RESEARCH

A hybrid Fourier-Prony method

Matteo Briani*†, Annie Cuyt and Wen-shin Lee

Abstract

The discrete Fourier transform is one of the top ten algorithms of the 20-th century, frequently used in the analysis of signals. It is known to be very stable and of low computational complexity. When the bandwidth is large or the spectrum is sparse, sampling at the Shannon-Nyquist rate may however be prohibitive.

We propose a new hybrid technique that exploits the spectral sparsity. Several shifted and decimated but uniformly sampled data sets are collected. The aliased discrete Fourier spectra are corrected for the aliasing effect by a combination with Prony's method. In our approach both the decimation factor and the time shift can be large and thus violate the Shannon-Nyquist theorem.

The resulting technique delivers the same accuracy as the non-decimated Fourier transform, uses fewer samples and achieves a smaller computational cost. Also, the method is suitable for parallellization.

Keywords: Fourier transform; Prony's method; sparse representation

AMS Subject Classification: Primary 65T50

1 Introduction

The Discrete Fourier Transform (DFT) is widely used ever since the beginning of the digital era, especially in digital signal processing. It is now possible to analyze large streams of samples using procedures of low computational complexity. Traditionally signals are analyzed that exhibit a dense structure in the frequency domain. Recently the interest has shifted to the analysis of signals that have a sparse frequency representation [1].

*Correspondence: matteo.briani@uantwerpen.be
Universiteit Antwerpen, Middelheimlaan 1, Antwerp, BE
Full list of author information is available at the end of the article
†This research is supported by the Instituut voor Wetenschap en Technology - IWT

In several applications a serious need arises to exploit this sparsity. For instance, high speed analog-to-digital converters go hand in hand with a high cost and high consumption but low bit resolution [2]. In order to capture inputs within a wide spectral range, the sparsity of the signal has to be exploited. Another domain is that of inverse source problems related to physics driven fields governed by linear partial differential equations [3, 4]. These fields are often non-bandlimited and therefore require extremely dense sampling to be in accordance with the classical Shannon-Nyquist sampling theorem [5, 6].

So-called parametric methods [7, 8, 9], based on Prony's result [10], are an obvious choice to make good use of sparsity in the frequency domain. The frequency resolution of these methods is not bound by the amount and time span of the available samples, unlike with Fourier methods. However, some extra care has to be taken because of their sensitivity to noise [11]. We restrict ourselves here to a discussion of sub-Nyquist sampling techniques that adhere to some underlying uniform sampling scheme, which is in the focus of our new approach. We do not compare our contribution to methods containing a probabilistic element such as [12], where (pseudo)-random permutation is used, or [13] which is based on the compressed sensing paradigm.

Several papers venture into uniform or regular sub-Nyquist sampling. In [14] the DFT is computed from fewer samples than required, collected at a rate that does not obey the Shannon-Nyquist theorem, but the spectral analysis suffers from aliasing. This is corrected by means of an additional set of samples that is close enough to the first one, meaning that the time shift is smaller than the Nyquist sampling step. A result of the same kind is presented in [15] where an alias free DFT results from applying the Chinese remainder theorem to two sets of samples collected at sub-Nyquist rates that are coprime. The concept of coprime samplers is also used in [16] where two DFT filter banks using different sampling rates are combined to achieve a higher spectral accuracy than achieved by each filter bank separately. In [17, 18] a sparse DFT is computed from two sub-sampled signals differing by a time shift of one Shannon-Nyquist step. In case a collision of frequencies is detected as a result of possible aliasing,

the implementation resolves it by using other subsamplings at coprime sampling rates. In [19] several DFT are computed again with an offset of a single sample respecting the Shannon-Nyquist bound. A parametric method is involved to output the final spectral analysis.

What underlies all of the above is that either coprime sampling rates are used or a second set of time shifted samples is collected where the time shift is small enough. We generalize this approach in the sense that we use only one sub-Nyquist sampling rate and allow for a large time shift [20, 21] that need only be coprime with the decimation factor used in the subsampling. This larger offset evidently imposes less restrictions on the data acquisition process. If not addressed properly however, this extension introduces additional aliasing in the time shifted spectral analysis. We combine the undersampled DFT result with a Prony-related method because the latter arises naturally when analyzing the time shifted signals. The Prony systems are small in size and allow to resolve the aliasing and identify the non-aliased frequencies. Also, the method offers an easy way to deal with collisions. This has mostly remained an open problem in other approaches: here both the number of collisions and the individual components can be unravelled. Also note that the use of multiple decimation factors, as sometimes proposed in other sub-Nyquist methods, implies that different collisions arise in the different decimated sample sets, a problem that we can avoid with the current technique.

