OPTIMAL SUBSAMPLING OF MULTICHANNEL DAMPED SINUSOIDS

Gilles Chardon and Laurent Daudet

Institut Langevin
UPMC - Univ. Paris 6, Université Paris Diderot - Paris 7, ESPCI ParisTech, CNRS UMR 7587,
10, rue Vauquelin, 75231 Paris Cedex 05, France
firstname.lastname@espci.fr

ABSTRACT

In this paper, we investigate the optimal ways to sample multichannel impulse responses, composed of a small number of exponentially damped sinusoids, under the constraint that the total number of samples is fixed - for instance with limited storage / computational power. We compute Cramér-Rao bounds for multichannel estimation of the parameters of a damped sinusoid. These bounds provide the length during which the signals should be measured to get the best results, roughly at 2 times the typical decay time of the sinusoid. Due to bandwidth constraints, the signals are best sampled irregularly, and variants of Matching Pursuit and MUSIC adapted to the irregular sampling and multichannel data are compared to the Cramér-Rao bounds. In practical situation, this method leads to savings in terms of memory, data throughput and computational complexity.

***Index Terms*—** compressed sensing, spectral analysis, array signal processing

1. INTRODUCTION

Array processing techniques have proved extremely powerful for many tasks in signal processing. To name but a few of these techniques, we can mention the remarkable achievements of beamforming [1], holography [2], Synthetic Aperture Radar [3] etc. However, there are a many practical situations, typically when a high bandwidth is needed, where such techniques lead to an extremely large number of samples, which in turn leads to stringent experimental constraints on the acquisition devices that need to cope with high data throughput and large storage capabilities. Processing these signals can also involve very large computational costs.

In most of these techniques, such as in the Nearfield Acoustic Holography (NAH) [2] that motivates this study, the acquired signals have a very specific structure, as the different sensors typically acquire different "views" of the same source, therefore sharing a lot of characteristics. For instance, in NAH used to visualize resonating modes of vibrating structures with an array of microphones, every signal is a sum of a few sinusoids (one per mode, in the simplest non-degenerated case), whose frequencies (and damping factor if relevant) are shared across all microphones. Here, individual microphone signals only differ by the amplitude and phase of each sinusoid, and that is precisely the information that is needed to image the vibration of the structure. It should be noted too that, in frequent experimental cases, these signals are corrupted by strong noise.

This work is partly supported by the French Agence Nationale de la Recherche, under project ECHANGE ANR-08-EMER-006

In view of this remarkable multichannel signal structure, this paper investigates the use of non-uniform sampling techniques to get the best estimation of these parameters for a fixed total number of samples. Reducing the number of samples not only reduces the global data throughput and the storage space needed to record the signals, in some cases it also reduces the computational complexity of the processing, as shown in this paper. It can be noted that incorporating prior information on the structure of the signal at the sampling stage is very much in the spirit of compressed sensing techniques, although the sinusoidal nature of the signals allows the use of specific resolution techniques.

In order to recover the parameters of the sinusoids from the non-uniform samples, we have studied two algorithms. The first one is an adaptation of Simultaneous Orthogonal Matching Pursuit (SOMP) [4], a generic multichannel sparse recovery algorithm. This extension is rather straightforward in the case of undamped sinusoids. The second one is based on the high resolution spectral estimator MUSIC (MUltiple SIgnal Classification) [5]. As we shall demonstrate, this technique can be extended in a natural way to our multichannel, non-uniformly sampled problem, with possibly damped sinusoids.

To summarize the original contributions of this paper, it presents a full study on how to handle non-uniform sampling for multichannel damped sinusoids (data model introduced section 2):

- the computation of Cramér-Rao bounds for the estimation of the signal components (section 3),
- the investigation of the optimal observation length, that as we shall demonstrate should be roughly 2 times the typical decay time of the signal (section 4),
- The generalization of two estimation algorithms, SOMP (section 5) and MUSIC (section 6), and the comparison of their performance in terms of precision, resolution, and computational cost (section 7).

