Super-Resolution From Unregistered and Totally Aliased Signals Using Subspace Methods

Patrick Vandewalle, *Associate Member, IEEE*, Luciano Sbaiz, *Senior Member, IEEE*, Joos Vandewalle, *Fellow, IEEE*, and Martin Vetterli, *Fellow, IEEE*

Abstract—In many applications, the sampling frequency is limited by the physical characteristics of the components: the pixel pitch, the rate of the analog-to-digital (A/D) converter, etc. A low-pass filter is usually applied before the sampling operation to avoid aliasing. However, when multiple copies are available, it is possible to use the information that is inherently present in the aliasing to reconstruct a higher resolution signal. If the different copies have unknown relative offsets, this is a nonlinear problem in the offsets and the signal coefficients. They are not easily separable in the set of equations describing the super-resolution problem. Thus, we perform joint registration and reconstruction from multiple unregistered sets of samples. We give a mathematical formulation for the problem when there are M sets of N samples of a signal that is described by L expansion coefficients. We prove that the solution of the registration and reconstruction problem is generically unique if $MN \geq L + M - 1$. We describe two subspace-based methods to compute this solution. Their complexity is analyzed, and some heuristic methods are proposed. Finally, some numerical simulation results on one- and two-dimensional signals are given to show the performance of these methods.

Index Terms—Aliasing, sampling, offset estimation, shift estimation, image registration, super-resolution.

P. Vandewalle and L. Sbaiz are with the School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mail: Patrick.Vandewalle@epfl.ch; Luciano.Sbaiz@epfl.ch).

J. Vandewalle is with the Department of Electrical Engineering (ESAT), Katholieke Universiteit Leuven, BE-3001 Leuven, Belgium (e-mail: Joos.Vandewalle@esat.kuleuven.be).

M. Vetterli is with the School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland, and also with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 USA (e-mail: Martin.Vetterli@epfl.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2007.894257

I. INTRODUCTION

ALIASING, caused by undersampling of a signal, is generally considered as a nuisance and has to be avoided. To this effect, an antialiasing low-pass filter is often placed in front of the actual sampling operation, such that the sampled signal is not aliased. However, in super-resolution signal reconstruction, aliased components contain valuable high-frequency information that can be used to recover a higher resolution reconstruction. In this paper, we study methods to reconstruct a signal, including its high-frequency information, from multiple aliased sampled signals with relative offsets. The offsets are unknown and need to be computed first. For the estimation of those offsets, we explicitly use the information available in the aliased part of the spectrum and do not need extra measurements.

Applications can be found in various domains, such as super-resolution imaging. A camera is held manually while taking multiple images of the same scene. In satellite imaging, a satellite acquires images of approximately the same region at different moments. For one-dimensional signals, we find a similar setup in high rate analog-to-digital (A/D) converters. A single high-rate A/D converter is replaced by multiple lower rate converters that operate with small offsets. These low-rate converters are hard to synchronize precisely and can therefore be considered as converters with unknown offsets.

Aliasing is most often formulated in frequency domain. It is seen as the replication of frequencies above half the sampling frequency in the base spectrum of the sampled signal. In this way, a frequency above half the sampling frequency is mapped onto a frequency below this limit after sampling, and the two cannot be distinguished from their samples. In images, aliasing typically appears as artificial low-frequency patterns or jagged edges. Examples are given in Fig. 1.

The aliasing in the sampled signals has two effects: it creates artificial low-frequency patterns and makes image registration more difficult, but it also allows reconstructing an image with higher resolving power. Suppose that an antialiasing lowpass filter is applied before sampling: the sampled signals are then aliasing-free [see Fig. 2(a)]. It is much easier to register such a set of images, as all the images contain the same low-frequency information. However, as they all contain the same information, no additional resolving power can be obtained by combining the different images. In such a case, a larger image can as well be reconstructed by interpolating a single image [Fig. 2(b)]. However, if the signals are not lowpass filtered prior to sampling [Fig. 2(c)], multiple images can be used to recover the high frequency information that was aliased. The reconstructed image has a higher resolving power than any of the original images

Fig. 1. Examples of aliasing in images. Artificial low-frequency patterns (a) and jagged edges (b) appear due to undersampling.


Fig. 2. Value of aliasing. If an antialiasing lowpass filter is applied to images prior to sampling (a), no additional high-frequency information can be recovered by combining multiple images. One can as well interpolate a single image (b). If no filter is applied, the captured images are aliased (c), and high-frequency information can be recovered by combining multiple images (d). However, new image registration methods are required to obtain sufficiently high precision.

Fig. 3. Classification of sampling methods. Sampling methods can be divided into uniform and nonuniform methods. The nonuniform sampling methods can be subdivided depending on whether the locations are known and whether the samples are grouped in uniform sets with only unknown offsets. In this paper, we discuss multichannel sampling methods with unknown offsets.

[Fig. 2(d)]. The drawback is that more sophisticated image registration methods are required to achieve subpixel motion estimates. Two such methods will be presented in this paper. To our knowledge, no other registration methods exist that deal with such totally aliased signals (as opposed to partially aliased signals, which have an aliasing-free part in their spectrum, and are addressed for example in [4]).

Sampling methods can be classified into different categories, according to the type of sampling that is performed (see Fig. 3). A first important distinction is between uniform and nonuniform sampling methods. With uniform sampling, a signal is sampled periodically at constant time intervals. This is the standard setup, which is described by the Shannon–Nyquist sampling theorem [5]. An overview of some more recent results is given by Unser [6]. Among the nonuniform sampling methods, a distinction needs to be made between methods using known sampling instants [7], [8], and other methods where the sampling locations are unknown. If the sampling locations are unknown and completely arbitrary, the problem cannot be solved. This can be shown using a simple counting argument. Assume that the signal to be reconstructed has L unknown parameters. For every additional sample, there is also an additional unknown (its location). Therefore, the number of unknowns is always larger than the number of measurements, and this problem is unsolvable. However, for discrete time (or space) signals, where the sampling locations can only take a finite number of values, a combinatorial solution can be found, as described by Marziliano and Vetterli [9].

An important subset of the nonuniform sampling methods is formed by multichannel sampling methods. In these methods, multiple sets of uniformly spaced samples are measured. Each of the sets of samples is uniform, but the different sets have a

nonuniform offset. Papoulis [10] described a solution for multichannel sampling with known sampling locations. He showed that a bandlimited signal can be perfectly reconstructed from M sets of samples that are sampled at $1/M$ the Nyquist sampling rate. This result was extended by Unser and Zerubia in their generalized sampling approach [11], [12]. The problem with multiple sets and unknown sampling locations was solved for discrete-time signals by Marziliano and Vetterli [9]. In this paper, we will study the continuous-time case: multichannel sampling with unknown, real-valued offsets.

A typical application of such a setup is super-resolution imaging. The goal in super-resolution imaging is to reconstruct a high resolution image from a set of multiple images that are taken from almost the same point of view. Most of the times, the images are assumed to differ by small, subpixel shifts. This problem was first described by Tsai and Huang [13] in 1984. A good overview of existing super-resolution methods is given by Borman and Stevenson [14], Farsiu *et al.* [15], and in the special issues on super-resolution imaging in the *IEEE Signal Processing Magazine* [16] and in the EURASIP *Journal on Applied Signal Processing* [17]. In most cases, super-resolution image reconstruction is separated into two main tasks, an image registration task followed by a reconstruction task, and these are addressed separately [18]–[21]. However, a joint consideration of the two problems offers opportunities for a better, global solution. Such an approach will be used here and is also presented in a frequency-domain setup by Tsai and Huang [13] and in a maximum *a posteriori* framework by Hardie *et al.* [22].

Different methods exist to compute a high-resolution image from a set of low-resolution images. They mainly differ by the amount of prior information that is assumed to be available about the images. The most generally applicable methods only require that the images belong to a broad signal class, such as bandlimited signals [13], or signals with finite rate of innovation [23]. Other methods, such as the hallucination method by Baker and Kanade [24], use more specific knowledge about the image type (faces, printed text, etc.) to reconstruct a high-resolution image. In general, there is a tradeoff between the amount of prior information and the number of images required to produce good results. The methods described in this paper belong to the first category and only use basic assumptions about the signal class. Most other super-resolution methods either consider the aliasing as part of the noise or, at most, include it in their model for the reconstruction part. As opposed to these methods, we will include the aliasing into our signal model and use its information for both registration and reconstruction. In this way, our algorithm can be used on signals with an arbitrary undersampling factor. The reconstruction part of our algorithm uses a frequency domain description of the aliased signal, which is the same as the setup presented earlier by Tsai and Huang [13] and Kaltenbacher and Hardie [19]. Some recent spatial domain methods are proposed by Farsiu *et al.* [20] and Pham *et al.* [21]. For an algorithmic solution of the equations described in this paper, one can use a special case of the variable projection method for separable nonlinear least-squares problems [25]. Such a method is similar to the algorithm we present in Section V.

This paper is structured as follows. First, a mathematical description of the problem is given in Section II. Next, the uniqueness of the solution to this problem is discussed in Section III. Two subspace-based solution methods are described in Sections IV and V, respectively. We discuss the advantages and disadvantages of the different methods in Section VI, and we also present heuristic solution methods. In Section VII, we derive the computational complexity of the different algorithms. Results are presented in Section VIII, and finally some conclusions are drawn in Section IX. All the results and figures presented in this paper are reproducible [26]. The Matlab code and data to reproduce the results are available online at http://lcavwww.epfl.ch/reproducible_research/VandewalleSVV06.

II. MATHEMATICAL DESCRIPTION

We aim to reconstruct a continuous-time signal from multiple uniform sets of samples with unknown offsets. The sampling frequency of each set is too low to allow direct reconstruction, but reconstruction is possible from the joint set of all samples together. In this section, we formulate the mathematical framework for this problem and specify notations. Assume we have a periodic bandlimited signal[1] with period 1. Such a signal can be expressed as

$$f(t) = \sum_{l=-K}^{K} \alpha_l e^{j2\pi lt} \qquad (1)$$

where α_l are the $L = 2K + 1$ Fourier coefficients of $f(t)$. Assume now that we sample $f(t)$ uniformly at a rate N over the signal period $[0, 1]$. This results in a first set of N samples taken at times

$$\mathcal{T}_0 = \begin{pmatrix} 0 & \dfrac{1}{N} & \dfrac{2}{N} & \cdots & \dfrac{N-1}{N} \end{pmatrix}. \qquad (2)$$

We can write these samples as

$$y_0(n) = f\left(\frac{n}{N}\right) = \sum_{l=-K}^{K} \alpha_l W^{ln} \qquad (3)$$

with $W = e^{j2\pi/N}$. We can write this more compactly using vectors and matrices as (4), shown at the bottom of the next page. In this equation, $\mathbf{y}_0$ is the $N \times 1$ sample vector, and $\boldsymbol{\alpha}$ is the $L \times 1$ vector of unknown Fourier coefficients. The matrix $\mathbf{F}^*$ is an $N \times L$ inverse discrete Fourier transform (IDFT) matrix, where the notation $\mathbf{F}^*$ is used for the Hermitian transpose of the forward discrete Fourier transform matrix $\mathbf{F}$. Due to the undersampling ($L > N$), some of the columns in $\mathbf{F}^*$ are repeated, and the set of (4) is thus underdetermined. We take another set of samples at times

$$\mathcal{T}_m = \begin{pmatrix} t_m & \dfrac{1}{N} + t_m & \dfrac{2}{N} + t_m & \cdots & \dfrac{N-1}{N} + t_m \end{pmatrix} \ (5)$$

These are the sampling times of the first set, shifted by an unknown offset t_m ($t_m \in \mathcal{R}$). As the offsets are relative to the first

[1]Note that most of the theory developed in this paper can be extended to signals in an arbitrary finite-dimensional Hilbert space. See [27] for details. Also, in most practical cases, the signals to reconstruct are real. Our algorithms work both for real and complex signals.

set of samples, we have $t_0 = 0$. We take M such sets of samples $\mathbf{y}_m$. The samples from the mth set can be written as

$$
\begin{aligned}
y_m(n) = f\left(\frac{n}{N} + t_m\right) &= \sum_{l=-K}^{K} \alpha_l e^{j2\pi l(n/N + t_m)} \\
&= \sum_{l=-K}^{K} \alpha_l W^{ln} W_{t_m}^l
\end{aligned}
\tag{6}
$$

with $W_{t_m} = e^{j2\pi t_m}$.[2] In matrix notation, this can be expressed as

$$
\mathbf{y}_m = \mathbf{F}^* \mathbf{D}_{t_m} \boldsymbol{\alpha} = \boldsymbol{\Phi}_{t_m} \boldsymbol{\alpha}
\tag{7}
$$

with $\mathbf{F}^*$ the $N \times L$ IDFT matrix defined above, $\mathbf{D}_{t_m}$ an $L \times L$ diagonal matrix with elements $\mathbf{D}_{t_m}(l, l) = W_{t_m}^l$ ($-K \leq l \leq K$), and $\boldsymbol{\Phi}_{t_m} = \mathbf{F}^* \mathbf{D}_{t_m}$.