The paper is organized as follows. The required notation for subsampled and time shifted data is introduced in Section 2 and their effect is discussed. In Section 3 the core of our new technique is presented, while we address the issue of frequency collision in Section 4. In Section 5 we explain the effect of noise on the data and discuss the computational complexity of the new method. Everything is illustrated on numerical examples in Section 6, where we deal with synthetic and noisy signals suffering from several collisions due to the subsampling strategy.

2 Undersampling and time shifts

Let us consider a finite vector of samples $x = (x_0, \dots, x_{N-1})$ of a complex-valued function $\mathcal{X}$, collected on a uniform time grid. The Discrete Fourier Transform (DFT) coefficients $X = (X_0, \dots, X_{N-1})$ are given by

$$X_j := \sum_{l=0}^{N-1} x_l \exp(-2\pi \mathrm{i} lj/N), \qquad j = 0, \dots, N-1, \qquad \mathrm{i}^2 = -1.$$

The Fourier coefficient X_j represents the amount of the specific complex exponential $\exp(-2\pi \mathrm{i} j/N)$ that is present in the discrete signal x. We say that X_j is associated to the frequency j, meaning to the associated complex exponential. The samples x can be reobtained from the Fourier coefficients X using the Inverse Discrete Fourier Transform (IDFT), as

$$x_l := \frac{1}{N} \sum_{j=0}^{N-1} X_j \exp(2\pi \mathrm{i} lj/N), \qquad l = 0, \dots, N-1. \quad (1)$$

The Shannon-Nyquist theorem [5, 6] states that one can exactly reconstruct bandlimited signals that have the same bandwidth as the sampling rate. If the analyzed signal contains higher frequencies, then aliasing is encountered. In Section 3 we explain how to tackle this problem. If not stated otherwise, we assume that the signal x satisfies the Shannon-Nyquist theorem.

We also assume that the signal x has a sparse representation in the frequency domain, meaning that it has only $K \ll N$ non-vanishing Fourier coefficients $X_{j_1}, \dots, X_{j_K}, 1 \le k \le K$. In case the signal is perturbed by noise, we set a threshold $T \in \mathbb{R}$ and we assume that only K Fourier coefficients have an amplitude larger than T,

$$|X_{j_k}| \ge T, \qquad k = 1, \dots, K. \quad (2)$$

We now consider the following (possibly sub-Nyquist) undersampled version of the signal x consisting only of samples x_{lu},

$$\begin{aligned} {}_u x &:= (x_0, 0, \dots, 0, x_u, \dots, x_{N-u}, 0, \dots, 0) \\ &= ({}_u x_0, {}_u x_1, \dots, {}_u x_{N-1}), \qquad N/u \in \mathbb{N} \end{aligned}$$

and denote the DFT of ${}_u x$ by ${}_u X$. A larger distance between the samples implies that the highest retrievable frequency is now smaller. Some coefficients

$${}_u X_j = \sum_{l=0}^{N/u-1} x_{ul} \exp(-2\pi \mathrm{i} luj/N), \qquad j = 0, \dots, N/u-1$$

may have been associated with smaller frequencies. Given a specific Fourier coefficient, there is no way to know if it has been affected by aliasing or not: the Fourier coefficient X_{j_k} may now appear at a different index $\tilde{j_k}$. In particular

$$X_{j_k} = {}_u X_{\tilde{j_k}}, \qquad j_k = u\tilde{j_k} \bmod N. \quad (3)$$

Moreover, it can happen that a non-vanishing Fourier coefficient is mapped down to an exponential already associated with another non-vanishing coefficient. In that case the coefficient ${}_uX_{\tilde{j_k}}$ is the sum of two or more coefficients of the full-length DFT. We refer to this phenomenon as frequency collision. In order to leave the notation light, from here on we do not explicitly put a tilde over the aliased indices. More on frequency collision and the role of the Fourier leakage effect in this respect is found in Section 4.

We denote a shifted version of the signal x by

$$\begin{aligned} {}^sx &:= \left(x_s, x_{1+s}, \ldots, x_{(N-1)+s}\right) \\ &= \left({}^sx_0, {}^sx_1, \ldots, {}^sx_{N-1}\right), \qquad s \in \mathbb{Z} \end{aligned}$$

where we take $x_{l+\lambda N} = x_l, \lambda \in \mathbb{N}$. With X and ${}^sX$ respectively denoting the noise-free Fourier coefficients computed from x and ${}^sx$, we have

$$\begin{aligned} {}^sX_j &= \sum_{l=0}^{N-1} {}^sx_l \exp(-2\pi \mathrm{i} lj/N) \\ &= \sum_{l=-s}^{N-1-s} {}^sx_l \exp(-2\pi \mathrm{i} lj/N) \\ &= \exp(2\pi \mathrm{i} sj/N) \times \\ &\quad \sum_{l=-s}^{N-1-s} {}^sx_l \exp(-2\pi \mathrm{i} lj/N) \exp(-2\pi \mathrm{i} sj/N) \\ &= \exp(2\pi \mathrm{i} sj/N) \times \\ &\quad \sum_{l+s=0}^{N-1} x_{l+s} \exp(-2\pi \mathrm{i}(l+s)j/N) \\ &= \exp(2\pi \mathrm{i} sj/N) X_j, \qquad j = 0, \ldots, N-1. \end{aligned} \tag{4}$$

