16\%\}$, see (4). In this paper, each simulation repeats 1000 times with randomly generated original signal and measurement matrix in each iteration and all numerical results shown in following figures use the average values. The open source optimizers SPGL1 [19], [20] and CVX [21], [22] are used in our simulations.

Fig. 2 shows the reconstruction quality, in terms of NMSE, versus saturation rate using the existing and extended rejection approaches. Four different quantizer resolutions are tested for both approaches. The simulation results show that the extended rejection approach has a comparable reconstruction performance to the existing rejection approach for high quantizer resolutions. In Fig. 2a, the maximum difference of reconstruction quality between the two approaches is only approximately 6% for the case of $q = 4, 6$ or 8 bits/meas. Relatively big difference of reconstruction quality between the two approaches is observed for the case of $q = 2$ bits/meas. This is due to that the quantization error in low quantizer resolutions becomes relatively larger when there is one quantization partition less in the extended rejection approach. In Fig. 2b, the curves for the two approaches hold a consistent similarity. The maximum difference is approximately 12% when $q = 4$ or 6 bits/meas. When $q = 8$ bits/meas., the extended rejection approach even provides a better reconstruction quality than the existing rejection approach. This is an added benefit for the extended approach. It should be noticed that quantization error is only one of the factors which influence the signal reconstruction quality [3], [6]. Due to the non-linearity of reconstruction algorithms in compressed sensing [2], it is possible that the proposed extended approaches exceed the existing approaches in some cases.

Fig. 3 shows the reconstruction quality, in terms of NMSE, versus saturation rate using the existing and extended consistent approaches. The number of quantization indices for unsaturated measurements is two fewer in the extended consistent approach. However, like results in Fig. 2, the extended consistent approach with high quantization resolutions provides a comparable reconstruction performance to the existing consistent approach. In Fig. 3a, the maximum difference of reconstruction quality between the two consistent approaches

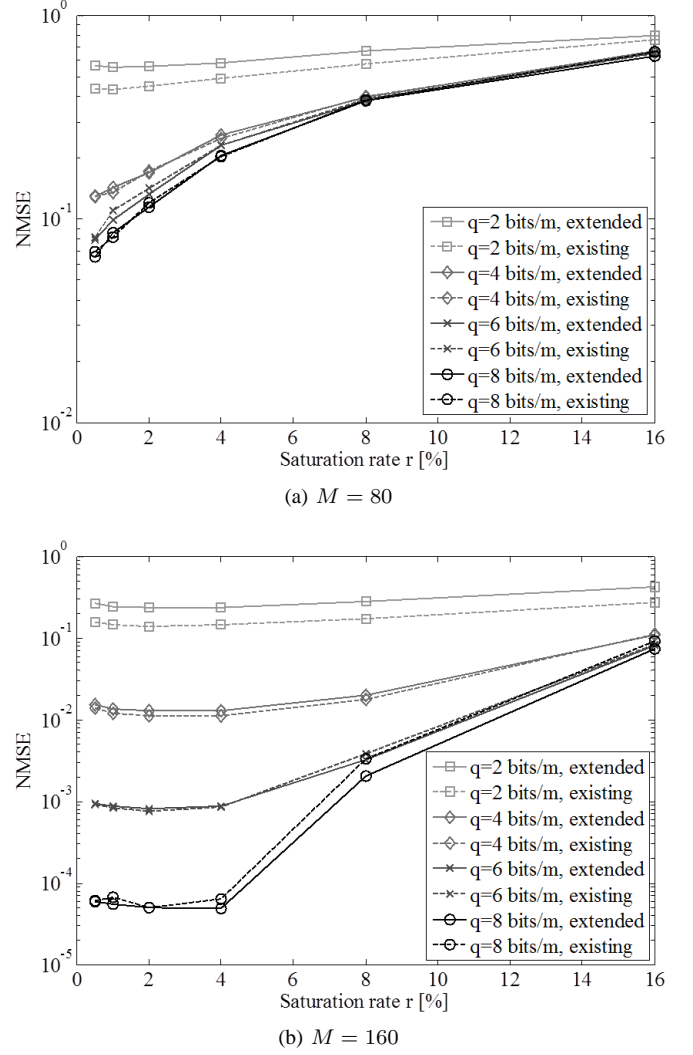


Fig. 2. Simulated NMSE versus saturation rate r based on both the existing and extended rejection approaches, q is quantizer resolution, M is number of measurements.

is below 11% when $q = 6$ or 8 bits/meas. In Fig. 3b, the maximum difference is below 6% when $q = 6$ or 8 bits/meas.

It is noticed that both proposed extended approaches provide more accurate signal reconstruction with larger numbers of compressed measurements M or higher quantizer resolutions q . Choosing the optimum saturation rate may significantly improve the signal reconstruction quality for both extended approaches. The extended consistent approach is more robust to larger saturation rates than the extended rejection approach. The main decline of reconstruction quality in the proposed extended approaches occurs when using low resolution quantizers. However, since the two existing approaches cannot provide accurate signal reconstruction with low resolution quantizers, e.g., $q = 2$ bits/meas., it is not a significant drawback for the extended approaches. In practical implementations, low resolution quantizers are not recommended to be applied to the extended rejection or consistent approach.

