19th Symposium IMEKO TC 4 Symposium and 17th IWADC Workshop
Advances in Instrumentation and Sensors Interoperability
July 18-19, 2013, Barcelona, Spain

On compressed sensing and super-resolution in DFT-based spectral analysis

M. Bertocco, G. Frigo, C. Narduzzi

Università di Padova, Department of Information Engineering, via G. Gradenigo, 6/b, I-35131 Padova, Italy
ph.: +39 049 8277500, fax: +39 049 8277699, e-mail: {mat, frigogug, narduzzi}@dei.unipd.it

Abstract-The paper discusses a novel frequency interpolation and super-resolution method for multitone waveform analysis, where a compressive sensing algorithm is employed to process data. Each signal acquisition involves a short data record, whose DFT coefficients are computed. A set of compressed measurements is obtained by taking records with different known starting instants, and employed to determine, by solving an orthogonal matching pursuit problem, the set of frequency components of the analysed waveform. Interpolation is presented as a compressed sensing problem and algorithm performances discussed.

Keywords-Fourier analysis, super-resolution, compressed sensing

I. Introduction

A classical problem for spectral analysis is the accurate measurement of a multitone waveform by the estimation of the individual sinusoidal component parameters. With algorithms based on discrete Fourier transform (DFT), spectral leakage is known to affect the accuracy of amplitude estimation as well as the ability to resolve closely spaced frequency components [1]. As long as the frequency separation of waveform components is large enough, several forms of frequency interpolation of DFT coefficients allow to significantly improve estimation accuracy, approaching the theoretical limit set by the Cramér-Rao bound [2], [3], [4]. However, if sinusoid frequencies are close to each other, interpolation is likely to fail and different approaches are required, among which are a variety of super-resolution methods [5].

This paper analyses the results obtained by the use of a compressive sensing (CS) approach. This data acquisition and signal processing paradigm exploits the fact that several signals of interest are *sparse*, that is, they can be described by a comparatively small number of parameters in a suitable domain [6], [7]. The problem outlined above clearly has a sparse nature, as the large number of non-zero DFT coefficients is a consequence of spectral leakage, whereas the original sinusoidal components are a much smaller number. The preliminary issue is to describe interpolation and super-resolution using the formal framework of CS problems.

It should be observed that, although CS approaches have already been proposed in the literature for image super-resolution [8], [9], application of a CS algorithm for frequency interpolation and super-resolution in multitone waveform analysis has not, to the authors' knowledge, been discussed before.

II. Frequency interpolation as a Compressive Sensing problem

The measurement model underlying CS problems is given by a simple linear sensing equation of the form:

$$y = \boldsymbol{A}\, x + n \tag{1}$$

where $\boldsymbol{y}$ is the column vector of measurements and $\boldsymbol{x}$ is the vector of unknown quantities that generate the observations. The basic assumption is that $\boldsymbol{x}$ is a sparse vector, that is, only a few of its elements are non-zero. The vector $\boldsymbol{n}$ accounts for noise and uncertainty that, in general, will affect measurements. The *a priori* assumption that $\boldsymbol{x}$ is sparse provides the main constraint to the solution of (1). It is assumed that vector $\boldsymbol{y}$ has size N, while the dimension H of the vector $\boldsymbol{x}$ represents the size of a candidate index set, only a few elements of which are associated to non-zero vector components. This sparsity assumption prevents the problem from becoming underdetermined.

Of course, because of the presence of noise, no exact solution is possible and the best estimate is the sparsest vector that satisfies the condition:

$$\| y - \boldsymbol{A} \cdot x \|_2^2 \;\leq\; \epsilon \;. \tag{2}$$

A multitone waveform is the sum of a number of complex exponential components (with suitable symmetries if the waveform is assumed to be real). Indicating by S_H the set of values of the index h corresponding to a waveform component, the signal can be written in sampled form as:

$$x[n]=\sum_{h\in S_H} A_h e^{j\left(\phi_h+2\pi\lambda_h n\right)},\ -\infty<n<+\infty \quad , \tag{3}$$

where component frequencies have been expressed in the normalised form $\lambda_h = f_h T$, and T is the sampling interval. The starting point for the problem at hand is the equation that defines the DFT of a multitone waveform, as calculated from N time-domain samples acquired with sampling interval T:

$$X\left(\frac{k}{N}\right)=\sum_{h\in S_H} A_h e^{j\phi_h}\left[\frac{\sin\pi\left(k-\lambda_h N\right)}{N\sin\frac{\pi}{N}\left(k-\lambda_h N\right)}e^{-j\pi\frac{N-1}{N}\left(k-\lambda_h N\right)}\right]e^{-j\frac{2\pi}{N}\left(k-\lambda_h N\right)n_0} \quad , \tag{4}$$

where n_0 is the starting index of the time-domain sequence.