Now, all the sets of samples $\mathbf{y}_m$ are combined into a single vector $\mathbf{y}$ and, similarly, the basis matrices $\boldsymbol{\Phi}_{t_m}$ are combined into $\boldsymbol{\Phi}_\mathbf{t}$, with $\mathbf{t} = (t_0 \; t_1 \; \cdots \; t_{M-1})$ denoting the offset vector. This can be written as

$$
\mathbf{y} = \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_{M-1} \end{pmatrix} = \begin{pmatrix} \mathbf{F}^* \\ \mathbf{F}^* \mathbf{D}_{t_1} \\ \vdots \\ \mathbf{F}^* \mathbf{D}_{t_{M-1}} \end{pmatrix} \boldsymbol{\alpha} = \begin{pmatrix} \boldsymbol{\Phi}_0 \\ \boldsymbol{\Phi}_{t_1} \\ \vdots \\ \boldsymbol{\Phi}_{t_{M-1}} \end{pmatrix} \boldsymbol{\alpha} = \boldsymbol{\Phi}_\mathbf{t} \boldsymbol{\alpha}.
\tag{8}
$$

The matrix $\boldsymbol{\Phi}_\mathbf{t}$ has size $MN \times L$. Assuming that the total number of samples is larger than or equal to the number of Fourier coefficients, or $MN \geq L$, this set of equations is in general well determined or overdetermined if $\mathbf{t}$ is known. If, in addition

$$
MN \geq L + M - 1
\tag{9}
$$

the number of equations is also larger than or equal to the number of unknowns (L expansion coefficients and $M - 1$ offsets), and it should be possible to remove the uncertainty of the unknown offsets.[3] As we will show in the next sections, these additional equations allow us in general to compute the unknown offsets. Note that the challenging part of the problem

[2]The W defined in (3) can therefore also be written in this notation as $W_{1/N}$. For compactness, we keep W as a separate symbol.

[3]This supposes of course that the MN equations are independent, which is generally the case if offsets with different non-integer parts are used between the sets of samples. Remark that in practical situations, numerical instability might occur if the offsets are chosen very close together [12].


Fig. 4. Illustration of the different variables with $M = 2$ and a Fourier basis. (a) Time-domain representation of the signal $f(t)$ and its sets of samples y_0 (—) and y_1 (— —). (b) Frequency-domain representation of the absolute values of the signal spectrum (—) and its aliased copies after sampling (— —).

is that it is nonlinear in the offsets and that $\boldsymbol{\alpha}$ and $\mathbf{t}$ have to be found simultaneously. In summary, the most important variables in this reconstruction problem are listed here (see also Fig. 4):

N the number of samples in each set $\mathbf{y}_m$;
$\mathbf{y}_m$ the length N vector of the mth set of samples;
L the number of unknown Fourier coefficients ($L = 2K + 1$);
$\boldsymbol{\alpha}$ the length L vector of the expansion coefficients α_l to be reconstructed;
M the number of sets of samples;
$\mathbf{t}$ the length M vector of the offsets t_m between the different sets of samples.

The unknown variables are the Fourier coefficients $\boldsymbol{\alpha}$ and the offsets $\mathbf{t}$. We assume that all the other variables are known. This is evident for the sets of samples $\mathbf{y}_m$, the number of samples per set N and the number of sets M, as they form the input of the problem. We will also require that the number of Fourier coefficients L, or at least an estimate for L, is available.

The DFT of a set of samples $\mathbf{y}_m$ can be written as

$$
\mathbf{Y}_m = \frac{1}{N} \mathbf{F}_N \mathbf{y}_m = \frac{1}{N} \mathbf{F}_N \mathbf{F}^* \mathbf{D}_{t_m} \boldsymbol{\alpha}
\tag{10}
$$

where $\mathbf{F}_N$ is a square $N \times N$ DFT matrix, and $\mathbf{F}^*$ is the $N \times L$ matrix defined in (3). The resulting vector $\mathbf{Y}_m$ has length N

$$
\mathbf{y}_0 = \begin{pmatrix} 1 & \cdots & 1 & 1 & 1 & \cdots & 1 \\ W^{-K} & \cdots & W^{-1} & 1 & W & \cdots & W^K \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ W^{-(N-1)K} & \cdots & W^{-(N-1)} & 1 & W^{N-1} & \cdots & W^{(N-1)K} \end{pmatrix} \begin{pmatrix} \alpha_{-K} \\ \vdots \\ \alpha_{-1} \\ \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_K \end{pmatrix} = \mathbf{F}^* \boldsymbol{\alpha}.
\tag{4}
$$

Fig. 5. Example of the problem setup. Bandlimited function with two sets of two samples and unknown offset between the sets.

and is an aliased and phase shifted version of $\boldsymbol{\alpha}$. If we take, for example, $L = 3N$, we get

$$\mathbf{Y}_m = \frac{1}{N}\mathbf{F}_N\mathbf{F}^*\mathbf{D}_{t_m}\boldsymbol{\alpha} = \frac{1}{N}\mathbf{F}_N(\mathbf{F}_N^* \quad \mathbf{F}_N^* \quad \mathbf{F}_N^*)\mathbf{D}_{t_m}\boldsymbol{\alpha}$$

$$= (\mathbf{I} \quad \mathbf{I} \quad \mathbf{I})\mathbf{D}_{t_m}\boldsymbol{\alpha} = \sum_{i=-1}^{1} W_{t_m}^{iN}\mathbf{D}'_{t_m}\boldsymbol{\alpha}_i \quad (11)$$

where $\mathbf{D}'_{t_m}$ is the $N \times N$ central part of the $L \times L$ matrix $\mathbf{D}_{t_m}$, and $\boldsymbol{\alpha}_i$ is the ith block of length N from the vector $\boldsymbol{\alpha}$. In general, if L is not a multiple of N, we can still perform the same decomposition by adding zeros to $\boldsymbol{\alpha}$ up to the next multiple of N.

Before tackling the problem in general, let us first analyze a very simple and intuitive example.

Example 2.1: Take $f(t)$ a bandlimited function (see Fig. 5), with $M = 2, N = 2$, and $L = 3$. Using (8), we can write

$$\mathbf{y} = \begin{pmatrix} y_0(0) \\ y_0(1) \\ y_1(0) \\ y_1(1) \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & -1 \\ \hline W_{t_1}^{-1} & 1 & W_{t_1} \\ -W_{t_1}^{-1} & 1 & -W_{t_1} \end{pmatrix} \begin{pmatrix} \alpha_{-1} \\ \alpha_0 \\ \alpha_1 \end{pmatrix} = \boldsymbol{\Phi}_\mathbf{t}\boldsymbol{\alpha}. \quad (12)$$

The unknown offset in the matrix $\boldsymbol{\Phi}_\mathbf{t}$ is multiplied with the expansion coefficients in $\boldsymbol{\alpha}$ and makes the problem nonlinear. As described in (7), we can write $\boldsymbol{\Phi}_{t_1}$ (the lower part of $\boldsymbol{\Phi}_\mathbf{t}$) as an IDFT matrix multiplied with a diagonal matrix

$$\boldsymbol{\Phi}_{t_1} = \begin{pmatrix} W_{t_1}^{-1} & 1 & W_{t_1} \\ -W_{t_1}^{-1} & 1 & -W_{t_1} \end{pmatrix}$$

$$= \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & -1 \end{pmatrix} \begin{pmatrix} W_{t_1}^{-1} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & W_{t_1} \end{pmatrix} = \mathbf{F}^*\mathbf{D}_{t_1}. \quad (13)$$

$\square$

From the above example, we see that the unknown offsets and signal coefficients are mixed up, which makes the problem

challenging. Once the offsets are known, (8) is reduced to a set of linear equations in the unknown signal coefficients. This can be solved straightforwardly using a least-squares method. From this description, we clearly see that the hardest part of the reconstruction problem is the estimation of the offsets. In the next sections, we will concentrate on this problem, and present methods to compute these offsets based on properties of subspaces.

III. UNIQUENESS OF THE SOLUTION

In this section, we discuss the existence and uniqueness of a solution to the super-resolution problem with unregistered sets of samples. It follows straightforwardly from the description above that a solution exists (at least in the ideal, noiseless case). The uniqueness of such a solution is less trivial.

Let us start with an intuitive statement about uniqueness. A unique solution exists if there is a unique mapping from each set of samples to a single space generated by the columns of $\boldsymbol{\Phi}_\mathbf{t}$. In other words, the intersection of the spaces generated for two different offset vectors $\mathbf{t}$ only contains the trivial zero vector, except for possibly some degenerate cases. This argument follows the one made by Marziliano for discrete sampling with unknown locations [28]. In general, the intersection of two such spaces has a lower dimension than the spaces themselves and thus, the set of sample vectors for which a unique solution exists is dense in the signal space. We shall therefore ignore the degenerate cases.

As we discussed in Section II, the matrix $\boldsymbol{\Phi}_{t_m}$ can be written as the product of an $N \times L$ IDFT matrix $\mathbf{F}^*$ with an $L \times L$ diagonal matrix $\mathbf{D}_{t_m}$. For aliased signals, some of the columns in $\mathbf{F}^*$ are repeated, and the corresponding columns of $\boldsymbol{\Phi}_{t_m}$ only differ by their multiplication factor $W_{t_m}^l$. These columns correspond to overlapping frequencies in the sampled spectrum.