2. DATA MODEL

The signals measured by the K sensors are combinations of P exponentially modulated sinusoids (impulse responses). A given sinusoidal component, indexed by p, has the same frequency ω_p and damping factor α_p across sensors, but different phases ϕ_{pk} and amplitudes A_{pk}. The measurements are distorted by a complex noise w_k, assumed white gaussian of variance $2\sigma^2$ and uncorrelated between sensors. The expression of the signal received by the k-th sensor at time t is then :

$$X_k(t) = \sum_{p=1}^{P} A_{pk} e^{-\alpha_p t} e^{i(\phi_{pk}+\omega_p t)} + w_k(t) \qquad (1)$$

Such signals, sums of the same atoms with different coefficients, can be described by the Joint Sparsity Model 2 (JSM-2) introduced in [6]. They are found in nearfield acoustical holography of normal modes of freely-vibrating plates, where each normal mode of the plate radiates a damped sinusoid, solution of a second-order differential equation. Amplitudes and phases of each mode, varying in space, are used to recover the modal shape.

These signals are sampled at M times t_m, identical for each channel.

3. CRAMER-RAO BOUNDS

Cramér-Rao bounds (CRBs) for monochannel damped sine waves can be found in [7]. Following the same computation, these results can be generalized to derive bounds for the multichannel case (K channels) of a single damped sinusoid affected by white Gaussian noise :

$$\begin{aligned} CRB_\omega &= \frac{1}{KR}\frac{a}{ac-b^2} \\ CRB_\alpha &= \frac{1}{KR}\frac{a}{ac-b^2} \\ CRB_{\phi_i} &= \frac{1}{R_i a} + \frac{1}{KR}\frac{b^2}{a(ac-b^2)} \\ CRB_{A_i} &= \frac{\sigma^2}{a} + \frac{A_i^2}{KR}\frac{b^2}{a(ac-b^2)} \end{aligned}$$

where $a=\sum_{m=1}^M e^{-2\alpha t_m}$, $b=\sum_{m=1}^M t_m e^{-2\alpha t_m}$, $c=\sum_{m=1}^M t_m^2 e^{-2\alpha t_m}$ and $R=\frac{1}{K}\sum_{k=1}^K \frac{A_k^2}{\sigma^2}$. These bounds are obviously lower than the monochannel ones, especially for the damping and the frequency parameters which are shared among the channels : whereas the monochannel bounds for frequency and damping are inversely proportional to the ratio of the amplitude squared at $t=0$ over the noise variance (ratio we will call Signal-to-Noise Ratio (SNR) in the remainder of the paper, with a slight abuse of language) the multichannel bounds are inversely proportional to KR, i.e. the sum of the SNRs in each channel. Phases and amplitudes, although different for each channel, also have better bounds : they are sums of two terms, one inversely proportional to the sum of the SNRs, and one inversely proportional to the SNR of the channel. This can be explained by the fact that a better estimation of the frequency and the damping gives a better estimation of the phases and the amplitudes. Both algorithms used in this paper show a similar behavior : phases and amplitudes are estimated by projecting the signal on the subspace spanned by the chosen atoms. A better estimation of this subspace will obviously lead to a better estimation of the parameters of the components.

4. OPTIMAL OBSERVATION LENGTH

In this section, we investigate how these Cramér-Rao bounds can be used to compute the optimal observation length T, as a function of the damping parameter α, for a fixed number of samples M.

Here, in order to obtain closed-form solutions, we assume that the signals are uniformly subsampled, i.e. $t_m = \frac{m}{M}T = m\Delta t$, for $0 \le m < M$. We can then rewrite the frequency and damping Cramér-Rao bounds as

$$CRB_\omega = CRB_\alpha = \frac{1}{KR}\frac{\alpha^2}{M} f(M, \alpha T) \qquad (2)$$

where

$$f(M,\beta) = \frac{M\sum_{m=0}^{M-1} f_1\left(\frac{\beta m}{M}\right)}{\sum_{m=0}^{M-1} f_1\left(\frac{\beta m}{M}\right)\sum_{m=0}^{M-1} f_3\left(\frac{\beta m}{M}\right) - \left(\sum_{m=0}^{M-1} f_2\left(\frac{\beta m}{M}\right)\right)^2}$$