The normalised frequency granularity of a N-point DFT is $1/N$ and the target of this work is to estimate frequency components whose separation may be close to, or even lower than the granularity. Accordingly, a denser frequency grid is defined, whose minimum step is $\Delta\lambda = 1/H$, with $H >> N$ (one can assume $H = NP$, where P is the interpolation factor), which allows to express normalised component frequencies as integer multiples of the step $\Delta\lambda$, yielding: $\lambda_h = h/H$ with $h \in S_H$ and $S_H \subset \{0, \ldots, H\text{-}1\}$. Of course, the calculation of a H-point DFT could provide the desired granularity, however this would be achieved at the cost of a longer observation interval spent acquiring the samples.

The idea discussed in this paper is to process a small set of M data records, each having length N and with different ***known*** starting indexes n_m defined with respect to a common time reference, or trigger point. These are arranged into a "compressive" measurement matrix $\boldsymbol{Y}$ composed of M complex column vectors, each containing the DFT coefficients of a single record. Varying n_m allows to exploit the corresponding changes induced on the phases of individual DFT coefficients in different records.

Using (4), if no additive noise is present each element of the measurement matrix $\boldsymbol{Y}$, indicated as $y_{k,m}$, can be written as:

$$y_{k,m} = \sum_{h\in S_H} A_h e^{j\left(\phi_h+2\pi\frac{h}{H}n_m\right)}\left[\frac{\sin\pi\left(k-\frac{h}{H}N\right)}{N\sin\frac{\pi}{N}\left(k-\frac{h}{H}N\right)}e^{-j\pi\frac{N-1}{N}\left(k-\frac{h}{H}N\right)}\right]. \tag{5}$$

It should be noticed that the complex exponential term $e^{-j\frac{2\pi}{N}\left(k-\frac{h}{H}N\right)n_m}$, that appears in the general expression (4), has been split in two terms: $e^{j2\pi\frac{h}{H}n_m}$ has been added as a constant known phase term to the corresponding signal component, while $e^{-j\frac{2\pi}{N}kn_m}$ is simply set to zero by always using $n = 0$ as the starting index for DFT computation, regardless of the actual start time. This equation shows that the frequency interpolation problem can be seen as the problem of solving the matrix equation:

$$\boldsymbol{Y}=\boldsymbol{D}\cdot\boldsymbol{X}+\boldsymbol{N} \tag{6}$$

where the columns of matrix $\boldsymbol{D}$ contain the complex samples of the Dirichlet kernel (within square brackets in (5)), each progressively shifted by fractional frequency amounts $\Delta\lambda$. Thus, the row index is k while the column index is h. The matrix $\boldsymbol{N}$ accounts for measurement noise and uncertainty.

In Fig.1 the peculiar structure of matrix $\boldsymbol{D}$ is presented: a three-dimensional representation is combined with a row and column trend.

ISBN-10: 84-616-5438-2 | ISBN-13: 978-84-616-5438-3

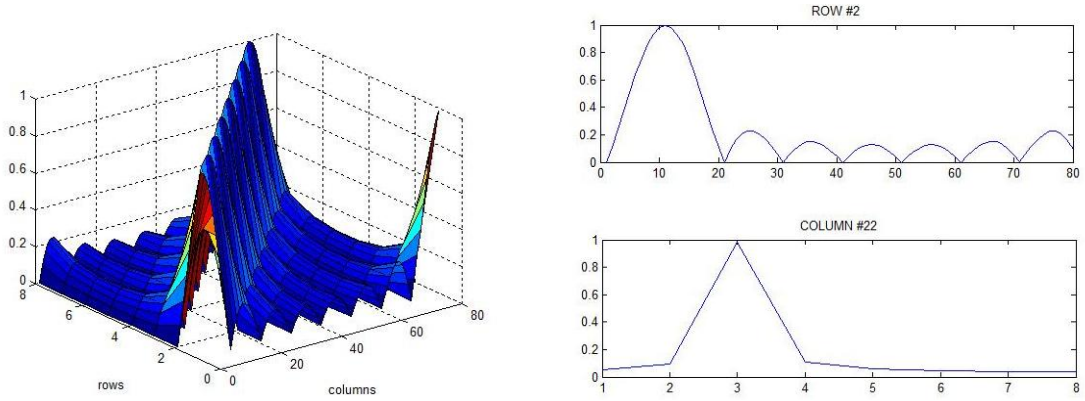

Fig. 1: a) three-dimensional representation of matrix $\boldsymbol{D}$ with $N = 8$ and $P = 10$; b) typical row (top) and column (down) trends.