If all the sets of samples are considered together, the lth column $\boldsymbol{\phi}_\mathbf{t}^l$ of the matrix $\boldsymbol{\Phi}_\mathbf{t}$ is a repetition of M times the same basis vector $\mathbf{f}_l$, the lth column of $\mathbf{F}^*$, multiplied by the different factors $W_{t_m}^l$:

$$\left(\boldsymbol{\phi}_\mathbf{t}^l\right)^T = \left(\mathbf{f}_l^T \quad W_{t_1}^l\mathbf{f}_l^T \quad \cdots \quad W_{t_{M-1}}^l\mathbf{f}_l^T\right). \quad (14)$$

Example 3.1: For example, if we take $M = 2, L = 5$, and $N = 4$, there are two overlapping frequency components (see Fig. 6). We have

$$\boldsymbol{\Phi}_\mathbf{t} = \begin{pmatrix} \boldsymbol{\Phi}_0 \\ \boldsymbol{\Phi}_{t_1} \end{pmatrix} = \begin{pmatrix} \mathbf{F}^* \\ \mathbf{F}^*\mathbf{D}_{t_1} \end{pmatrix}$$

$$= \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ W^2 & W^3 & 1 & W & W^2 \\ 1 & W^2 & 1 & W^2 & 1 \\ W^2 & W & 1 & W^3 & W^2 \\ \hline W_{t_1}^{-2} & W_{t_1}^{-1} & 1 & W_{t_1} & W_{t_1}^2 \\ W_{t_1}^{-2}W^2 & W_{t_1}^{-1}W^3 & 1 & W_{t_1}W & W_{t_1}^2W^2 \\ W_{t_1}^{-2} & W_{t_1}^{-1}W^2 & 1 & W_{t_1}W^2 & W_{t_1}^2 \\ W_{t_1}^{-2}W^2 & W_{t_1}^{-1}W & 1 & W_{t_1}W^3 & W_{t_1}^2W^2 \end{pmatrix}, \quad (15)$$

Fig. 6.   Signal spectrum before (a) and after sampling (b), for $L = 5$ and $N = 4$. The base spectrum (—) and aliased spectrum (−−) overlap for the first and last spectral component, corresponding to the first and fifth column in (15).

with $W = e^{j2\pi/4}$. The first and the fifth column of $\mathbf{F}^*$ are equal, such that the corresponding columns of $\mathbf{\Phi_t}$ only differ by their factors $W_{t_1}^{-2}$ and $W_{t_1}^2$ from $\mathbf{D}_{t_1}$. $\square$

Because the Fourier basis vectors $\mathbf{f}_l$ of the $N \times L$ matrix $\mathbf{F}^*$ are orthogonal (note that we use $W = e^{j2\pi/N}$ in this matrix, and not $e^{j2\pi/L}$), the vectors $\boldsymbol{\phi}_\mathbf{t}^l$ are also orthogonal to each other for any set of offset values $\mathbf{t}$. Only the vectors corresponding to overlapping spectrum coefficients (like the first and fifth column in Example 3.1) are not orthogonal, because they are composed by the same Fourier vector $\mathbf{f}_l$ but have different coefficients. Each vector describes a trajectory in $\mathcal{C}^{MN}$ for varying values of $\mathbf{t}$. This trajectory is defined more precisely in the following theorem.

Theorem 3.1: For varying $\mathbf{t}$, any vector $\boldsymbol{\phi}_\mathbf{t}^l$ (with $l \neq 0$) describes a trajectory in the M-dimensional subspace $\mathcal{V}_l$ of $\mathcal{C}^{MN}$:

$$\mathcal{V}_l = \mathrm{span}\{\mathbf{A} \otimes \mathbf{f}_l\}$$
$$\text{with } \mathbf{A} = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & -1 & & 1 \\ \vdots & & \ddots & \vdots \\ 1 & 1 & \cdots & -1 \end{pmatrix} \quad (16)$$

where $\mathbf{A}$ is an $M \times M$ matrix, and $\otimes$ represents the Kronecker product. The vectors $\boldsymbol{\phi}_\mathbf{t}^{l+iN}$ corresponding to overlapping spectrum coefficients belong to the same space $\mathcal{V}_l$, while other vectors $\boldsymbol{\phi}_\mathbf{t}^k$ belong to orthogonal subspaces $\mathcal{V}_k \perp \mathcal{V}_l$.

Proof: The trajectory of $\boldsymbol{\phi}_\mathbf{t}^l$ as a function of $\mathbf{t}$ is in $\mathcal{V}_l$ iff we can write any arbitrary $\boldsymbol{\phi}_\mathbf{t}^l$ as a linear combination of the columns of $\mathbf{A} \otimes \mathbf{f}_l$. From (14) and (16), we can write the vector $\boldsymbol{\phi}_\mathbf{t}^l$ as a linear combination of the columns of $\mathbf{A} \otimes \mathbf{f}_l$ by solving the set of linear equations

$$\begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & -1 & & 1 \\ 1 & & \ddots & \vdots \\ 1 & 1 & \cdots & -1 \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{M-1} \end{pmatrix} = \begin{pmatrix} 1 \\ W_{t_1}^l \\ \vdots \\ W_{t_{M-1}}^l \end{pmatrix}. \quad (17)$$

As this set of M equations is of full rank (because $\mathbf{A}$ is nonsingular), it always has a unique solution, and our vector is therefore part of $\mathcal{V}_l$. Vectors corresponding to overlapping Fourier coefficients are composed from the same Fourier basis vectors $\mathbf{f}_l$, and therefore belong to the same space $\mathcal{V}_l$.

The orthogonality between two subspaces $\mathcal{V}_k$ and $\mathcal{V}_l$ for $k \neq l$ can easily be seen. As each of the subspaces is generated by the columns of the matrices defined in (16), it is sufficient that we

prove that any arbitrary column of $\mathbf{A} \otimes \mathbf{f}_k$ is orthogonal to any column of $\mathbf{A} \otimes \mathbf{f}_l$. Denoting the ith column of $\mathbf{A}$ as $\mathbf{A}^i$, we can write the inner product between two such vectors as

$$\langle \mathbf{A}^i \otimes \mathbf{f}_k, \mathbf{A}^j \otimes \mathbf{f}_l \rangle = \sum_{n=0}^{M-1} (\pm 1)\langle \mathbf{f}_k, \mathbf{f}_l \rangle = 0 \quad \text{if } k \neq l. \quad (18)$$

This is valid because $\langle \mathbf{f}_k, \mathbf{f}_l \rangle = 0$ from the orthogonality of the Fourier basis.   ■

From this analysis, we can see that the problem can be considered in each of the N subspaces separately by projecting $\mathbf{y}$ onto the different M-dimensional subspaces. The projection of $\mathbf{y}$ onto the subspace $\mathcal{V}_l$ is called

$$\mathbf{y}^{(l)} = \mathbf{P}_{\mathcal{V}_l}\mathbf{y}. \quad (19)$$

An illustration for two overlapping vectors in a three-dimensional space is sketched in Fig. 7.

Assuming that $MN \geq L + M - 1$, at most M columns of $\mathbf{\Phi_t}$ belong to the same subspace $\mathcal{V}_l$. If there are M vectors $\boldsymbol{\phi}_\mathbf{t}^{l+iN}$ from $\mathbf{\Phi_t}$ in a particular subspace $\mathcal{V}_l$, these vectors form a basis for $\mathcal{V}_l$. As the vector $\mathbf{y}^{(l)}$ is also in $\mathcal{V}_l$, it belongs to the subspace spanned by these vectors $\boldsymbol{\phi}_\mathbf{t}^{l+iN}$ for any value of $\mathbf{t}$. However, for the subspaces containing less than M vectors $\boldsymbol{\phi}_\mathbf{t}^{l+iN}$, the span of these vectors does not cover the entire space $\mathcal{V}_l$. Hence, the projection of $\mathbf{y}$ onto $\mathcal{V}_l$ generally only belongs to $\mathrm{span}(\boldsymbol{\phi}_\mathbf{t}^{l+iN})_{0 \leq i < M-1}$ for a single offset vector $\mathbf{t}$. The offset vector $\mathbf{t}$ can then be uniquely determined from these subspaces, and the signal coefficients $\boldsymbol{\alpha}$ are thus also uniquely determined.

There are certain cases for which it is not sufficient to require $MN \geq L + M - 1$. The vectors in the space $\mathcal{V}_0$ (the columns $l = iN$) do not depend on $\mathbf{t}$, and therefore, no information about the offsets can be derived from these vectors. So if only the space $\mathcal{V}_0$ contains less than M vectors, it is not possible to derive the offsets from the sets of samples. This happens when the number of overlapping parts of the frequency spectrum is even (and therefore also the number of sets of samples M is even). In that case, we need $MN \geq L + M$.

Example 3.2: Looking back at Example 3.1, we can see that the first and fifth column of $\mathbf{\Phi_t}$ span the subspace $\mathcal{V}_2$, except in the degenerate case when $t_1 = 1$. The projection of $\mathbf{y}$ onto these two columns will therefore always be the same as its projection onto $\mathcal{V}_2$. No information about t_1 can be derived from this subspace. However, for $\mathcal{V}_{-1}$ and $\mathcal{V}_1$, there is only one column of $\mathbf{\Phi_t}$ in each of these subspaces, and the projections onto that column or onto the space $\mathcal{V}_l$ only coincide for the correct value of t_1.

Fig. 7. Illustration of the trajectory of the span of two columns ϕ_t^l (the planes in this drawing) in a three-dimensional space $\mathcal{V}_l$ for different offset values (0, $\mathbf{t}$ and $\mathbf{t}'$). Only for the correct offsets $\mathbf{t}$, the vector $\mathbf{y}^{(l)}$ belongs to the space spanned by the two columns.


Fig. 8. Difference between the sample vector projected onto the spaces $\mathcal{V}_n$ and the corresponding columns of $\mathbf{\Phi_t}$. The space $\mathcal{V}_2$ contains two vectors, and $\|\mathbf{y}^{(n)} - \hat{\mathbf{y}}^{(n)}\|^2$ is zero for any value of t_1. Similarly, $\mathcal{V}_0$ does not depend on t_1, and thus it gives no information about the offset either. Both $\mathcal{V}_{-1}$ and $\mathcal{V}_1$ clearly indicate the correct value of t_1.

The subspace $\mathcal{V}_0$ is a one-dimensional subspace that does not depend on t_1, and does therefore not give any information about t_1. The difference between the projection of the sample vector $\mathbf{y}$ onto $\mathcal{V}_n$ and its projection onto the corresponding columns of $\mathbf{\Phi_t}$, $\|\mathbf{y}^{(n)} - \hat{\mathbf{y}}^{(n)}\|^2$, is plotted for varying values of t_1 in Fig. 8. We need at least one higher dimensional subspace containing only 1 column from $\mathbf{\Phi_t}$ in order to be able to compute t_1. Therefore, we require that $2N \geq L + 2$, or $N \geq 4$. $\quad\square$

This can be summarized as follows. If $MN \geq L + M - 1$ for M odd, or $MN \geq L + M$ for M even, a bandlimited signal with L Fourier coefficients can be uniquely reconstructed from M uniform sets of N samples with unknown offsets.

IV. SOLUTION USING MATRIX RANK

In this section, we describe a first solution method for the problem of finding relative offsets described in Section II.