In other words, the shifted vector has the same Fourier coefficients X_j multiplied by the complex exponential $\exp(2\pi \mathrm{i} sj/N)$. In the time domain this corresponds to a phase change of the signal. Analogously, in a collision free context,

$${}^s_uX_j = \exp(2\pi \mathrm{i} sj/N)\ {}_uX_j, \qquad j = 0, \ldots, N/u-1, \tag{5}$$

meaning that each decimated shifted DFT equals ${}_uX$ times a phase component. Moreover, (5) indicates that in case of a frequency collision, the collision index remains the same after the shift. Of course, if noise is added to the Fourier coefficients, the relation between ${}^sX_j$ and X_j is not exact anymore due to the non-periodicity of the noise. In Section 3 we use (5) to resolve the aliasing issue.

Let us take a look at the particular samples x and ${}^1x$ and compute their respective non-vanishing Fourier coefficients X_{j_k} and ${}^1X_{j_k}$ for $k = 1, \ldots, K$. When the sampling was performed at Nyquist rate, then the division ${}^1X_{j_k}/X_{j_k}$ delivers the associated exponential $\exp(2\pi \mathrm{i} j_k/N)$, as in [14] and [19]. More generally, for $s > 1$ the division ${}^sX_{j_k}/X_{j_k}$ delivers $\exp(2\pi \mathrm{i} s j_k/N)$, from which only a set of plausible solutions for the associated exponential can be obtained, namely

$$\begin{aligned} S_{j_k} &:= \{S_{j_k,l}, l = 0, \ldots, s-1\}, \\ S_{j_k,l} &:= \exp\left(2\pi \mathrm{i} s j_k/N + 2\pi \mathrm{i} l/s\right). \end{aligned} \tag{6}$$

It is not possible to know which of the exponentials in S_{j_k} is the correct one. However, the use of particular values of u and s allows to answer this question.

3 Fixing aliasing

Each Fourier coefficient of an undersampled vector ${}_ux$ may be affected by aliasing. A smart way of using the parameters u and s can fix the aliasing. The key point is to choose u and s to be coprime (for more general choices we refer to [22]). The procedure goes as follows.

Calculate the DFT of ${}_ux$ and compute for each non-vanishing Fourier coefficient ${}_uX_{j_k}$ the set

$$\begin{aligned} U_{j_k} &:= \{U_{j_k,l}, l = 0, \ldots, u-1\}, \\ U_{j_k,l} &:= \exp\left(2\pi \mathrm{i} u j_k/N + 2\pi \mathrm{i} l/u\right). \end{aligned}$$

Then consider the time-shifted vector of samples ${}^s_ux$ and compute the Fourier coefficients ${}^s_uX_{j_k}$. For each non-vanishing j_k the division ${}^s_uX_{j_k}/{}_uX_{j_k}$ returns the exponential $\exp(2\pi \mathrm{i} s j_k/N)$ from which the set (6) follows. Now both U_{j_k} and S_{j_k} contain the non-aliased exponential associated to the Fourier coefficient. Since u and s are coprime, these two sets share one and only one element which is the correct non-aliased exponential (see also [20, 21, 23] where the time shift has appropriately been termed an identification shift). Figure 1 summarizes the procedure.

Note that the set S_{j_k} can also be obtained without computing the DFT coefficients ${}^s_uX$. From (1) the vector of ${}^s_uX_{j_k}/N, k = 1, \ldots, K$ can be obtained as the solution of a Vandermonde structured linear system of equations with coefficient matrix

$$V = \begin{pmatrix} 1 & \cdots & 1 \\ \exp(2\pi \mathrm{i} u j_1/N) & \cdots & \exp(2\pi \mathrm{i} u j_K/N) \\ \exp(2\pi \mathrm{i} 2 u j_1/N) & \cdots & \exp(2\pi \mathrm{i} 2 u j_K/N) \\ \vdots & & \vdots \\ \exp(2\pi \mathrm{i}(K-1) u j_1/N) & \cdots & \exp(2\pi \mathrm{i}(K-1) u j_K/N) \end{pmatrix}$$

and right hand side $(x_s, x_{u+s}, \ldots, x_{(K-1)u+s})$.