The amplitude and phase of sinusoidal components forming the multitone waveform are represented by the non-zero complex elements of a matrix $\boldsymbol{X}$, composed of M vectors $\boldsymbol{x_m}$ whose size H agrees with the desired, denser, frequency grid. It should be noticed that each vector $\boldsymbol{x_m}$ is associated to a specific value of n_m, since from (5) one has:

$$x_{h,m} = A_h e^{j\left(\phi_h + 2\pi \frac{h}{H} n_m\right)} . \tag{7}$$

Thus, the matrix $\boldsymbol{X}$ can be factored as:

$$X = \begin{bmatrix} A_1 e^{j\phi_1} & \dots & 0 \\ 0 & \ddots & 0 \\ 0 & \dots & A_{H-1} e^{j\phi_{H-1}} \end{bmatrix} \begin{bmatrix} e^{j2\pi \frac{0}{H} n_0} & \dots & e^{j2\pi \frac{0}{H} n_{M-1}} \\ \vdots & \ddots & \vdots \\ e^{j2\pi \frac{H-1}{H} n_0} & \dots & e^{j2\pi \frac{H-1}{H} n_{M-1}} \end{bmatrix} \tag{8}$$

The fundamental assumption that allows to formulate frequency interpolation as a CS problem and achieve super-resolution is that, whatever the individual value of n_m, the frequencies of the non-zero components are the same for all measurements. Hence, vectors $\boldsymbol{x_m}$ are sparse and, additionally, they are known to share the ***same*** set S_H of frequency indexes [10], corresponding to the non-zero waveform components. In CS terms, this is called a multiple measurement vectors problem.

III. Compressed Sensing based Fourier analysis

From a measurement viewpoint, the proposed CS-based data acquisition allows to reduce the length of the acquisition interval (and of the total sample record), since the sets of samples from which each DFT vector in $\boldsymbol{Y}$ is computed can be broadly overlapped. If the average difference between consecutive starting indexes is Δn, the total number of samples required is $N+\Delta n(M-1)$, with $\Delta n << N$. Assuming the CS algorithm allows an interpolation factor P, this should be compared with the length $H = NP$ required to achieve the same frequency granularity by a standard DFT algorithm.

The first step in solving the compressed sensing problem is *support recovery*, which has the purpose of determining which elements in $\boldsymbol{x_m}$ are non-zero. In practice, the aim is to determine the set $S_H \subset \{0, \dots, H-1\}$, as defined above. As far as Fourier analysis is concerned, this is in fact the key part of the algorithm, since S_H contains the frequency locations of the signal components.

Following the approach proposed in [10], this step is carried out via a singular-value decomposition of the matrix $\boldsymbol{YY}^H$, where the superscript denotes transposition and complex conjugation. Decomposition allows to obtain the matrix $\boldsymbol{V}$, whose columns are the eigenvalues of $\boldsymbol{YY}^H$ multiplied by the square roots of the corresponding eigenvectors. Collapsing the columns of $\boldsymbol{V}$ into a single vector $\boldsymbol{v}$, and solving the equation $\boldsymbol{v}=\boldsymbol{Du}$ by an orthogonal matching pursuit algorithm allows to find the support for $\boldsymbol{u}$, which can be shown to be the same as S_H. By construction, $\boldsymbol{V}$ and $\boldsymbol{D}$ are known: the first is obtained from the measurements result, the second defines the measuring scheme.

To test the approach, a signal formed by three known complex exponential components, according to (3), was

ISBN-10: 84-616-5438-2 | ISBN-13: 978-84-616-5438-3

considered and different sets of M records of length $N = 256$ samples were generated. In the multitone waveform given in Fig. 2 spectral components have normalised frequencies $\lambda_1 = 72.6$, $\lambda_2 = 74$, $\lambda_3 = 103.4$, so that the second component contributes no spectral leakage, but is affected by leakage from the first one. In Fig. 2.a sampling rate is set to 500 Hz, while component frequencies in Fig. 2.b correspond to 141.80, 145.70 and 201.95 Hz. The third component is employed as a control element, being sufficiently far away from the first two to be only affected by scalloping loss.