A. Method

As shown in (10), the discrete Fourier transform $\mathbf{Y}_m$ of the mth set of samples can be written as

$$\mathbf{Y}_m = \frac{1}{N}\mathbf{F}_N\mathbf{F}^*\mathbf{D}_{t_m}\boldsymbol{\alpha} \qquad (20)$$

with $\mathbf{F}_N$ a square $N \times N$ DFT matrix, and $\mathbf{F}^*$ an $N \times L$ IDFT matrix. If we extend $\boldsymbol{\alpha}$ to a length that is a multiple of N, we can split the Fourier coefficient vector $\boldsymbol{\alpha}$ in blocks $\boldsymbol{\alpha}_i$ of length N, and obtain (see also (11))

$$\mathbf{Y}_m = \mathbf{D}'_{t_m} \sum_{i=\lceil -(S-1)/2\rceil}^{\lceil (S-1)/2\rceil} W_{t_m}^{iN}\boldsymbol{\alpha}_i. \qquad (21)$$

The vectors $\boldsymbol{\alpha}_i$ represent the overlapping parts of the Fourier spectrum due to undersampling, and there are $S = \lceil L/N\rceil$ such parts. This means that for any set of samples, $\mathbf{D}'^{-1}_{t_m}\mathbf{Y}_m$ is a linear combination of the S parts of the Fourier spectrum $\boldsymbol{\alpha}_i$

$$\mathbf{D}'^{-1}_{t_m}\mathbf{Y}_m = \sum_{i=\lceil -(S-1)/2\rceil}^{\lceil (S-1)/2\rceil} W_{t_m}^{iN}\boldsymbol{\alpha}_i. \qquad (22)$$

In other words, each vector $\mathbf{D}'^{-1}_{t_m}\mathbf{Y}_m$ belongs to the S-dimensional subspace $\text{span}(\{\boldsymbol{\alpha}_i\}_{i=\lceil -(S-1)/2\rceil\ldots\lceil (S-1)/2\rceil})$. If we have more than S sets of samples ($M > S$), and $N \geq S$, the rank of the matrix containing all the sets of samples should therefore be S:

$$\text{rank}\left(\mathbf{Y}_\mathbf{t}^D\right) = S$$

with $\mathbf{Y}_\mathbf{t}^D = \begin{pmatrix} \mathbf{Y}_0 & \mathbf{D}'^{-1}_{t_1}\mathbf{Y}_1 & \cdots & \mathbf{D}'^{-1}_{t_{M-1}}\mathbf{Y}_{M-1} \end{pmatrix}. \quad (23)$

Fig. 9. Signal reconstruction algorithm using the matrix rank method from Section IV. The estimate for the offsets $\hat{\mathbf{t}}$ and the corresponding sample matrix $\mathbf{Y}_{\hat{\mathbf{t}}}^D$ are updated iteratively. Once the estimate is good enough, the signal parameters $\boldsymbol{\alpha}$ are reconstructed.

In general, if the estimated offsets $\hat{\mathbf{t}}$ do not have the correct values, rank($\mathbf{Y}_{\hat{\mathbf{t}}}^D$) $> S$. For most offset values $\hat{\mathbf{t}}$, the matrix $\mathbf{Y}_{\hat{\mathbf{t}}}^D$ will be of full rank M (and $M > S$). The correct values of the offsets $\mathbf{t}$ can then be found as the values for which the rank of the matrix becomes S. A schematic description of the complete reconstruction algorithm built on this method is given in Fig. 9.

B. Discussion

As there are S overlapping parts in the sampled Fourier spectrum, we need at least $M = S + 1 = \lceil L/N \rceil + 1$ sets of samples to find the offsets using this method. If more sets of samples are available, we still have rank($\mathbf{Y}_{\hat{\mathbf{t}}}^D$) $= S$, and this adds some robustness to both the offset estimation and the signal reconstruction. However, each new set of samples also adds a new unknown offset, thereby increasing the complexity of the estimation.

Remark that the operation performed by multiplying a sample vector $\mathbf{Y}_m$ with a matrix $\mathbf{D}'^{-1}_{t_m}$ does not change its norm $\|\mathbf{Y}_m\|_2^2$. It merely performs a rotation to align the sample vector into the S-dimensional subspace.

The computation of the rank of a matrix has a rather 'binary' outcome: either it is S, or it has an integer value larger than S. Hence, if the measurements are noisy, this test is very likely to fail even for the correct values of $\mathbf{t}$. It is therefore much better to evaluate the $S + 1$th singular value of the matrix $\mathbf{Y}_{\hat{\mathbf{t}}}^D$, or the determinant of the square matrix $\mathbf{Y}_{\hat{\mathbf{t}}}^{D*}\mathbf{Y}_{\hat{\mathbf{t}}}^D$. This can also give an indication about the quality of the current approximation. While the determinant requires less computations, it is also numerically less stable than a singular value decomposition. The $S + 1$th singular value could be computed using the inverse power method described by Strang [29]. This is an iterative method to approximate the smallest singular value of a matrix. It requires the solution of a linear $M \times M$ system of equations in each iteration. The number of iterations needed will generally be very small, because the $S + 1$th singular value is typically much smaller than the first S singular values.

The objective function based on (23) can thus be written as

$$\min_{\hat{\mathbf{t}}} \sigma_{S+1}\left(\mathbf{Y}_{\hat{\mathbf{t}}}^D\right) \tag{24}$$

where the operation $\sigma_{S+1}(\mathbf{A})$ stands for computing the $S + 1$th singular value of the matrix $\mathbf{A}$.

Example 4.1: Let us consider a bandlimited function with $L = 81$ unknown Fourier coefficients. It is sampled with two sets of 90 samples, with offsets $\mathbf{t} = (0\ \ 0.6)$. The objective


Fig. 10. Example of the objective function in (24). A signal with 81 unknown Fourier coefficients is sampled with two aliased sets of 90 samples. The exact offset is $t_1 = 0.6$. Next to the global minimum, it also contains many local minima.

function from (24) is shown as a function of t_1 in Fig. 10.[4] Similarly, in a more realistic setup, we sample the same function with $M = 3$ (aliased) sets of $N = 41$ samples and offsets $\mathbf{t} = (0\ \ 0.2\ \ 0.6)$. The objective function is shown in Fig. 11(a) as a function of the offsets t_1 and t_2. For this signal, a slice of the objective function for $t_1 = 0.2$ is shown in Fig. 11(b). This figure shows the three singular values as a function of t_2, showing the clear minimum of σ_3 for $t = 0.6$, while the other singular values are much larger and remain almost constant.

In the above examples, the objective function (24) is very "flat," except for a small region around the optimal offset values. At an arbitrary value away from the correct solution, it is therefore not possible to predict what would be a better approximation. Standard minimization algorithms such as gradient descent will generally not converge to the right solution for such functions. However, from the example with three sets of samples, we can see that the objective function is not completely arbitrary. Horizontal, vertical, and diagonal lines appear. They correspond to relative alignments of two out of the three sets of samples. It is therefore interesting to use this information in the search for the minimum. Such a heuristic algorithm will be described in Section VI.

---

[4]Note that this example is not very interesting from a practical point of view, as there is actually no aliasing in the sampled signals ($N > L$). However, this is the setup that is required for this method if we have only two sets of samples ($M > S$).

the correctness of the offset vector by verifying if the projection $\hat{\mathbf{y}}$ of the sample vector $\mathbf{y}$ onto span($\mathbf{\Phi}_{\hat{\mathbf{t}}}$) gives the sample vector $\mathbf{y}$ again. Or mathematically

$$\begin{cases} \hat{\mathbf{y}} = \mathbf{P}_{\mathbf{\Phi}_{\hat{\mathbf{t}}}} \mathbf{y} = \mathbf{\Phi}_{\hat{\mathbf{t}}} (\mathbf{\Phi}_{\hat{\mathbf{t}}}^* \mathbf{\Phi}_{\hat{\mathbf{t}}})^{-1} \mathbf{\Phi}_{\hat{\mathbf{t}}}^* \mathbf{y} = \mathbf{y}, & \text{for } \hat{\mathbf{t}} = \mathbf{t} \\ \hat{\mathbf{y}} = \mathbf{P}_{\mathbf{\Phi}_{\hat{\mathbf{t}}}} \mathbf{y} \neq \mathbf{y}, & \text{for } \hat{\mathbf{t}} \neq \mathbf{t}. \end{cases} \quad (25)$$

Therefore, this can be used to build an optimization problem for finding the correct offsets $\mathbf{t}$. We search for the value of $\hat{\mathbf{t}}$ such that the following function is minimized:

$$\min_{\hat{\mathbf{t}}} \|\mathbf{y} - \hat{\mathbf{y}}\|^2. \quad (26)$$

Using this method to find the offsets between the sets of samples $\mathbf{y}_m$, we can derive an algorithm to reconstruct the signal $f(t)$ from its combined sets of samples $\mathbf{y}$. A block diagram of such an algorithm is given in Fig. 12. Note that this algorithm does not use any specific property of the Fourier transform. It can be applied to arbitrary signals in a finite-dimensional Hilbert space, like wavelets, uniform splines, etc. More details can be found in [27].

B. Discussion

For this method, we need to have at least as many samples as the total number of unknowns, or $MN \geq L + M - 1$ for M odd, and $MN \geq L + M$ for M even (see Section III). Therefore, we need $M = \lceil (L-1)/(N-1) \rceil$ sets of samples for M odd, or $M = \lceil L/(N-1) \rceil$ for M even. Let us now study some examples of the objective function $\|\mathbf{y} - \hat{\mathbf{y}}\|^2$ (see (26)) as a function of the offsets $\hat{\mathbf{t}}$.

Example 5.1: A bandlimited function with 81 unknown Fourier coefficients is sampled with two sets of 90 samples. The same function is used here as in Example 4.1. The offset vector that we used is $\mathbf{t} = (0 \ 0.6)$. The objective function $\|\mathbf{y} - \hat{\mathbf{y}}\|^2$ from (26) is shown as a function of the offset t_1 in Fig. 13(a). It is 0 at the correct offset value ($t_1 = 0.6$) and is larger at other values of $\hat{t}_1$. Next to this global minimum, the function also has many local minima.

Similarly, we also take three sets of 41 samples from a bandlimited function with again 81 Fourier coefficients. The offset vector for this example is $\mathbf{t} = (0 \ 0.2 \ 0.6)$. Fig. 13(b) shows the objective function as a function of the offsets t_1 and t_2. Smaller values are represented by darker gray levels. Again, the minimum can clearly be seen, and is at the intersection of the dark horizontal, vertical and diagonal lines corresponding to pairwise alignments between the sets of samples. However, the objective function also has many local minima. □

From these examples, it is clearly visible that (26) is not easy to minimize (just like (24) in Section IV). Next to the global minimum, the objective function has also many local minima. We cannot use a standard algorithm like gradient descent for minimization. However, as can also be seen from Fig. 13(b), the objective function shows the same structure of horizontal, vertical, and diagonal lines as in Fig. 11(a).

Let us now give some further interpretation of the above results. As it was discussed in Section III, the MN-dimensional sample space can be divided into N orthogonal subspaces of dimension M. It is therefore possible to split the function from
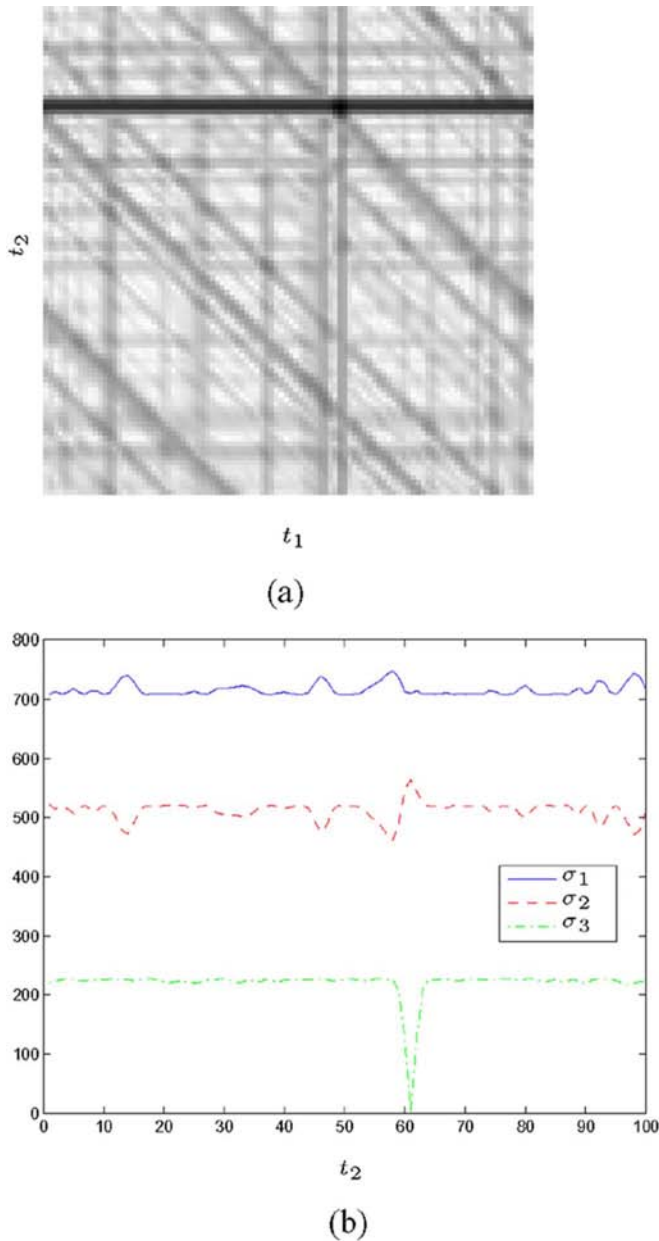

Fig. 11. Example of the objective function in (24). A signal with 81 unknown Fourier coefficients is sampled with three aliased sets of 41 samples. The exact offsets are $t_1 = 0.2$ and $t_2 = 0.6$. (a) Two-dimensional objective function has many local minima. (b) A slice of the error surfaces for $t_1 = 0.2$ is shown for the different singular values. The three singular values $\sigma_1, \sigma_2,$ and σ_3 do not vary much, except for σ_3 at the correct offset value.