To compute the intersection of the sets U_{j_k} and S_{j_k}, several ways can be used: either the distance between the elements in the set, or the Euclidean algorithm, or our robust matching technique. In the first approach a distance matrix is built, containing all distances between elements of U_{j_k} and elements in S_{j_k}. The pair for which the distance is minimal is then selected to represent the true non-aliased exponential. Of course, one can replace the distance $||U_{j_k,l} - S_{j_k,m}||$ by the inner product $< U_{j_k,l}, S_{j_k,m} >, l = 0, \ldots, u-1, m = 0, \ldots, s-1$ and look for the maximum value, which is 1 in the ideal case. But when u and s increase, the computational cost increases. In that case the Euclidean algorithm offers an alternative. It states that for integer coprime u and s there exist (non-unique) integers v and t such that $uv + st = 1$. Then

$$\begin{aligned}&\exp(2\pi \mathrm{i} uvj_k/N)\exp(2\pi \mathrm{i} stj_k/N)\\ &\quad = \exp\left(2\pi \mathrm{i}(uv+st)j_k/N\right) = \exp(2\pi \mathrm{i} j_k/N)\end{aligned}$$

returns the exponential we are looking for. The downside of this approach is the sensitivity to noise. In fact, when computing $\exp(2\pi \mathrm{i}(uv + st)j_k/N)$ we are powering $\exp(2\pi \mathrm{i} uj_k/N)$ and $\exp(2\pi \mathrm{i} sj_k/N)$. Since both these exponentials are corrupted by the noise present in the Fourier coefficients, the noise is also powered, except when the integers v and t have sufficiently small absolute values.

To counter the latter, a third approach can be used. Let ${}_uX_{j_k}$ denote a peak in the aliased DFT. Assume that we in addition have several time shifted DFT ${}^{s_h}_uX, h = 1, \ldots S$ at our disposal (usually $S \geq 3$). In Section 4 a similar need and use will arise. Also assume that all s_h are coprime with u. In order to identify the correct exponential $U_{j_k,l}$ in the set U_{j_k} we construct for $l = 0, \ldots, u-1$ the products

$${}_uX_{j_k}\left(U_{j_k,l}\right)^{s_h}, \qquad h = 0, 1, \ldots, S, \qquad s_0 = 0$$

and compare these with the available ${}^{s_h}_uX_{j_k}$, $h = 0, \ldots, S$. For the comparison the inner product

$$\langle {}^{s_h}_uX_{j_k}, {}_uX_{j_k}\left(U_{j_k,l}\right)^{s_h}\rangle$$

can be used. The correct $U_{j_k,l}$ is the one that returns the largest inner product modulus.

4 Frequency collision

When no frequency collision occurs, the proposed method guarantees the correct retrieval of the Fourier coefficients affected by aliasing. However, this ideal situation rarely happens due to the Fourier leakage effect. Leakage appears when a component in the signal x cannot exactly be represented by a single complex exponential of the form $\exp(2\pi \mathrm{i} j/N), j \in \mathbb{N}$. In this case, the specific frequency leaks to neighbouring frequencies and it affects all Fourier coefficients [24]. The effect is more evident in the Fourier coefficients near the location of the exact frequency.

The situation becomes problematic in the undersampled signal ${}_ux$. Because of aliasing the Fourier coefficients are mapped to different complex exponentials. Moreover, if we have leakage, frequency collisions may occur because a component may have leaked over all Fourier coefficients. This is limiting the applicability of the proposed method because the division ${}^s_uX_j/{}_uX_j$ does not yield one exponential. The same limitation exists in [14] and it is precisely our objective to show how to deal with frequency collisions.

Let X_{j_1} and X_{j_2} be two non-vanishing Fourier coefficients of the data vector x. We consider ${}_uX$ and we assume that the two coefficients are colliding at the index $\hat{j}$ in ${}_uX_{\hat{j}}$. We also compute ${}^s_uX_{\hat{j}}$ and recall from (5) that the shift factor s does not influence the index at which the frequencies are colliding. It only affects the phase of the Fourier coefficient. When dividing, we obtain

$$\begin{aligned}\frac{{}^s_uX_{\hat{j}}}{{}_uX_{\hat{j}}} &= \frac{{}^sX_{j_1} + {}^sX_{j_2}}{X_{j_1} + X_{j_2}}\\ &= \frac{\exp(2\pi \mathrm{i} sj_1/N)X_{j_1} + \exp(2\pi \mathrm{i} sj_2/N)X_{j_2}}{X_{j_1} + X_{j_2}}\end{aligned}$$

and we are unable to extract the complex exponentials $\exp(2\pi \mathrm{i} sj_1/N)$ and $\exp(2\pi \mathrm{i} sj_2/N)$. We are thus not able to fix the aliasing yet. But by combining with an exponential analysis method based on Prony's algorithm, we can continue.