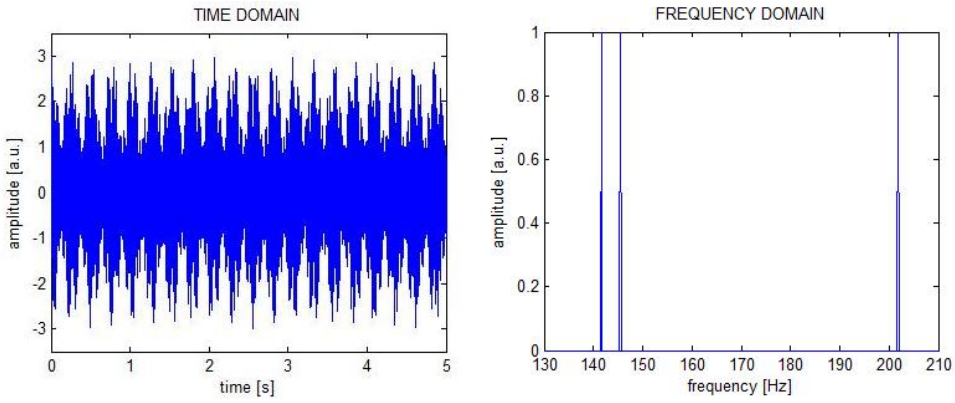

(a) (b)

Fig. 2: a) signal real part in the time domain; b) detail of signal Fourier transform amplitude.

In the test configuration $M = 64$ and the starting points of each sample record are separated by 3 samples. Component amplitudes were set to 1, their initial phases to 0. In the following, amplitude values are expressed in arbitrary units, while phase values are expressed in radians. Simulations were iterated 100 times for each test condition. Trials were first carried out with no noise. The desired interpolation factor is $P = 10$, ensuring that frequencies fall exactly on points of the denser grid (i.e., they are integer multiples of $\Delta\lambda$). With these values the support recovery step always provided the correct component frequencies.

The subsequent algorithm step involves the construction of a restricted matrix $\boldsymbol{D_S}$ by deleting columns whose index h is not in S_H. This turns (6) into an overdetermined matrix equation, that can be solved by computing the pseudo-inverse:

$$\hat{\boldsymbol{X}}_S = \left(\boldsymbol{D_S^H}\,\boldsymbol{D_S}\right)^{-1}\boldsymbol{D_S^H}\,\boldsymbol{Y} \tag{9}$$

which only contains the non-zero rows of $\boldsymbol{X}$. The final step is to invert the relevant rows according to Eq. (8); this accurately reconstructs component amplitudes and phases (to within machine error), provided the number of measurement vectors M is greater than the number of signal components.

When measurement noise is present, replacing Eq. (6) into (9) shows the estimate to be:

$$\hat{\boldsymbol{X}}_S = \boldsymbol{X} + \left(\boldsymbol{D_S^H}\,\boldsymbol{D_S}\right)^{-1}\boldsymbol{D_S^H}\,\boldsymbol{N} \tag{10}$$

provided the support recovery step can still be successfully completed. However, noise variability affects the support recovery algorithm, whose results are no longer deterministic, but susceptible to changes: each iteration could provide a different support, with obvious consequences also on amplitude and initial phase estimates. This phenomenon becomes more and more significant as the signal-to-noise ratio decreases. Then, it is reasonable to assume there exists a limiting value of the signal-to-noise ratio (SNR), below which the support is wrongly identified and accurate component estimation becomes impossible.

To find this limiting value, tests were carried out by repeating the CS algorithm 100 times for a number of different SNR values, recording the number of successful support recoveries. In this case, to avoid any interference contribution the normalised frequencies were set to $\lambda_1 = 72.6$, $\lambda_2 = 98$, $\lambda_3 = 103.4$, so that the second component, which still causes no leakage, is also far enough to be unaffected by leakage from any of the other two components.

Results are shown in Fig. 3, which evidences that successful support recovery cannot be guaranteed below 20 dB SNR. It should be noted that only an exact coincidence means a successful result.