V. SOLUTION USING PROJECTIONS

In this section, we describe a method that uses projections onto a subspace basis to find the relative offsets. This setup is the same as the setup used by Golub *et al.* in their variable projection method [25].

A. Method

From (8), we can see that $\mathbf{y}$ belongs to span($\mathbf{\Phi}_{\mathbf{t}}$). However, this is only true if the matrix $\mathbf{\Phi}_{\mathbf{t}}$ is constructed using the correct values for the offset vector $\mathbf{t}$ (except for degenerate cases). For another (incorrect) set of offset values $\hat{\mathbf{t}}$, (8) no longer holds and the sample vector $\mathbf{y}$ is not in span($\mathbf{\Phi}_{\hat{\mathbf{t}}}$) anymore. One can test

Fig. 12. Signal reconstruction algorithm using the projection method from Section V. The estimate for the offsets $\hat{\mathbf{t}}$ is updated iteratively. Once the estimate is considered good enough, the signal parameters $\boldsymbol{\alpha}$ are reconstructed.


Fig. 13. Examples of the objective function in (26). Next to the global minimum, it also contains many local minima. (a) Two aliased sets of 90 samples, with 81 unknown coefficients. The exact offset is $t_1 = 0.6$. (b) Three aliased sets of 41 samples, with 81 unknown coefficients. The exact offsets are $t_1 = 0.2$ and $t_2 = 0.6$. Small values are represented by dark pixels.

Fig. 14. Example of the decomposition of the objective function into its different components belonging to orthogonal M-dimensional subspaces.

(26) into independent terms according to these M-dimensional subspaces

$$\|\mathbf{y} - \hat{\mathbf{y}}\|^2 = \sum_{n=0}^{N-1} \left\| \mathbf{y}^{(n)} - \hat{\mathbf{y}}^{(n)} \right\|^2. \qquad (27)$$

In this equation, $\mathbf{y}^{(n)}$ is the projection of $\mathbf{y}$ onto the subspace $\mathcal{V}_n$ spanned by the nth vector for different values of $\hat{\mathbf{t}}$, and $\hat{\mathbf{y}}^{(n)}$ is the projection of $\mathbf{y}$ onto the vectors $\boldsymbol{\phi}_{\hat{\mathbf{t}}}^{n+iN}$ that belong to this space. An example of such a decomposition is given in Fig. 14.

As we also described in Section III, for the frequencies n that have M overlapping spectrum coefficients, the spaces $\mathcal{V}_n$ and span($\boldsymbol{\phi}_{\hat{\mathbf{t}}}^{n+iN}$) are the same, regardless of the value of $\hat{\mathbf{t}}$. For these frequencies, $\|\mathbf{y}^{(n)} - \hat{\mathbf{y}}^{(n)}\|^2 = 0$, and they do not contribute to the objective function in (27). The other terms each contribute a periodic term to the global objective function. The minimum of (27) can therefore be found by minimizing the different components individually (see Fig. 14). For each of the components, we can now minimize over an M-dimensional subspace instead of the original MN-dimensional space. In each of the subspaces, we still need to minimize the objective function jointly over the $M - 1$ offsets $\mathbf{t}$, but the vectors to project are now length M instead of MN previously. This approach also allows us to elim-

inate those subspaces from the search space that do not give information about the offsets. Such a decomposition reduces the complexity of our algorithm by a factor N^2, but even in such a case, the complexity is still slightly higher than the complexity of the rank-based method presented before.

Observe that this algorithm is not limited to bandlimited signals and sets of samples differing only by unknown offsets. As we use projections onto the space spanned by sampled basis functions, we can apply this algorithm to signals in an arbitrary (known) finite-dimensional Hilbert space. More complex motions like rotations or scalings can be taken into account by sampling the basis functions accordingly.

VI. PRACTICAL ISSUES

A. Sampling Kernel

In the different methods described above, no sampling kernel was considered. The signals are sampled using Diracs. Although this is not very realistic, it is an approximation that is often made to simplify the analysis. In most cases, it is not very difficult to take a different sampling kernel into account. The sampling operation can be considered as a convolution with a sampling kernel, followed by the actual sampling. In the frequency domain, this convolution can be seen as a multiplication with the Fourier transform of the sampling function. Therefore, as long as the sampling kernel does not remove frequencies, the reconstructed function can always be divided by the sampling kernel function again. This should cancel the effect of the sampling kernel. Such an operation can be performed after the reconstruc-

tion of an aliasing-free signal using the methods described in this paper. Of course, this supposes that the sampling kernel is known, space–time invariant, and not too ill-conditioned.

B. Images and Higher-Dimensional Signals

The above approaches were described for one-dimensional functions. However, all the methods can be extended to higher dimensional functions. Nothing in the descriptions is limited to a one-dimensional signal, so the algorithms can be directly applied to images or higher dimensional signals. The two-dimensional equations are derived in the Appendix. Of course, the complexity increases fast with the dimensionality. While a minimum of $\lceil L/N \rceil + 1$ sets of samples are required for a 1-D signal using the rank-based method, and $\lceil (L-1)/(N-1) \rceil$ sets of samples for the projection-based method, we need at least $\lceil L/N \rceil^2 + 1$ 2-D images for the rank-based method, and $\lceil (L^2-2)/(N^2-2) \rceil$ images for the projection-based method. As signals can be shifted along each of their dimensions, the number of offsets (which is also the dimensionality of the search space) also increases rapidly. Examples of the above methods on images are shown in Section VIII.

C. Minimization

The main difficulty with the methods described in Sections IV and V is to find the global minimum of the objective functions. A good approximation of the offsets $\mathbf{t}$ is required for an approach like gradient descent to converge to the global minimum. Such an approximation could be obtained by evaluating the error functions from (24) or (26) on a uniform (dense) grid of possible values $\hat{\mathbf{t}}$. The global minimum can then be found with high probability close to the smallest value obtained on this uniform grid. However, this is a computationally very intensive method. It only makes sense if the number of sets of samples M—and therefore also the number of offsets—is small, or when a first estimate of the offsets is available. Although this may sound like moving the problem to obtaining such a first estimate, it is a reasonable assumption. For example in super-resolution imaging, the shifts between the images are generally very small. Therefore, our search could be restricted to a small area around the origin (for example 10×10 pixels). Similarly, a first estimate could also be obtained from another registration method that does not use the aliasing. Such a method is described in the next subsection. Other approaches are heuristic methods that rely on the lines that can be seen in Figs. 11(a) and 13(b). They will also be described in the next subsection.

During the minimization, it would be interesting to know whether the global minimum has been found, or whether the current optimum is a local minimum. Such an indication can be obtained by evaluating (24) or (26) for the current value of $\hat{\mathbf{t}}$. Even for noisy measurements, there is a large difference between the average values of these functions and their value at the global minimum. This can therefore be used to check whether the algorithm has converged, and whether the result is reliable.

D. Heuristic Approaches

The high computational complexity of the algorithms from Sections IV and V is mainly due to the coupling between the offsets, i.e., the need to search the different offsets between the

signals jointly. The joint minimum is not necessarily located at the intersection of the minima from individual optimizations.

Algorithm 6.1 (Hierarchical Approach): If the sets of samples are images, we know that their coefficients are not arbitrary Gaussian random variables. In general, the amplitude of the Fourier transform of a natural image decays like $1/\omega$ (with ω the frequency) [30]. We can therefore assume that a good estimate for the offsets can be obtained from the low frequencies of the sampled sets. For these frequencies, the aliased coefficients are much smaller than the base spectrum coefficients, and can be neglected in a first estimate. We used the method from Vandewalle *et al.* [4] for this. Once we have such an initial estimate, we can use the methods from Sections IV or V using a gradient descent algorithm to obtain a more precise estimate for the motion parameters, taking the aliasing into account. $\qquad \square$

Algorithm 6.2 (Keeping the P Best Pairwise Alignments): From the horizontal, vertical, and diagonal lines that are often visible in figures like Figs. 11(a) and 13(b), it can be seen that an independent pairwise alignment often works. Such lines correspond to pairwise registrations of the first and second, first and third, and second and third signals, respectively. Therefore, we search a fixed number P of local minima for the pairwise registration between the first set of samples and each of the other sets. When performing such a pairwise registration between the first and the mth set of samples, all the other offset values are kept constant. For each pair, a vector of P possible offset values $(\hat{t}_{m_1} \ \hat{t}_{m_2} \ \cdots \ \hat{t}_{m_P})$ is obtained. Next, the best combination $\hat{\mathbf{t}}$ of these pairwise local minima is searched among all possible combinations. The global minimum can then be searched in the neighborhood of this value of $\hat{\mathbf{t}}$. Keeping only a single minimum for each of the pairwise alignments $(P = 1)$ is generally not sufficient, due to the different approximations made by pairwise registration. However, by trying all the combinations of the best $P > 1$ pairwise alignments, the algorithm typically converges to the correct result. In our simulations using five sets of 32×32 images, good results were obtained with $P = 5$, but the optimal number depends on the size of the problem. $\qquad \square$

We have also investigated a method that searches specifically for the lines in the objective function as they were shown in Figs. 11(a) and 13(b). However, because initially the objective function is unknown, it is not trivial to design an algorithm that finds these lines. Also, not all the lines indicating pairwise alignments are visible for every specific realization. Finally, such lines are less distinguishable for decreasing signal-to-noise ratio (SNR) values. Such an algorithm therefore typically has lower performance and/or a higher computational complexity than the uniform grid evaluation presented in the previous section.

VII. COMPLEXITY

In this section, we discuss the computational complexity of the different methods. For each of the methods described above, the complexity of computing the offsets $\mathbf{t}$ can be written as the number of times the error function has to be evaluated multiplied with the number of operations required to evaluate the error function. In this analysis, we will assume that the variables N and L grow at the same rate, and M grows at a lower rate. The number of operations required for the reconstruction is the same for all the algorithms and is thus not determinant.