Let us consider the discrete signals ${}_ux, {}^s_ux, \ldots, {}^{(M-1)s}_ux$ and their respective DFTs ${}_uX, {}^s_uX, \ldots, {}^{(M-1)s}_uX$. The frequencies j_1 and j_2 are still colliding at the index $\hat{j}$ of ${}^{ms}_uX$ and this for all m. Each Fourier coefficient ${}^{ms}_uX_{\hat{j}}$ equals

$$\begin{aligned}{}_uX_{\hat{j}} &= X_{j_1} + X_{j_2},\\ {}^s_uX_{\hat{j}} &= \exp(2\pi \mathrm{i} sj_1/N)X_{j_1}\\ &\quad + \exp(2\pi \mathrm{i} sj_2/N)X_{j_2},\\ &\vdots\\ {}^{ms}_uX_{\hat{j}} &= \exp(2\pi \mathrm{i} msj_1/N)X_{j_1}\\ &\quad + \exp(2\pi \mathrm{i} msj_2/N)X_{j_2},\\ &\vdots\\ {}^{(M-1)s}_uX_{\hat{j}} &= \exp(2\pi \mathrm{i}(M-1)sj_1/N)X_{j_1}\\ &\quad + \exp(2\pi \mathrm{i}(M-1)sj_2/N)X_{j_2}. \quad (7)\end{aligned}$$

From (7) we can extract the exponentials $\exp(2\pi \mathrm{i} s j_1/N)$ and $\exp(2\pi \mathrm{i} s j_2/N)$ using a Prony-like parametric method such as [7, 8, 9]. More precisely, given the sequence ${}_u^{ms}X_{\hat{j}}, m = 0, \ldots, M-1$, such a method returns the coefficients X_{j_1}, X_{j_2} and their associated exponentials $\exp(2\pi \mathrm{i} s j_1/N), \exp(2\pi \mathrm{i} s j_1/N)$. Also note that the frequency resolution of Prony-like methods is not restricted to a pre-assigned grid. On the other hand, a drawback of these methods is their susceptibility to noise. In Section 5 we explain how to deal with noise.

Besides being able to separate colliding terms and aliased frequencies, it is also important to have a test for collision. In [14] the authors do not make use of repeated shifts and cannot deal with collisions. They check the necessary condition for a coefficient ${}_uX_j$ to be collision free, namely

$$\left| \frac{{}_u^s X_j}{{}_u X_j} \right| = |\exp(2\pi \mathrm{i} s j/N)| = 1. \tag{8}$$

It is however more reliable to analyze (7). Using the ${}_u^{ms}X_{\hat{j}}, m = 0, \ldots, (M-1)s$, the number of meaningful components in ${}_uX_{\hat{j}}$ can be extracted from the singular value decomposition (SVD) of a Hankel matrix built with the ${}_u^{ms}X_{\hat{j}}$ [25], thus detecting a collision. This is illustrated in Section 6.

5 Handling noise

Let us summarize the algorithm. Given a signal x, we fix an undersampling factor u, a shift s coprime with u and a total number M of shift repetitions. Then we compute the DFT vectors

$${}_uX, {}_u^sX, \ldots, {}_u^{(M-1)s}X.$$

The peaks of ${}_uX$ correspond to a sum of one or more frequencies that may have collided because of possible aliasing. Each peak ${}_uX_{\hat{j}}$ generates a set $U_{\hat{j}}$. For a fixed peak index $\hat{j}$ we inspect the coefficients in (7) and extract the components

$$\exp(2\pi \mathrm{i} s j_1/N), \exp(2\pi \mathrm{i} s j_2/N), \ldots$$

using a Prony-like method. The number of components ν in the sum is also revealed by Prony's method, as indicated above, and the Fourier coefficients $X_{j_1}, X_{j_2}, \ldots$ are obtained from the Vandermonde structured linear system (7). Each of these components then brings forth a set $S_{j_1}, S_{j_2}, \ldots$ which needs to be intersected with the set of exponentials $U_{\hat{j}}$ to find the true location of the non-aliased frequencies.