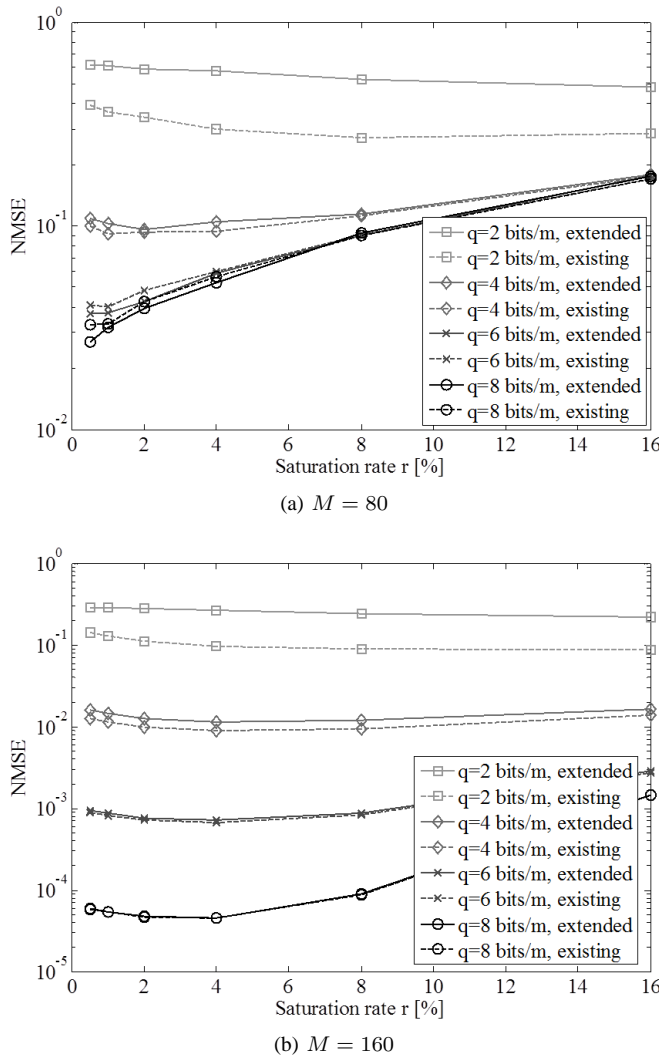


Fig. 3. Simulated NMSE versus saturation rate r based on both the existing and extended consistent approaches, q is quantizer resolution, M is number of measurements.

V. CONCLUSIONS

A reserved quantization indices method is proposed for saturation effects in compressed sensing. It is capable of identifying saturated measurements, which has not been included in the existing approaches. This is achieved by allocating a few of the available quantization indices to saturated measurements. Two extended approaches based on the reserved quantization indices method can be directly applied in the existing quantizer models. The investigation has been done by comparing the reconstruction quality of the extended approaches to the existing approaches. The simulation results indicate that the extended approaches are feasible and more practicable to realize with existing ADCs. The reserved quantization indices method holds the potential of practical quantizer implementations in compressed sensing.

VI. ACKNOWLEDGEMENT

The work is financed by The Danish Council for Strategic Research under grant number 09-067056.

REFERENCES

- [1] B. Hayes, "The best bits," *American Scientist*, vol. 97, pp. 276–280, 2009.
- [2] E. Candès and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, March 2008.
- [3] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [4] E. Candès and T. Tao, "Near-optimal signal recovery from random sensing: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.
- [5] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [6] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Processing Letters*, vol. 17, no. 2, pp. 149–152, Feb. 2010.
- [7] L. Jacques, D. Hammond, and J. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 559–571, Jan. 2011.
- [8] W. Dai, H. V. Pham, and O. Milenkovic, "Distortion-rate functions for quantized compressive sensing," in *IEEE Information Theory Workshop on Networking and Information Theory*, June 2009, pp. 171–175.
- [9] Wei Dai, Hoa Vinh Pham, and Milenkovic, O., "A comparative study of quantized compressive sensing schemes," in *IEEE International Symposium on Information Theory*, July 2009, pp. 11–15.
- [10] C. Güntürk, M. Lammers, A. Powell, R. Saab, and O. Yilmaz, "Sigma delta quantization for compressed sensing," in *44th Annual Conference on Information Sciences and Systems (CISS)*, March 2010, pp. 1–6.
- [11] P. T. Boufounos, "Reconstruction of sparse signals from distorted randomized measurements," in *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, March 2010, pp. 3998–4001.
- [12] J. Sun and V. Goyal, "Optimal quantization of random measurements in compressed sensing," in *IEEE International Symposium on Information Theory*, July 2009, pp. 6–10.
- [13] P. Boufounos, "Universal rate-efficient scalar quantization," *IEEE Transactions on Information Theory*, vol. 58, no. 3, pp. 1861–1872, March 2012.
- [14] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing," *Applied and Computational Harmonic Analysis*, vol. 31, no. 3, pp. 429–443, Nov. 2011.
- [15] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, ser. The Springer International Series in Engineering and Computer Science. Springer, 1991.
- [16] R. J. Bagby, "Calculating normal probabilities," *American Math Monthly* 102, p. 46C49, 1995.
- [17] S. Becker, J. Bobin, and E. J. Candès, "Nesta: A fast and accurate first-order method for sparse recovery," *SIAM Journal on Imaging Sciences*, vol. 4, no. 1, pp. 1–39, 2011.
- [18] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *Society for Industrial and Applied Mathematics*, vol. 43, no. 1, pp. 129–159, 2001.
- [19] E. van den Berg and M. P. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008. [Online]. Available: <http://link.aip.org/link/?SCE/31/890>
- [20] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction."
- [21] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21."
- [22] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.