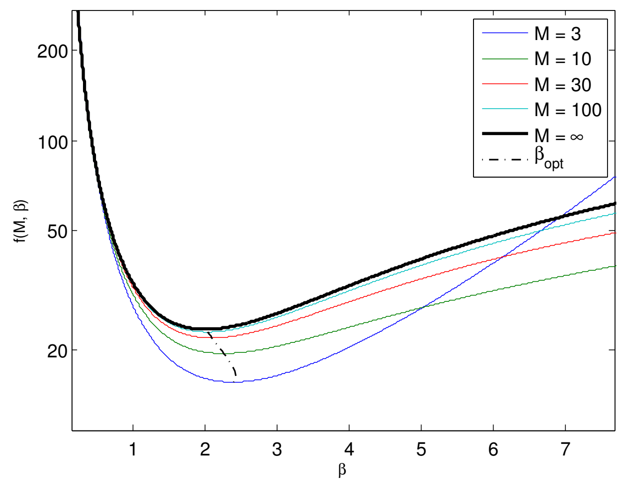

Fig. 1. Values of $f(M,\beta)$ that governs the CRB as a function of the observation duration $T=\beta/\alpha$. M is the number of samples in the interval $[0,T]$, α is the decay time of the sinusoid.

$$f_1(x) = e^{-2x}, f_2(x) = xe^{-2x}, f_3(x) = x^2 e^{-2x}$$

Here, the $\beta = \alpha T$ parameter represents the number of decay characteristic times $1/\alpha$ that is being used for the observation interval $[0,T] = [0,\beta/\alpha]$. The optimal observation length is the T_{opt} minimizing the lower bound. It should be noted that, as seen from Equation (2), this optimum is independent of the number K of channels and the SNR R. The limit of f as M goes to infinity can be computed (discrete sums are seen as Riemann sums of f_1, f_2 and f_3 and therefore converge to their integral on $[0,\beta]$) :

$$f(+\infty,\beta) = \frac{8\beta(e^{4\beta}-e^{2\beta})}{e^{4\beta}-2e^{2\beta}+4\beta^2+1}$$

Values of $f(M,\beta)$ for different M and $\beta=\alpha T$ and its limit for $M=\infty$ are plotted figure 1, where optimal values of β are indicated. We can see that the optimal value of β is approximately 2 for every M, thus the optimal observation length for a damped sinusoid is approximately $T_{opt} = \frac{2}{\alpha}$. The bottom line goes as follows: when analyzing exponentially damped sinusoids under a constraint on the total number M of samples, it is better not to take the first M samples, but rather to spread the samples across the signal. However, after too many characteristic decay times $\frac{1}{\alpha}$ the samples mostly represent noise and carry little information. For uniform sampling, a good trade-off is to sample within $[0, \frac{2}{\alpha}]$.

In practical estimation of sinusoidal components in a signal with uniform sampling, the bandwidth is limited by the Shannon-Nyquist theorem. A solution is to sample the signal irregularly. The following two sections describe different algorithms that can be used to recover the parameters of signals from irregular samples. Numerical tests in section 7 show that the optimal sampling duration $T \approx 2/\alpha$ still holds approximately for non-uniform sampling with a constant density on $[0,T]$.

5. SIMULTANEOUS ORTHOGONAL MATCHING PURSUIT

Simultaneous orthogonal matching pursuit [4], a greedy sparse approximation algorithm specialized in recovering JSM-2 signals, aims at identifying the components one at a time, by selecting the atom of a dictionary's most correlated, on average, with the signals, subtracting the contribution of this component, and iterating until the desired number of components are found.

An atom of frequency ω and damping α is $\psi_{\omega\alpha} = \gamma_{\omega\alpha}(e^{(i\omega-\alpha)t_1}, \dots, e^{(i\omega-\alpha)t_m}, \dots, e^{(i\omega-\alpha)t_M})$, with $\gamma_{\omega\alpha}$ such that $||\psi_{\omega\alpha}|| = 1$.