Now let the signal x be corrupted by complex white Gaussian noise $n = (n_0, \ldots, n_{N-1})$. We recall that the signal is still considered K-sparse if the threshold T introduced in Section 2 is such that (2) holds. Because of the noise, the DFT are all corrupted and then so are the peaks in (7). However, a connection between Prony's method and Padé approximation theory offers a way to separate the noise from the true values ${}_u^{ms}X_{\hat{j}}$. In fact, through the Z-transform the frequencies in the exponential terms of (7) correspond to the poles in the Padé approximant $[\nu - 1/\nu](z)$ of degree $\nu - 1$ in the numerator and ν in the denominator for the partial sum

$$\sum_{m=0}^{M-1} {}_u^{ms}X_{\hat{j}} z^{-m}, \qquad M \geq 2\nu.$$

From [26, 27, 28] we see that modelling (7) using more than ν poles hugely improves the result. The extra poles allow to model the noise and in doing so push the poles corresponding to the actual frequencies closer to their true locations. We thus approximate the ${}_u^{ms}X_{\hat{j}}$ with more exponential terms than needed, but the additional terms serve to model the noise n. Of course, this requires the collection of additional samples in (7), in other words, the need of a larger M when replacing ν by a larger value. The total number of required shifted DFT vectors ${}_u^{ms}X$ depends on the SNR and the number of components ν colliding in ${}_uX_{\hat{j}}$. An example illustrating the procedure is presented below.

Let us analyze the computational complexity of the new strategy. The signal length is N and the number of meaningful Fourier coefficients is $K \leq N$. The undersampling rate u, the time shift s and the number of shift repetitions $M-1$ are provided by the user. First the decimated DFTs ${}_u^{ms}X, m = 0, \ldots, M-1$ are computed with a computational cost of $O(MN/u \log(N/u))$. These DFTs contain at most K meaningful Fourier coefficients, possibly less when collisions occur (we neglect the fact that meaningless coefficients may have collided and exceeded the threshold T given in (2)). For each decimated Fourier coefficient ${}_uX_{j_k}, k = 1, \ldots, K$ a Prony system of M equations (7) is solved for the (still aliased) composing frequencies, using for instance the ESPRIT implementation of which the cost is dominated by the complexity $O(M^3)$ of the singular value decomposition. Additionally a Vandermonde linear system is solved in $O(M^2)$ operations when exploiting the structure of the linear system. The final identification of the intersection of the sets U_{j_k} and S_{j_k} costs $O(Kus)$ operations, leading us to a grand total of

$$O(KM^3 + MN/u \log(N/u) + Kus).$$

We remark that larger values of the decimation factor u allow for the computation of smaller N/u sized

DFTs, but at the cost of a possibly higher number of collisions, which necessitates the choice of larger values for M. On the other hand, smaller decimation factors u leed to larger sized DFTs with less chance for collision and hence no need for large M values.

Before proceeding with the numerical illustration we point out some interesting computational aspects. The method, as described above:

- immediately allows for M to be updated iteratively, and
- the Prony step in the procedure can easily be parallellized.

The former because increasing M merely implies that more shifted ${}_u^{ms}X$ need to be collected and the Prony step needs to be repeated. The latter because all the peaks ${}_uX_j$ for distinct j can be disentangled entirely by independent Prony steps.

6 Numerical illustration

We generate a few synthetic signals following the model

$$x_l = \sum_{i=1}^{L} \alpha_i \exp(2\pi \mathrm{i} \mu_i l/N), \qquad l = 0, \ldots, N-1. \quad (9)$$

In the first example, containing three test signals, we focus on the disentanglement of colliding frequencies. The model parameters are as in Table 1. In the second example we investigate the effect of a decreasing SNR, in other words an increasing noise level, in addition to disentangling a number of collisions. We have $L = 8$ with the phases and amplitudes randomly generated in the interval $[1, 3]$ and the other parameters as in Table 2.

For the three signals in Table 1, we set $N = 1000$ and the circular white Gaussian noise level to 30 dB. If we choose $u = 100$, the actual sampling rate becomes $1000/100 = 10$ which makes all frequencies collide in the same Fourier bin at 5 Hz. But these collisions can be resolved. We take $s = 3$ and $M = 12$. With $N = 1000$ each DFT ${}_u^{ms}X, m = 0, \ldots, M-1$ is computed up to a length of $N/u = 10$ samples. In Figure 2 at the left, one finds all 12 DFTs ${}_u^{ms}X$ for each of the three signals. At the right the results using on the one hand the standard DFT computed from $N = 1000$ samples and on the other hand the output of the proposed method are found. Figure 3 depicts for each of the three signals the computed amplitudes and the number of collided frequencies as indicated by a singular value decomposition of the $(M/2) \times (M/2)$ matrices involved in the Prony method [28].

In the second example $N = 20000, u = 400$ and the level of the white Gaussian noise ranges up to -10 dB. We investigate the use of increasing M values, namely $M = 16, 24, 32$. The $L = 8$ frequencies are chosen so that they collide in 4 Fourier bins. For the time offset we take $s = 7$. The results of the experiment are reported in Figure 4. All the exponential terms are retrieved correctly up to an SNR of 10 dB. Note that the method becomes less prone to noise as the value for M increases.