ISBN-10: 84-616-5438-2 | ISBN-13: 978-84-616-5438-3

SNR	Success probability
20	100
15	98
10	92
5	70
0	28
-5	8
-10	0

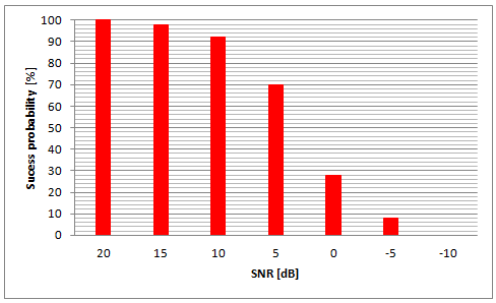

Fig. 3: success probability with progressively reduced SNR.

IV. Super-resolution experimental results

So far, it has been shown that the proposed CS algorithm can accurately estimate the frequency position of signal components on a dense (interpolated) grid, thus allowing to compensate for the effects of spectral leakage. It remains to show that components that are placed very close in frequency can be picked out and correctly estimated. Then, achievable performances should be characterized.

Given the initial test configuration described above, in the absence of noise the second component was shifted in frequency, by steps of 0.1 bins, towards the first component. Results were found to depend on signal parameter settings, in particular on the initial phases. If the phase shift between the first and second components is equal to 0, $\pi/2$ or π the resolution threshold at which support recovery is successful is equal to 1.1 bins. By raising the threshold to just 1.2 bins, the two equal amplitude components are correctly resolved with any phase configuration. For any random set of phase values, reconstruction error is negligible.

Spectral interference becomes worse if one component is larger than the other. Keeping the first component amplitude constant, we progressively reduced the second one, up to a ratio of 1:1000. In this case, a minimum separation of 1.4 bins is required to ensure a correct resolution for any possible setting. This condition was verified by simultaneously varying the amplitudes and phases of both components with a random pattern.

It is also important to find out which is the minimum required number of measurements. While correct support recovery is possible even for $M = 1$, at least $M = 4$ is necessary to provide good accuracy in waveform reconstruction, as shown in Fig. 4.

Number of measurements	Amplitude mean error component #1	Amplitude mean error component #2	Amplitude mean error component #3
64	1,11E-016	4,44E-016	8,88E-016
32	6,66E-016	4,44E-016	4,44E-016
16	1,11E-015	4,44E-016	2,22E-016
8	8,88E-016	0,00E+000	1,22E-015
4	4,44E-016	2,33E-015	3,33E-016
2	4,84E-001	5,13E-001	3,13E-003
1	6,67E-001	6,67E-001	6,67E-001

(a)

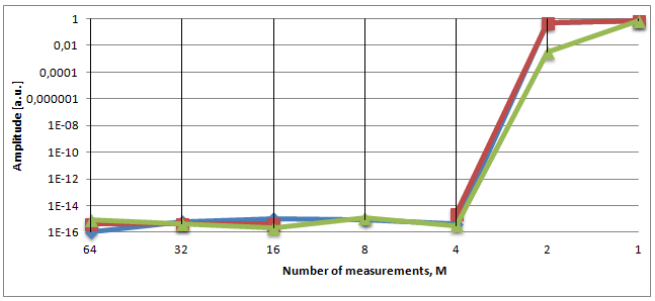

(b)

Fig. 4: a) amplitude means errors according to different number of measurements; b) logarithmic plot of mean error trends.

On the contrary, the relative spacing between the initial points n_m does not affect significantly the results, even though by reducing the superposition of two sequences, their statistical correlation is less significant. On the other hand, by a proper choice of the number and position of the measurement initial points the total observation interval can be markedly reduced. A good compromise was found to be $M = 4$ and $n_m = [0\ 3\ 6\ 9]$. In this way, only 265 samples are acquired to obtain an estimate over a grid of 2560 frequency points.

In the presence of noise, the method's effectiveness depends critically on support recovery. Adopting a SNR equal to 20 dB, method's performances are comparable with the noiseless case. A distance of 1.4 bins represents a resolution threshold, once more not dependent on component amplitudes and initial phases. For components with different amplitudes, the minimum ratio has to be limited to 1:10 (i.e., -20 dB), otherwise the second component is hidden by noise fluctuations.

To understand the impact of support recovery in comparison to noise effects on component parameter estimation with known support, two sets of simulations were carried out, where SNR varied but support was

known a priori. Results are presented in Fig. 5.a for $A_1 = A_2 = 1$ a.u. and in Fig. 5.b for $A_1 = 1$ and $A_2 = 0.1$ a.u.. In both cases, the mean deviation of the amplitude estimation error for A_2 is almost one order of magnitude below the standard deviation. It is important to note that, when the support is known, estimation variance is proportional to SNR, as implied by Eq. (10).