The algorithm works as follows :

1. Initialize the residuals $r_{k,0} = (X_k(t_1), \dots, X_k(t_M))^T$, set the iteration counter $l = 1$
2. Select the pair (ω_l, α_l) maximizing $\sum_{k=1}^{K} |\langle \psi_{\omega\alpha}, r_{k,l}\rangle|^p$.
3. Let $q_{k,l}$ be the orthogonal projection of $r_{k,l}$ on $\mathrm{span}(\psi_{\omega_1\alpha_1}, \dots, \psi_{\omega_l\alpha_l})$ and $r_{k,l+1} = r_{k,l} - q_{k,l}$
4. Increment l and iterate until enough components are found
5. Estimate phases and amplitudes with the coordinates of the projection of X_k in the space spanned by $(\psi_{\omega_1\alpha_1}, \dots, \psi_{\omega_l\alpha_l})$.

The parameter $p \geq 1$ controls the way information is integrated between channels. In this paper, p will always be equal to 2.

As this algorithm uses a Fourier approach to detect sinusoidal component, its resolution is limited, restricting its use to signals with clearly separated partials. When the signals contains sinusoids with very close frequencies, which indeed occurs in some experimental setups, e.g. holography of plates of particular geometries, high resolution methods can be helpful, as described in the next section.

6. MULTICHANNEL MUSIC

High resolution spectral estimation methods, such as MUSIC [5], are often used to go beyond the standard limitations of plain Fourier-based methods. Based on the computation of a signal subspace and a noise subspace, MUSIC requires an estimation of the covariance of the signal. This covariance is, in the traditional monochannel and regularly sampled case, estimated by a average over overlapping windows. In our case, this averaging is not done in this way, made impossible by the irregular sampling, but over the different signals available. The signals can be arranged in a matrix $\mathbf{X}$ of general term $x_{ij} = X_i(t_j)$. Our signal model 1 allows us to write $\mathbf{X}$ as:

$$\mathbf{X} = \mathbf{VA} + \mathbf{W}$$

where $\mathbf{V}$ is the matrix containing the damped sinusoids sampled at the times t_m, $v_{ij} = e^{(i\omega_j - \alpha_j)t_i}$, and $\mathbf{A}$ contains the amplitudes and phases of each partial, $a_{ij} = A_{ij}e^{i\phi_{ij}}$.

The first step is to estimate the matrix $\mathbf{V}$. The autocorrelation of the noiseless signals $\mathbf{R_u} = \mathbf{VA}(\mathbf{VA})^H$ has the same image as the matrix $\mathbf{V}$, under the condition that the matrix $\mathbf{A}$ has sufficient rank (this, among other conditions, requires more channels than components). This image, the signal subspace, is also the space related to the P strictly positive eigenvalues of $\mathbf{R_u}$. The nullspace is called noise subspace. The eigenvalues of the autocorrelation of the noisy signals $\mathbf{R_n} = E[\mathbf{XX}^H] = \mathbf{R_u} + \sigma^2\mathbf{I}$ are the eigenvalues of $\mathbf{R_u}$ augmented by σ^2. The signal subspace is thus the space related to the P largest eigenvalues of $\mathbf{R_n}$.

Frequencies and damping are found by searching the vectors $\psi_{\omega\alpha} = (e^{(i\omega-\alpha)t_1}, \dots, e^{(i\omega-\alpha)t_M})$ contained by the signal subspace. These vectors are orthogonal to the noise subspace, the length of their projection $\Pi_n(\psi_{\omega\alpha})$ in the noise subspace is 0. In practice, as we only have an estimate of $\mathbf{R_n}$, we pick the values of ω and α where the pseudospectrum $S(\omega, \alpha) = \frac{1}{||\Pi_n(\psi_{\omega\alpha})||^2}$ reaches a maximum.

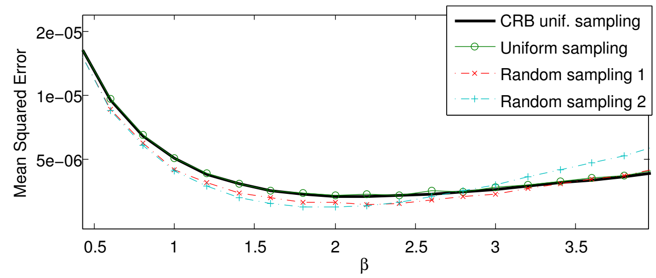

Fig. 2. Mean square error on the frequency estimation, SOMP, for different observations duration. SNR = 30dB, 20 samples, $\alpha = 0.1$.