7 Conclusion

The newly proposed method belongs to the family of sparse Fourier techniques, which are especially useful when dealing with a sparse signal spectrum. Input are several uniform u-fold downsampled data sets, which are shifted with respect to each other by the same time shift s. Here the time shift need not be bounded by the Shannon-Nyquist interval, which poses less restrictions on the data acquisition process. Also, all the downsamplings are done at the same possibly sub-Nyquist rate, which implies that across the downsampled spectra the same aliasing effects arise.

First, the aliased DFTs are computed from the shifted downsampled data. Then, these DFT coefficients are reinterpreted in (7) as samples of a Prony system. The aliasing is fixed and possible frequency collisions are resolved by intersecting the frequency output associated to the aliased DFT coefficient with that associated to the aliased Prony system. To this end u and s need to be chosen coprime. The size of the Prony system, in other words the number of shift repetitions, can be tuned to the noise level in the data.

Since the Prony systems are used to disentangle collisions in a particular Fourier coefficient, the Prony step in the procedure can easily be parallellized. In the near future we plan to generalize the technique for the computation of a multidimensional sparse Fourier transform [29, 20, 21].

Declarations

Ethics approval and consent to participate
Not applicable

Consent for publication
None of the authors have any competing interests in the manuscript

Availability of data and material
Not applicable

Competing interests
Not applicable

Funding
Matteo Briani's research is supported by the Instituut voor Wetenschap en Technology - IWT, Belgium

Authors' contributions
Not applicable

Acknowledgements
Not applicable

Authors' information (optional)
Matteo Briani matteo.briani@uantwerpen.be
Annie Cuyt annie.cuyt@uantwerpen.be
Wen-shin Lee wen-shin.lee@uantwerpen.be