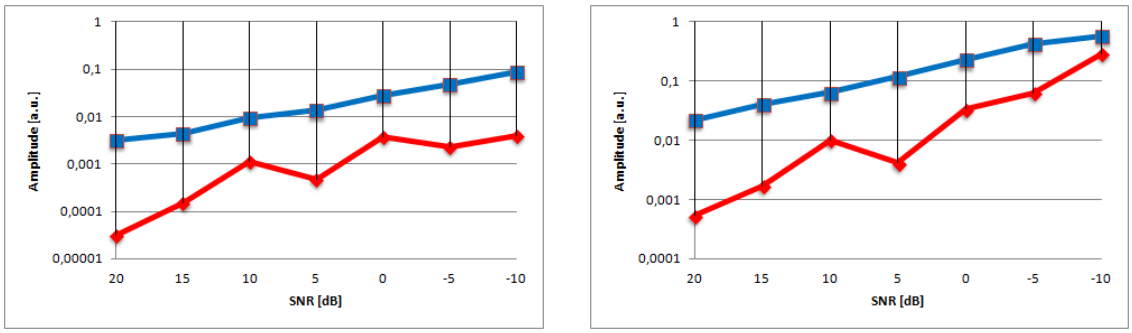

(a) (b)

Fig. 5: mean (diamonds) and standard deviation (squares) of amplitude estimation error vs. SNR on second signal component: a) $A_1 = A_2 = 1$ a.u.; b) $A_1 = 1$, $A_2 = 0.1$ a.u..

In conclusion, obtained results suggest that a resolution threshold equal to 1.4 bins can be considered valid until SNR value is not lower than 20 dB. Within these constraints, the proposed method recovers the original support and provides accurate estimates of component amplitudes and initial phases.

IV. Conclusions

The use of CS algorithms has become extensive in recent years, but application details and a number of subtle tradeoffs are not always clear, until a specific implementation is tackled. Frequency interpolation and super-resolution is a classic problem in measurement: the investigation presented in the paper shows how this can be approached within the CS framework and provides some characterisation of the algorithm accuracy. The latter is an essential side of the analysis, but in this case it is also a rather complex problem, involving an extensive study of various algorithmic aspects. Results obtained so far support the idea that the proposed approach can be considered an interesting alternative for high-resolution multitone waveform analysis and can attain super-resolution performance, provided the computing power necessary for this class of algorithms is available.

References

[1] D. C. Rife and R. R. Boorstyn, “Multiple tone parameter estimation from discrete-time observations”, *Bell Syst. Tech. J.*, vol. 55, pp. 1389-1410, 1976.
[2] V. H. Jain, W. L. Collins and D. C. Davis, “High accuracy analog measurements via interpolated FFT”, *IEEE Trans. Instrum. Meas.*, vol. 28, pp. 113-122, 1979.
[3] C. Offelli and D. Petri, “The influence of windowing on the accuracy of multifrequency signal parameter estimation”, *IEEE Trans. Instrum. Meas.*, vol. 41, pp. 256-261, 1992.
[4] D. Agrez, “Weighted multipoint DFT to improve amplitude estimation of multifrequency signal”, *IEEE Trans. Instrum. Meas.*, vol. 51, pp. 287-292, 2002.
[5] S. M. Kay, *Modern Spectral Estimation*, Prentice-Hall, 1987.
[6] E. J. Candès and M. B. Wakin, ``An introduction to Compressive Sampling'', *IEEE Signal Process. Mag.*, vol. 25, pp.~21-30, 2008.
[7] D. L. Donoho, “Compressed sensing”, *IEEE Trans. Information Theory*, vol. 52, pp. 1289-1306, 2006
[8] A. Fannjiang, W. Lao, “Super-resolution by compressive sensing algorithms”, Proc. IEEE 46th Ann. Asilomar Conf., Pacific Grove, CA, USA, 4-7 Nov. 2012.
[9] E. J. Candès and C. Fernandez-Granda, “Towards a mathematical theory of super-resolution”, arXiv:1203.5781v3 [cs.IT].
[10] M. Mishali, Y. C. Eldar, “Reduce and Boost: Recovering Arbitrary Sets of Jointly Sparse Vectors”, *IEEE Trans. Signal Process.*, vol. 56, pp. 4692-4702, 2008.