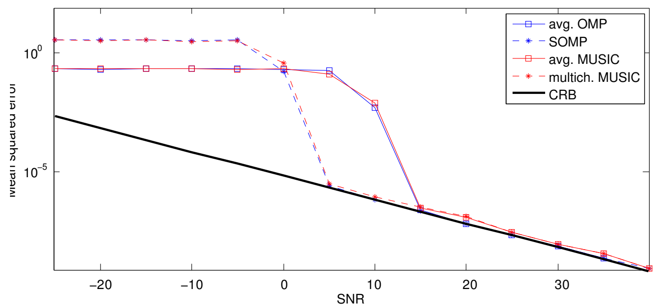

Fig. 3. Mean square error on the frequency estimation, SOMP compared to Cramér-Rao bounds

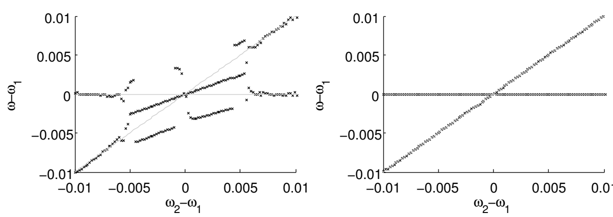

Fig. 4. Frequency estimations for two partials with close frequencies, 100 signals, no noise, 50 samples per signal. left: SOMP, right: nonuniform multichannel MUSIC

Amplitudes and phases are then found, as in SOMP, with the coordinates of the projection of the signals in the space spanned by the selected atoms.

It should be noted that the use of the MUSIC algorithm is standard in antenna processing, but in a framework where regularly sampled signals (in time) are gathered by different sensors (possibly irregularly sampled in space) to infer directions of arrival. The novelty is to use here the dual view of the same problem, where the signals are irregularly sampled in time, under the constraint that the sampling times t_m are identical across channels to determine frequencies and dampings of sinusoidal components of signals.

7. PERFORMANCES

7.1. Accuracy

The accuracy of SOMP and MUSIC is compared to the Cramér-Rao bound in two cases. For both cases, the signals are sum of a unique damped sinusoid with varying amplitudes and phases over the channels, and gaussian white noise.

First, with fixed SNR (30 dB), number of samples ($M = 20$) damping ($\alpha = 0.1$) and number of channels ($K = 128$), we estimate the variance of the estimators for different observation length T, with uniform sampling and nonuniform sampling (sampling times drawn from a uniform distribution on $[0, T]$). These estimated variances are plotted on Fig. 2 in the case of SOMP (MUSIC exhibits a similar behavior). Estimation with uniform sampling has good performances compared to the CRBs, and has the same behaviour, with a minimum at approximately $\beta = 2$. Variances for nonuniform sampling are similar, with minima close to $\beta = 2$. It should be noted that some of the variances are lower than the Cramér-Rao bounds for uniform sampling, which is not surprising as some random choices from nonuniform sampling times may better catch the structure of the signal. These results justify the arguments of section 4, as the optimal observation length $T \approx 2/\alpha$ given by the CRB in the uniform case is close to the actual optimal length for nonuniform sampling.

The second case is with varying SNRs, with fixed number of channels ($K = 10$), damping ($\alpha = 0.01$), number of samples ($M = 20$), and quasi optimal observation length ($T = 200$). We compare on Fig. 3 the joint frequency estimation (SOMP, multichannel MUSIC), and the averaged monochannel estimation (OMP, monochannel MUSIC). For low SNRs, joint and averaged estimations have similar performances, close the Cramér-Rao bounds, but between SNRs of 5 dB and 15 dB, separate estimations fails, whereas joint estimation still recovers the frequency of the sinusoids.