References

1. Gilbert, A.C., Indyk, P., Iwen, M., Schmidt, L.: Recent developments in the sparse Fourier transform: A compressed Fourier transform for big data. IEEE Signal Processing Magazine **31**(5), 91–100 (2014). doi:10.1109/MSP.2014.2329131
2. Murmann, B.: A/D converter trends: Power dissipation, scaling and digitally assisted architectures. In: 2008 IEEE Custom Integrated Circuits Conference, pp. 105–112 (2008). doi:10.1109/CICC.2008.4672032
3. Murray-Bruce, J., Dragotti, P.L.: Estimating localized sources of diffusion fields using spatiotemporal sensor measurements. IEEE Transactions on Signal Processing **63**(12), 3018–3031 (2015). doi:10.1109/TSP.2015.2419187
4. Murray-Bruce, J., Dragotti, P.L.: A sampling framework for solving physics-driven inverse source problems. IEEE Transactions on Signal Processing **65**(24), 6365–6380 (2017). doi:10.1109/TSP.2017.2742983
5. Nyquist, H.: Certain topics in telegraph transmission theory. Trans. Am. Inst. Electr. Eng. **47**(2), 617–644 (1928). doi:10.1109/T-AIEE.1928.5055024
6. Shannon, C.E.: Communication in the presence of noise. Proc. IRE **37**, 10–21 (1949)
7. Schmidt, R.: Multiple emitter location and signal parameter estimation. IEEE Transactions on Antennas and Propagation **34**(3), 276–280 (1986). doi:10.1109/TAP.1986.1143830
8. Roy, R., Kailath, T.: ESPRIT-estimation of signal parameters via rotational invariance techniques. IEEE Trans. Acoust. Speech Signal Process. **37**(7), 984–995 (1989). doi:10.1109/29.32276
9. Hua, Y., Sarkar, T.K.: Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. IEEE Trans. Acoust. Speech Signal Process. **38**, 814–824 (1990). doi:10.1109/29.56027
10. de Prony, R.: Essai expérimental et analytique sur les lois de la dilatabilité des fluides élastiques et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de l'alkool, à différentes températures. J. Ec. Poly. **1**, 24–76 (1795)
11. Batenkov, D.: Stability and super-resolution of generalized spike recovery. Applied and Computational Harmonic Analysis, in press. doi:10.1016/j.acha.2016.09.004
12. Hassanieh, H., Indyk, P., Katabi, D., Price, E.: Simple and practical algorithm for sparse Fourier transform. In: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1183–1194. SIAM. http://dl.acm.org/citation.cfm?id=2095209 Accessed 2014-09-24
13. Mishali, M., Eldar, Y.C., Dounaevsky, O., Shoshan, E.: Xampling: Analog to digital at sub-Nyquist rates. IET Circuits, Devices Systems **5**(1), 8–20 (2011). doi:10.1049/iet-cds.2010.0147
14. Christlieb, A., Lawlor, D., Wang, Y.: A multiscale sub-linear time Fourier algorithm for noisy data. Appl. Comput. Harmon. Anal. **40**(3), 553–574 (2016). doi:10.1016/j.acha.2015.04.002
15. Iwen, M.A.: Combinatorial sublinear-time Fourier algorithms. Found. Comput. Math. **10**(3), 303–338 (2010). doi:10.1007/s10208-009-9057-1
16. Vaidyanathan, P.P., Pal, P.: Sparse sensing with co-prime samplers and arrays. IEEE Transactions on Signal Processing **59**(2), 573–586 (2011). doi:10.1109/TSP.2010.2089682
17. Abari, O., Hamed, E., Hassanieh, H., Agarwal, A., Katabi, D., Chandrakasan, A.P., Stojanovic, V.: A 0.75-million-point Fourier-transform chip for frequency-sparse signals. In: 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), pp. 458–459 (2014). doi:10.1109/ISSCC.2014.6757512
18. Hassanieh, H., Shi, L., Abari, O., Hamed, E., Katabi, D.: GHz-wide sensing and decoding using the sparse Fourier transform. In: IEEE INFOCOM 2014 - IEEE Conference on Computer Communications, pp. 2256–2264 (2014). doi:10.1109/INFOCOM.2014.6848169
19. Potts, D., Tasche, M., Volkmer, T.: Efficient spectral estimation by MUSIC and ESPRIT with application to sparse FFT. Frontiers in Applied Mathematics and Statistics **2**, 1 (2016). doi:10.3389/fams.2016.00001
20. Cuyt, A., Lee, W.-s.: Smart Data Sampling and Data Reconstruction. US Patent 9,690,749. Patent US 61/611,899
21. Cuyt, A., Lee, W.-s.: Smart Data Sampling and Data Reconstruction. PCT/EP2012/066204. Patent PCT/EP2012/066204
22. Cuyt, A., Lee, W.-s.: An analog Chinese Remainder Theorem. Technical report, Universiteit Antwerpen (2017)
23. Cuyt, A., Lee, W.-s.: How to get high resolution results from sparse and coarsely sampled data. Technical report, Universiteit Antwerpen (2017)
24. Thompson, J.K., Tree, D.R.: Leakage error in fast Fourier analysis. Journal of Sound and Vibration **71**(4), 531–544 (1980). doi:10.1016/0022-460X(80)90725-7
25. Cuyt, A., Lee, W.-s.: Sparse interpolation and rational approximation. In: Hardin, D., Lubinsky, D., Simanek, B. (eds.) Contemporary Mathematics, vol. 661, pp. 229–242. American Mathematical Society, Providence, RI (2016). doi:10.1090/conm/661/13284.
26. Bessis, D.: Padé approximations in noise filtering. J. Comput. Appl. Math. **66**, 85–88 (1996). doi:10.1016/0377-0427(95)00177-8
27. Gilewicz, J., Pindor, M.: Padé approximants and noise: rational functions. J. Comput. Appl. Math. **105**, 285–297 (1999). doi:10.1016/S0377-0427(99)00041-2
28. Cuyt, A., Tsai, M., Verhoye, M., Lee, W.-s.: Faint and clustered components in exponential analysis. Applied Mathematics and Computation (2017). To appear
29. Cuyt, A., Lee, W.-s.: Multivariate exponential analysis from the minimal number of samples. Adv. Comp. Math. (2017). To appear

Figures

Figure 1 Schematic diagram of how to fix possible aliasing in a DFT computed from sub-Nyquist sampling, where both $u > 1$ and $s > 1$.

Figure 2 For the signals in Table 1: ${}_u^{ms}X, m = 0, \ldots, 11$ (left) and output of the proposed hybrid method and the standard DFT (right).

Figure 3 For the signals in Table 1: amplitudes retrieved by the hrybrid method (top) and number of colliding components indicated by the SVD (bottom).

Tables

Figure 4 For the signal in Table 2: output of the hybrid scheme for increasing noise levels (right to left) and growing M values (top to bottom) compared to a noiseless DFT (far left).

Table 1

Signal 1

μ_i	α_i
125	1

Signal 2

μ_i		α_i	
125	165	1	$e^{\mathrm{i}\pi/3}$

Signal 3

μ_i			α_i		
125	165	245	1	$e^{\mathrm{i}\pi/3}$	$e^{\mathrm{i}\pi/4}$

Table 2

μ_i	-6155.93	-4055.92	-2005.99	-47.97
	242.01	642.05	7310.07	8110.05
α_i	random in $[1,3]+\mathrm{i}[1,3]$			