TABLE I
COMPUTATIONAL COMPLEXITY FOR THE DIFFERENT METHODS ON 1-D SIGNALS. THE TOTAL COMPLEXITY IS
OBTAINED BY MULTIPLYING THE NUMBER OF FUNCTION EVALUATIONS WITH THE COMPLEXITY OF A SINGLE
FUNCTION EVALUATION. THE COST TO OBTAIN THE INITIAL ESTIMATE IN ALG. 6.1 IS DENOTED AS C_{init}

	matrix rank method $M^3 + M^2 N$	projection method $MNL^2 + L^3$	projection method (Fourier) $S^2 L$
regular grid N^{M-1} eval.	$N^{M-1}(M^3 + M^2 N)$	$N^{M-1}(MNL^2 + L^3)$	$N^M S^2 L$
hier. method (Alg. 6.1) E eval. + C_{init}	$(M^3 + M^2 N)E + C_{\mathrm{init}}$	$(MNL^2 + L^3)E + C_{\mathrm{init}}$	$S^2 L E + C_{\mathrm{init}}$
pairw. al. (Alg. 6.2) $MN + 5^{M-1}$ eval.	$(MN + 5^{M-1})(M^3 + M^2 N)$	$(MN + 5^{M-1})(MNL^2 + L^3)$	$(MN + 5^{M-1})S^2 L$

Let us start by analyzing the complexity of the objective function evaluation both for the matrix rank method (Section IV) and the projection method (Section V). In the matrix rank method, the smallest singular value of the matrix $\mathbf{Y}_{\hat{\mathbf{t}}}^D$ has to be computed for every error evaluation. The construction of $\mathbf{Y}_{\hat{\mathbf{t}}}^D$ for a specific set of offsets $\hat{\mathbf{t}}$ requires $O((M-1)N)$ operations. To compute the smallest singular value, we use the inverse power method on $\mathbf{Y}_{\hat{\mathbf{t}}}^{D*}\mathbf{Y}_{\hat{\mathbf{t}}}^D$ [29]. This matrix multiplication involves $O(M^2 N)$ operations, and the inverse power method itself requires $O(M^3 + aM^2)$ operations, where a is the number of iterations required for the method to converge. As we assume that the difference with the second smallest singular value is large, convergence will be fast. So the total number of operations becomes

$$C_{\mathrm{rank}} = O((M-1)N + M^2 N + M^3 + aM^2)$$
$$= O(M^3 + M^2 N). \tag{28}$$

For the projection method, we need to compute the projection of the sample vector $\mathbf{y}$ onto the space spanned by the columns of $\mathbf{\Phi}_{\hat{\mathbf{t}}}$. This can be done using (25). In general, such a projection has complexity

$$C_{\mathrm{proj}} = O(MNL^2 + L^3). \tag{29}$$

However, if enough storage space is available, these projection matrices can be precomputed, because they do not depend on the actual signal. In that case, the complexity is reduced to

$$C'_{\mathrm{proj}} = O(M^2 N^2). \tag{30}$$

For the Fourier basis, the complexity of such a projection can be strongly reduced. The blocks $\mathbf{\Phi}_{\hat{\mathbf{t}}_m}$ of the matrix $\mathbf{\Phi}_{\hat{\mathbf{t}}}$ can be divided into an $N \times L$ IDFT inverse Fourier transform block $\mathbf{F}^*$, multiplied by an $L \times L$ diagonal offset matrix $\mathbf{D}_{t_m}$: $\mathbf{\Phi}_{\hat{\mathbf{t}}_m} = \mathbf{F}^*\mathbf{D}_{t_m}$ (see also (8)). Using this decomposition, we can simplify the projection formula (25). The multiplication $\mathbf{\Phi}_{\hat{\mathbf{t}}}^*\mathbf{y}$ can be decomposed into M multiplications $\mathbf{D}_{\hat{t}_m}^*\mathbf{Fy}$, where $\mathbf{Fy}$ is the Fourier transform of $\mathbf{y}$. This Fourier transform can be computed outside of the iteration loop, because it is independent of the offsets $\mathbf{t}$. Only the multiplications with the diagonal matrices

$\mathbf{D}_{\hat{t}_m}$ remain, which require ML operations. The matrix $\mathbf{\Phi}_{\hat{\mathbf{t}}}^*\mathbf{\Phi}_{\hat{\mathbf{t}}}$ has nonzero elements only on the main diagonal and on the iNth diagonals, due to the orthogonality of the Fourier vectors. Its inverse, $(\mathbf{\Phi}_{\hat{\mathbf{t}}}^*\mathbf{\Phi}_{\hat{\mathbf{t}}})^{-1}$, can therefore be computed efficiently in $O(L^3/N^2) = O(S^2 L)$ operations. Finally, the multiplication with the first matrix $\mathbf{\Phi}_{\hat{\mathbf{t}}}$ can be decomposed into M multiplications with diagonal matrices, and M inverse DFTs. As the error can as well be computed in the Fourier domain (using Parseval's theorem), we only need to multiply with the diagonal matrices, which requires again ML operations. The overall complexity can then be approximated as

$$C_{\mathrm{proj},F} = O(ML + S^2 L + ML) = O(S^2 L) \tag{31}$$

where S is the maximum number of overlapping spectral components, which is typically of the same order as M, or $C_{\mathrm{proj},F} = O(M^2 L)$.

For the standard algorithm described in Section VI-C, we first evaluate the objective function on a uniform grid of N^{M-1} points, requiring N^{M-1} error evaluations. Next, a standard minimization algorithm is applied near the minimum value that was obtained on the uniform grid from the first part. Let us call the number of error evaluations in this second part of the minimization E (typical values in our simulations are around $E = 30$). This is negligible compared to the number of evaluations on the regular grid, and the total number of evaluations with this standard algorithm can therefore be approximated as

$$O(N^{M-1} + E) = O(N^{M-1}). \tag{32}$$

For the hierarchical method from Algorithm 6.1, we need only E evaluations. Of course, the complexity for computing the initial estimate needs to be added to this. Algorithm 6.2, using pairwise alignments, requires

$$O((M-1)N + 5^{M-1} + E) = O(MN + 5^{M-1}) \tag{33}$$

error function evaluations. The complexity of the different algorithms is summarized in Table I.

For the reconstruction, the set of linear equations from (8) has to be solved. The complexity of this operation is $O(MNL^2 + L^3)$.
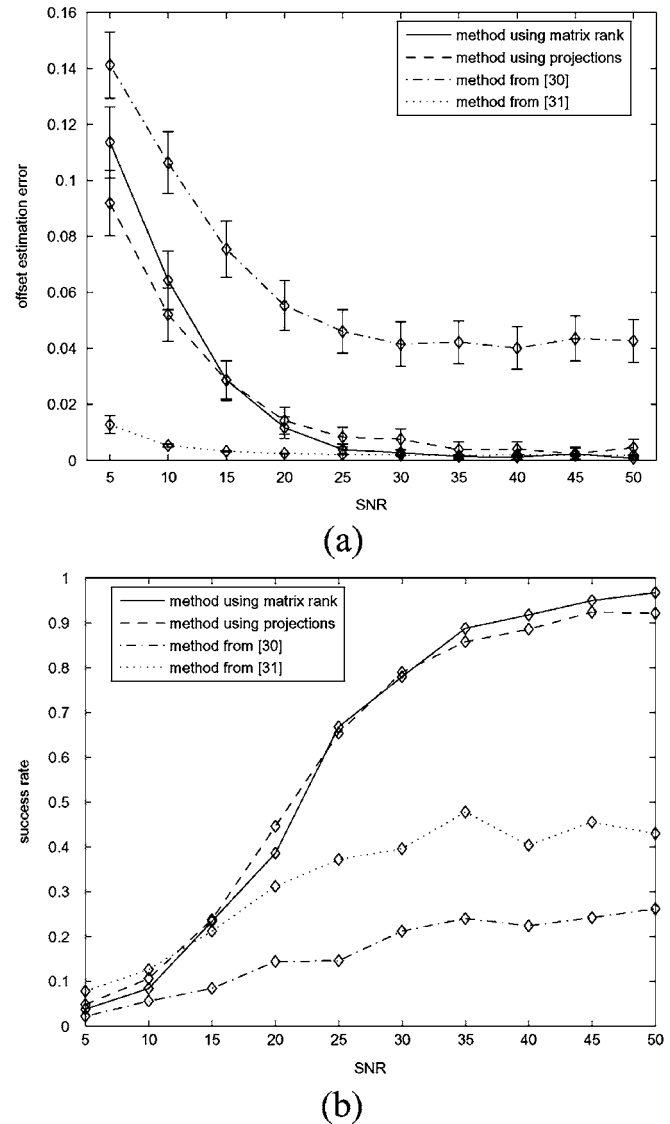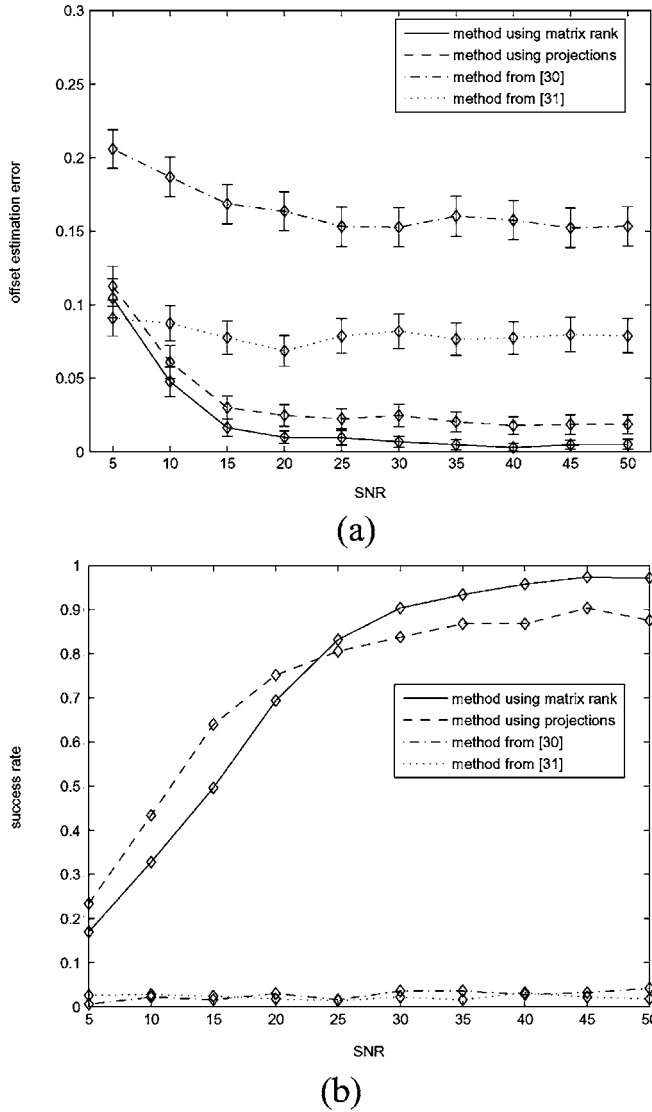
Fig. 15. Results of the different algorithms as a function of the signal-to-noise ratio (SNR) for white Gaussian noise signals. The matrix rank algorithm from Section IV performs slightly better than the projection algorithm from Section V. Both algorithms perform clearly better than the algorithms that do not use the aliasing for registration (method from [4] and from [31]). Parameter values of $M = 3, L = 81$, and $N = 41$ were used. (a) Offset estimation error as a function of the SNR of the sampled signals. An offset error of 1 corresponds to a shift over the entire signal period. (b) Success rate of the registration and reconstruction as a function of the SNR.

Fig. 16. Results of the different algorithms as a function of the SNR for pink Gaussian noise signals (with $1/\omega$ behavior). Parameter values of $M = 3, L = 81$, and $N = 41$ were used. (a) Offset estimation error as a function of the SNR of the sampled signals. An offset error of 1 corresponds to a shift over the entire signal period. The algorithm from [31] has a lower average absolute error than the rank based and projection based algorithms presented here. They still outperform the algorithm from [4], however. (b) Success rate of the registration and reconstruction as a function of the SNR. Both algorithms presented here outperform the other algorithms and compute the offsets up to a precision of 10^{-3} in a larger proportion of the simulations.