7.2. Resolution

The resolution of SOMP and MUSIC are compared with signals containing two components, with the same amplitude and random phase for each channel. We have here 50 samples picked from 1024 regularly spaced times, and 100 channels, with no noise. Figure 4 shows that SOMP ($p = 2$) does not recover the two components when their frequencies are close, whereas MUSIC only fails when frequencies coincide. Furthermore, the frequencies estimated by SOMP are biased even with well separated frequencies, where MUSIC does not show such a bias.

7.3. Computational complexity

Here, we compare the complexity of the different algorithms for a given target accuracy, tuning for each of them the number of samples in such a way that the error variance stays below 3×10^{-10}. The fastest algorithm is multichannel MUSIC with nonuniform sampling, as it needs the lowest number of samples per channel. Table 1 shows computation times for 100 channels and 1 channel, uniform and nonuniform sampling, using Matlab on a 3.2 GHz processor. MUSIC is more efficient for multichannel estimation, while SOMP (actually OMP) is faster for monochannel estimation. This can be explained by the fact that the complexity of MUSIC is dominated by the computation of the pseudo-spectrum of complexity $\mathcal{O}(NM^2)$, identical for multichannel and monochannel estimation, while SOMP is dominated by the computation of the correlations, of complexity $\mathcal{O}(KNM)$, linear in the number of channels.

	nonuniform sampling		uniform sampling	
	100 ch.	1 ch.	100 ch.	1 ch.
samples per channel	10	1000	350	1600
MUSIC	0.06	70.2	1.9	170.89
SOMP ($p = 1$ or 2)	0.21	0.96	3.2	1.48

Table 1. Computation time (seconds) for one sinusoidal component, and similar accuracy.

8. CONCLUSION

There are many experimental cases where antennas of sensors record a small number of sinusoidal components, possibly damped, varying only in phase and amplitude on every sensor. In this paper we have shown that in these cases, it is possible to use non-uniform samples which brings a significant advantage in terms of complexity, and memory requirements. In the damped case, CRBs show that the samples are best spread over $[0, 2/\alpha]$, i.e. over 2 times the typical decay time of the sinusoids.

To simultaneously recover the sinusoidal parameters in all the channels, we have proposed and compared two algorithms: a generic distributed compressed sensing algorithm, SOMP, and a variant of a spectral estimation algorithm, MUSIC. Both algorithms use the common sparsity of the channels and the larger period of time spanned by the samples to give a sharper accuracy for the frequency estimation, not so far from the Cramér-Rao bounds for reasonable SNRs. However, both algorithm have drawbacks : SOMP, based on a Fourier approach, cannot resolve close components, and MUSIC can have a large complexity for a small number of channels.

9. ACKNOWLEDGEMENTS

The authors would like to thank Bertrand David and Gérard Blanchet for useful discussions on high-resolution methods.

10. REFERENCES

[1] B.D. Van Veen and K.M. Buckley, "Beamforming: a versatile approach to spatial filtering," *ASSP Magazine, IEEE*, vol. 5, no. 2, pp. 4–24, April 1988.

[2] J. D. Maynard, E. G. Williams, and Y. Lee, "Nearfield acoustic holography: I. theory of generalized holography and the development of nah," *The Journal of the Acoustical Society of America*, vol. 78, no. 4, pp. 1395–1413, 1985.

[3] John C. Curlander and Robert N. McDonough, *Synthetic Aperture Radar: Systems and Signal Processing*, Wiley-Interscience, 1991.

[4] R. Gribonval, H. Rauhut, K. Schnass, and P. Vandergheynst, "Atoms of all channels, unite! average case analysis of multichannel sparse recovery using greedy algorithms," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 655–687, 2008.

[5] R. Schmidt, "Multiple emitter location and signal parameter estimation," *Antennas and Propagation, IEEE Transactions on*, vol. 34, no. 3, pp. 276–280, Mar 1986.

[6] Dror Baron, Marco F. Duarte, Michael B. Wakin, Shriram Sarvotham, and Richard G. Baraniuk, "Distributed compressive sensing," *CoRR*, vol. abs/0901.3403, 2009.

[7] T. Wigren and A. Nehorai, "Asymptotic cramer-rao bounds for estimation of the parameters of damped sine waves in noise," *Signal Processing, IEEE Transactions on*, vol. 39, no. 4, pp. 1017–1020, Apr 1991.