VIII. RESULTS

The above algorithms are tested and compared in a number of simulations. We also compared our algorithms to our earlier approach [4] and the algorithm by Pham *et al.* [31]. The algorithm from [4] estimates the offsets from the phase difference between two signals for low frequency values, as it was designed for partially aliased signals. The algorithm from [31] is an iterative Taylor-based algorithm, and we used the implementation by Pham *et al.* [32]. The simulations are performed on one-dimensional random bandlimited signals. The Fourier coefficients of the signals are generated as a white Gaussian random process. Hence, the resulting (time domain) signals also form white Gaussian random processes. A number of simulations are performed with different random offset values, and different amounts of additive white Gaussian noise. The performance of the different algorithms on 1-D signals is compared in Figs. 15, –17. Fig. 15(a) and (b) shows the mean absolute error in the shift estimates and the success rate as a function of the SNR, respectively. The success rate of the methods is defined as the relative number of simulations in which the error on the registration is smaller than 10^{-3}. All the results were averaged over 250 simulations. Parameter values of $M = 3, L = 81$, and $N = 41$ were used in all the simulations. An offset value of 1 corresponds to a shift over the whole signal period. We can see that the method using matrix rank from Section IV performs slightly better than
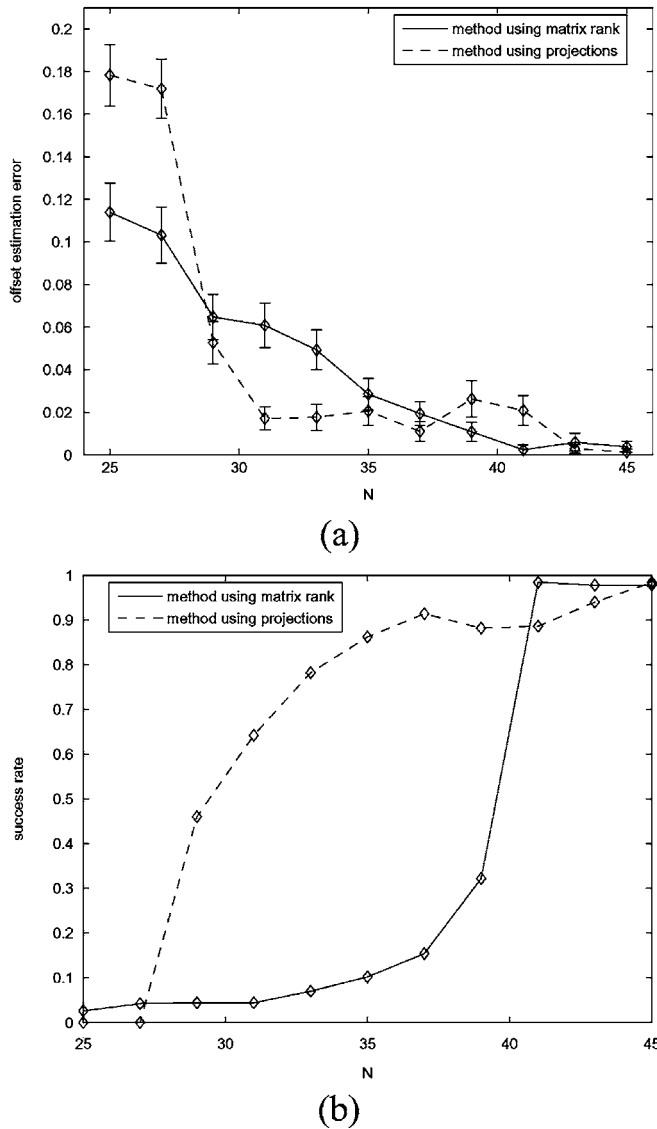
(a)

(b)

Fig. 17. Results of the different algorithms as a function of the number of samples N. Parameter values $M = 3$ and $L = 81$ were used in all simulations, and for N, values from 25 to 45 were used. No noise was added in this setup. The algorithm from Section V performs better than the algorithm from Section IV if the number of samples per set N is not sufficiently high. (a) Offset estimation error as a function of the number of samples per set N. An offset error of 1 corresponds to a shift over the entire signal period. (b) Success rate of the registration and reconstruction as a function of the number of samples per set N.

the method using projections described in Section V. It can also be clearly seen that both the rank-based and the projection-based algorithm perform much better on such signals than the methods that do not use the aliasing that is present in the sampled signals (from [4] and [31]).

Fig. 16 shows the results of similar simulations as those from Fig. 15. However, to simulate the typical behavior of images, we performed the simulations on pink noise signals, that have a $1/\omega$ decay as a function of the frequency ω. We can see that in such a case, the influence of the aliasing is decreased, and the methods from [4] and [31] perform much better than in the previous simulation. The method by Pham *et al.* [31] even has a better average performance than our methods described here. However, from Fig. 16(b), we can see that our algorithms still computed

the offsets up to a precision of 10^{-3} in a larger proportion of the cases. Moreover, if the sampling frequency N is further reduced (and the number of sets M increased), more aliasing is present in the sampled signals, and the performance of the algorithms from [4] and [31] will typically decrease faster than the algorithms presented here.

The absolute error in the shift estimates and the success rate are plotted as a function of the number of samples per set N in Fig. 17. In these simulations, parameters $M = 3$ and $L = 81$ were used, and for N, values from 25 to 45 were taken. No noise was added for these simulations. Again, the results were averaged over 250 simulations. The performance of the different algorithms increases with increasing number of samples. This is what we expected, as an increasing number of samples per set gives increasing amounts of information for the same number of unknowns. For the matrix rank algorithm from Section IV, the minimum number of samples per set that are required in this setup is $N = 41$, while for the projection algorithm from Section V, $N = 28$ samples per set are needed. This explains the better performance of the projection algorithm for $28 \leq N < 41$. With more than 41 samples per set, both algorithms perform very well. When $N < 28$ samples are available per set, all algorithms have low accuracy.

Our algorithms are also applied to some 2-D signals (images). In this case, the signals were undersampled by only two, because of the high complexity (see the discussion in Section VI). The results can be seen in Fig. 18. A double resolution image is accurately reconstructed from a set of low resolution images, both using the rank-based algorithm from Section IV [Fig. 18(a) and (b)] and the projection-based algorithm from Section V [Fig. 18(c) and (d)].

In all the simulations, we use periodic signals. This is not the case in most real applications, but we can generally assume that the shift between the sets of samples is small. The differences between the signals due to their aperiodicity are therefore small, and can be neglected.

IX. CONCLUSION

In this paper, we presented a theory of super-resolution from unregistered aliased sets of samples. We formulated the problem mathematically, and proved that the solution is unique if enough sets of samples are available. We described two methods to compute this unique solution. In the first method, the rank of a modified sample matrix is checked. The second method uses projections onto subspaces to compute the offsets between the different sets of samples. Both methods can be used for the reconstruction of one or two-dimensional signals from multiple sets of aliased samples. The main limitation of these methods is their computational complexity. They are therefore mainly applicable in domains that do not require real-time reconstruction. For example in satellite imaging, a very high resolution is desired, even if this requires some computational effort. Future work is mainly oriented towards reducing the computational complexity of the methods. In many cases, this super-resolution problem can be written as a set of polynomial equations. It is then possible to reconstruct the signal using Gröbner basis methods [33], [34], which are, however, computationally very expensive.

Fig. 18. Results of the different algorithms on images. (a) One of the five 32×32 images used as input for Algorithm 6.2, with the rank-based method from Section IV. (b) Reconstructed 63×63 image from the images from (a). (c) One of the five 8×8 images used as input for Algorithm 6.1, with the projection-based method from Section V. (d) Reconstructed 15×15 image from the images from (c).

APPENDIX

In this Appendix, we extend the above theory to 2-D signals. Assume we have a periodic, bandlimited image

$$f(\boldsymbol{x}) = \sum_{k_1,k_2=-K}^{K} \alpha_{\boldsymbol{k}} e^{j2\pi \boldsymbol{k}^T \boldsymbol{x}}$$

$$\text{with } \boldsymbol{k} = \begin{pmatrix} k_1 \\ k_2 \end{pmatrix}, \quad \boldsymbol{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}. \quad (34)$$

It is sampled with M images at horizontal and vertical frequencies N_1 and N_2

$$\begin{aligned}
\boldsymbol{y}_m(\boldsymbol{n}) &= f\left(\frac{n_1}{N_1} + t_{m,1}, \frac{n_2}{N_2} + t_{m,2}\right) \\
&= \sum_{k_1,k_2=-K}^{K} \alpha_{\boldsymbol{k}} e^{j2\pi(k_1(\frac{n_1}{N_1}+t_{m,1})+k_2(\frac{n_2}{N_2}+t_{m,2}))} \\
&= \sum_{k_1,k_2=-K}^{K} \alpha_{\boldsymbol{k}} e^{j2\pi(k_1 t_{m,1}+k_2 t_{m,2})} e^{j2\pi(k_1 \frac{n_1}{N_1}+k_2 \frac{n_2}{N_2})}.
\end{aligned}$$
$$(35)$$

The discrete Fourier transform (DFT) of such an image can be expressed as

$$\begin{aligned}
&\boldsymbol{Y}_m(\boldsymbol{l}) \\
&= \sum_{n_1,n_2=0}^{\boldsymbol{N}-1} \boldsymbol{y}_m(\boldsymbol{n}) e^{-j2\pi(l_1 \frac{n_1}{N_1}+l_2 \frac{n_2}{N_2})} \\
&= \sum_{n_1,n_2=0}^{\boldsymbol{N}-1} \sum_{k_1,k_2=-K}^{K} \alpha_{\boldsymbol{k}} e^{j2\pi(k_1 t_{m,1}+k_2 t_{m,2})} e^{j2\pi(k_1 \frac{n_1}{N_1}+k_2 \frac{n_2}{N_2})} \\
&\quad \times e^{-j2\pi(l_1 \frac{n_1}{N_1}+l_2 \frac{n_2}{N_2})} \\
&= \sum_{k_1,k_2=-K}^{K} \alpha_{\boldsymbol{k}} e^{j2\pi(k_1 t_{m,1}+k_2 t_{m,2})} \\
&\quad \times \sum_{n_1,n_2=0}^{\boldsymbol{N}-1} e^{j2\pi((k_1-l_1)\frac{n_1}{N_1}+(k_2-l_2)\frac{n_2}{N_2})} \\
&= \sum_{i_1,i_2} \alpha_{l_1+i_1 N_1, l_2+i_2 N_2} \\
&\quad \times e^{j2\pi((l_1+i_1 N_1)t_{m,1}+(l_2+i_2 N_2)t_{m,2})}. \quad (36)
\end{aligned}$$

In other words, $\boldsymbol{Y}_m(\boldsymbol{l})$ is a weighted sum of the overlapping Fourier coefficients at frequency $\boldsymbol{l}$. Combining all the DFT coefficients $\boldsymbol{Y}_m(\boldsymbol{l})$ into a vector $\mathbf{Y}_m$, we can write (see also (11))

$$\mathbf{Y}_m = \sum_{i_1,i_2} e^{j2\pi(i_1 N_1 t_{m,1}+i_2 N_2 t_{m,2})} \mathbf{D}'_{t_m} \boldsymbol{\alpha}_{i_1,i_2} \quad (37)$$

with now $\boldsymbol{Y}_m(l_1 + N_1 l_2) = \boldsymbol{Y}_m(l_1, l_2)$, $\mathbf{D}'_{t_m}(l_1 + N_1 l_2, l_1 + N_1 l_2) = e^{j2\pi(l_1 t_{m,1}+l_2 t_{m,2})}$, and $\boldsymbol{\alpha}_{i_1,i_2}(k_1, k_2) = \boldsymbol{\alpha}(k_1 + i_1 N_1, k_2 + i_2 N_2)$. The analysis from Section IV is therefore still valid for 2-D signals.

Similarly, we can combine all the samples of an image into a vector $\mathbf{y}_m$ (with $\mathbf{y}_m(n_1 + N_1 n_2) = \boldsymbol{y}_m(n_1, n_2)$)

$$\mathbf{y}_m = \boldsymbol{\Phi}_{t_m} \boldsymbol{\alpha}, \quad (38)$$

with $\boldsymbol{\Phi}_{t_m}$ an $N_1 N_2 \times L_1 L_2$ sampled basis matrix with elements $\boldsymbol{\Phi}_{t_m}(n_1 + N_1 n_2, l_1 + L_1 l_2) = e^{j2\pi(n_1 l_1/N_1+n_2 l_2/N_2)}$. We can then repeat the analysis from Section V using (38) instead of (7).

REFERENCES

[1] P. Vandewalle, L. Sbaiz, J. Vandewalle, and M. Vetterli, "How to take advantage of aliasing in bandlimited signals," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, May 2004, vol. 3, pp. 948–951.

[2] P. Vandewalle, L. Sbaiz, M. Vetterli, and S. Süsstrunk, "Super-resolution from highly undersampled images," in *Proc. IEEE Int. Conf. Image Processing*, Sep. 2005, vol. 1, pp. 889–892.

[3] P. Vandewalle, L. Sbaiz, J. Vandewalle, and M. Vetterli, "Aliasing Is Good for You: Joint Registration and Reconstruction for Super-Resolution," Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, Tech. Rep., Jul. 2006, .

[4] P. Vandewalle, S. Süsstrunk, and M. Vetterli, "A frequency domain approach to registration of aliased images with application to super-resolution," *EURASIP J. Appl. Signal Processing (Special Issue on Super-Resolution Imaging)*, vol. 2006, p. 14, 2006, article ID 71459.

[5] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, Jul. 1948.

[6] M. Unser, "Sampling—50 years after Shannon," *Proc. IEEE*, vol. 88, no. 4, pp. 569–587, Apr. 2000.

[7] F. Marvasti, Ed., *Nonuniform Sampling: Theory and Practice*. New York: Springer, 2001.

[8] T. Strohmer, "Computationally attractive reconstruction of bandlimited images from irregular samples," *IEEE Trans. Image Process.*, vol. 6, no. 4, pp. 540–548, Apr. 1997.

[9] P. Marziliano and M. Vetterli, "Reconstruction of irregularly sampled discrete-time bandlimited signals with unknown sampling locations," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3462–3471, Dec. 2000.

[10] A. Papoulis, "Generalized sampling expansion," *IEEE Trans. Circuits Syst.*, vol. 24, no. 11, pp. 652–654, Nov. 1977.

[11] M. Unser and J. Zerubia, "Generalized sampling without bandlimiting constraints," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process.*, Apr. 1997, vol. 3, pp. 2113–2116.

[12] ——, "Generalized sampling: Stability and performance analysis," *IEEE Trans. Signal Process.*, vol. 45, no. 12, pp. 2941–2950, Dec. 1997.

[13] R. Y. Tsai and T. S. Huang, "Multiframe image restoration and registration," in *Adv. Computer Vision Image Process.*, T. S. Huang, Ed. : JAI Press, 1984, vol. 1, pp. 317–339.

[14] S. Borman and R. Stevenson, "Spatial resolution enhancement of low-resolution image sequences—A comprehensive review with directions for future research," Univ. Notre Dame, Notre Dame, IN, Tech. Rep., 1998.

[15] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in super-resolution," *Int. J. Imag. Syst. Technol.*, vol. 14, no. 2, pp. 47–57, Aug. 2004.

[16] M. G. Kang and S. Chaudhuri, Eds., *IEEE Signal Process. Mag. (Special Issue on Super-Resolution)*, vol. 20, no. 3, May 2003.

[17] M. Ng, T. Chan, M. G. Kang, and P. Milanfar, Eds., *EURASIP J. Appl. Signal Process. (Special Issue on Super-Resolution)*, vol. 2006, 2006.

[18] R. C. Hardie, S. Cain, K. J. Barnard, J. Bognar, E. Armstrong, and E. A. Watson, "High resolution image reconstruction from a sequence of rotated and translated infrared images," in *Proc. SPIE*, 1997, vol. 3063, pp. 113–124.

[19] E. Kaltenbacher and R. C. Hardie, "High resolution infrared image reconstruction using multiple, low resolution, aliased frames," in *Proc. SPIE Hybrid Image Signal Processing V*, Jun. 1996, vol. 2751, pp. 142–152.

[20] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1327–1344, Oct. 2004.

[21] T. Q. Pham, L. J. van Vliet, and K. Schutte, "Robust fusion of irregularly sampled data using adaptive normalized convolution," *EURASIP J. Appl. Signal Process.*, vol. 2006, p. 12, 2006, article ID 83268.

[22] R. C. Hardie, K. J. Barnard, and E. E. Armstrong, "Joint MAP registration and high-resolution image estimation using a sequence of undersampled images," *IEEE Trans. Image Process.*, vol. 6, no. 12, pp. 1621–1633, Dec. 1997.

[23] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, Jun. 2002.

[24] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1167–1183, Sep. 2002.

[25] G. Golub and V. Pereyra, "Separable nonlinear least squares: The variable projection method and its applications," *Inverse Problems*, vol. 19, no. 2, pp. R1–R26, 2003.

[26] M. Schwab, M. Karrenbach, and J. Claerbout, "Making scientific computations reproducible," *Comput. Sci. Eng.*, vol. 2, no. 6, pp. 61–67, Nov. 2000.

[27] P. Vandewalle, "Super-resolution from unregistered aliased images," Ph.D. dissertation, School of Computer and Communication Sciences, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, Jul. 2006, no. 3591.

[28] P. Marziliano, "Sampling innovations," Ph.D. Dissertation, School of Computer and Communication Sciences, Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland, 2001, no. 2369.

[29] G. Strang, *Linear Algebra and Its Applications*, 3rd ed. : Saunders College Publishing, 1988.

[30] A. Torralba and A. Oliva, "Statistics of natural image categories," *Network: Comput. Neural Syst.*, vol. 14, no. 3, pp. 391–412, Aug. 2003.

[31] T. Q. Pham, M. Bezuijen, L. J. van Vliet, K. Schutte, and C. L. L. Hendriks, "Performance of optimal registration estimators," in *Proc. SPIE Visual Information Processing XIV*, May 2005, vol. 5817, pp. 133–144.

[32] M. van Ginkel, G. van Kempen, C. Luengo, and L. J. van Vliet, Dipimage, a Scientific Image Processing Toolbox for Matlab [Online]. Available: http://www.ph.tn.tudelft.nl/DIPlib/ 2006

[33] P. Vandewalle, L. Sbaiz, and M. Vetterli, "Signal reconstruction from multiple unregistered sets of samples using Groebner bases," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP)*, May 2006, vol. 3, pp. 604–607.

[34] L. Sbaiz, P. Vandewalle, and M. Vetterli, "Groebner basis methods for multichannel sampling with unknown offsets," *Appl. Comput. Harmon. Anal.*, 2006, submitted for publication.

Patrick Vandewalle (S'03–A'06) received the M.S. degree in electrical engineering from Katholieke Universiteit Leuven (K.U. Leuven), Belgium, in 2001 and the Ph.D. degree from Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland, in 2006.

From 2001 to 2002, he worked as a Research Assistant in the Medical Imaging Laboratory in the Department of Electrical Engineering (ESAT), K.U. Leuven. He is currently a Postdoctoral Researcher at the Audiovisual Communications Laboratory at EPFL. His research interests are in signal and image processing, sampling, and digital photography.

Luciano Sbaiz (M'98–SM'07) received the "Laurea in Ingegneria" degree in electronic engineering and the Ph.D. degree from the University of Padova, Padova, Italy, in December 1993 and June 1998, respectively.

Between 1998 and 1999, he was Postdoctoral Researcher at the Audiovisual Communications Laboratory at Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland, where he conducted research on the application of computer vision techniques to the creation of video special effects. In 1999, he joined Dartfish Ltd., Fribourg, Switzerland, as Project Manager. Within the company, he developed video special effects for television broadcasting and sport analysis. In 2004, he became a Senior Researcher at the Audiovisual Communications Laboratory at EPFL, where he conducts research on signal processing. His activities are in the field of image and audio processing, superresolution techniques, and tomography.

Joos Vandewalle (S'71–M'79–SM'82–F'92) received the Electrical Engineering degree and the Ph.D. degree in engineering from the Katholieke Universiteit Leuven (K.U. Leuven), Belgium, in 1971 and 1976, respectively.

From 1976 to 1978, he was Research Associate and from July 1978 to July 1979 he was Visiting Assistant Professor both at the University of California, Berkeley. Since July 1979, he has been with the Department of Electrical Engineering (ESAT) of the K.U. Leuven, where he has been a Full Professor since 1986 and the head of the SCD division at ESAT, which has more than 130 researchers. He held positions as Chairman of the Department of Electrical Engineering and Vice Dean of the Faculty of Engineering there. In the second semester of 2002–2003, he was on sabbatical leave at the I3S laboratory of CNRS, Sophia Antipolis, France. Since February 2005, he has been a member of the Governing Board of Exact Sciences at the K.U. Leuven. He teaches courses in linear algebra, linear and nonlinear system and circuit theory, system identification, and neural networks. His research interests are mainly in mathematical system theory and its applications in circuit theory, control, signal processing, cryptography, and neural networks. His recent research interests are in nonlinear methods (support vector machines, multilinear algebra) for

data processing. He has authored or coauthored more than 300 international journal papers in these areas. He is the coauthor of four books and co-editor of five books.

Dr. Vandewalle was elected Fellow of the IEEE in 1992 for contributions to nonlinear circuits and systems and in 2006 as Vice-President Technical Activities of the IEEE CAS Society. He is also Fellow of the Institution of Electrical Engineers (IEE), United Kingdom. He is a member of the editorial board of various scientific journals, and program chairman of various conferences. He received several best paper awards and research awards.

Martin Vetterli (S'86–M'86–SM'90–F'95) received the Engineering degree from Eidgenoessische Technische Hochschule (ETH), Zurich, Switzerland, in 1981, the M.S. degree from Stanford University, Stanford, CA, in 1982, and the Ph.D. degree from Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland, in 1986.

In 1986, he joined Columbia University, New York, first with the Center for Telecommunications Research and then with the Department of Electrical Engineering, where he was an Associate Professor of Electrical Engineering. In 1993, he joined the University of California at Berkeley, were he was Full Professor until 1997. Since 1995, he has been a Professor at EPFL, where he headed the Communication Systems Division (1996–1997) and heads the Audiovisual Communications Laboratory. From 2001 to 2004, he directed the National Competence Center in Research on mobile information and communication systems. He has also been Vice-President for International Affairs at EPFL since October 2004. He has held visiting positions at ETH Zurich (ETHZ) in 1990 and Stanford University in 1998. His research interests are in the areas of applied mathematics, signal processing, and communications. He is the coauthor of a textbook on *Wavelets and Subband Coding* (Englewood Cliffs, NJ: Prentice-Hall, 1995) and over 100 journal papers.