ELSEVIER

Signal Processing 63 (1997) 175–197

SIGNAL PROCESSING

Joint automatic design of prefilter and decimation grid for the reduction of spectral redundancy in 2D digital signals

Federico Pedersini, Augusto Sarti*, Stefano Tubaro

Dipartimento di Electtronica e Informazione, Politecnico di Milano, Piazza L. Da Vinci, 32, 20133 Milano, Italy

Received 28 March 1996; received in revised form 14 April 1997

Abstract

In this paper we propose and discuss a new technique for reducing the spectral redundancy of 2D digital signals through decimation on an arbitrary sublattice. Starting from a second-order analysis of the signal's spectral extension, the method automatically selects a decimation grid that minimizes the overlap between spectral replicas and defines an ideal hexagonal prefilter that maximizes the spectral energy to be preserved. Furthermore, the method automatically designs the hexagonal prefilter through a modified version of the frequency transformation method. In particular, it geometrically designs the transformation map in such a way to fit the prefilter's passband. The method has been tested over a variety of 2D test signals (non-synthetic digital images) in order to evaluate the performance of the method in terms of the impact of the spectral truncation on the overall quality of the reconstructed signal.

Zusammenfassung

In diesem Artikel stellen wir eine neue Technik vor, die die spektrale Redundanz zweidimensionaler Digitalsignale durch Dezimierung auf ein willkürliches Untergitter reduziert, und diskutieren sie. Ausgehend von einer Analyse zweiter Ordnung der spektralen Erweiterung des Signals wählt die Methode automatisch ein Dezimierungsgitter, das die Überlappung zwischen den spektralen Kopien minimiert, und definiert ein ideales hexagonales Vorfilter, das die spektrale Energie, die erhalten werden soll, maximiert. Außerdem entwirft sie durch eine modifizierte Version der Methode der Frequenztranformation automatisch das hexagonale Vorfilter. Insbesondere entwirft sie geometrisch eine an den Durchlaßbereich des Vorfilters angepaßte Übertragungsfunktion. Die Methode wurde an einer Reihe von zweidimensionalen Testsignalen (nichtsynthetische digitale Bilder) ausprobiert, um die Leistung der Methode hinsichtlich der Auswirkung der spektralen Beschneidung auf die Gesamtqualität des rekonstruierten Signals abschätzen zu können.

Résumé

Nous proposons et discutons dans cet article une technique nouvelle pour la réduction de la redondance spectrale des signaux numériques 2D par décimation sur un sous-treillis arbitraire. A partir d'une analyse d'ordre deux de l'extension spectrale du signal, la méthode sélectionne automatiquement une grille de décimation minimisant le recouvrement entre

* Corresponding author. Tel.: 39 2 2399 3647; fax: 39 2 2399 3413; e-mail: sarti@elet.polimi.it.

0165-1684/97/$17.00
PII S0165-1684(97)00155-2

les répliques spectrales et définissant un pré-filtre hexagonal idéal qui maximise l'énergie spectrale à préserver. De plus, la méthode conçoit automatiquement le pré-filtre hexagonal avec une version modifiée de la méthode de transformation fréquentielle. En particulier, elle conçoit géométriquement la fonction de transformation de telle manière à s'ajuster à la bande passante du pré-filtre. La méthode a été testée sur une variété de signaux 2D (images numériques non synthétiques) pour évaluer ses performances en termes d'impact de la troncation spectrale sur la qualité globale du signal reconstruit.

Keywords: Sampling lattice; Decimation; Prefilter; FIR filter

1. Introduction

The *spectral redundancy* of multidimensional digital signals can be thought of as the 'degree of sparseness' of the spectral replicas of the original analog signal, which are generated by the sampling process. The problem of selecting a regular sampling structure that results in the most densely packed replications of the spectrum of the original 2D analog signal has often been addressed in the literature. In fact, there exists a variety of applications, from array processing [9, 21] (radar, sonar and seismic) to image processing [5, 6, 8], that could greatly benefit from a rational selection of data samples and a careful elimination of negligible spectral components for the purpose of avoiding problems of aliasing. In fact, it is well-known that minimizing the gap among the replicas (generated by the sampling process) of the analog spectrum reduces information redundancy [5].

Spectral packing through decimation is not an easy task as it consists not just of a rational selection of data samples, but it also needs a careful spectral truncation for avoiding aliasing. In order to perform anti-aliasing prefiltering, in fact, knowing the area of the spectral extension (spectral occupancy) is not sufficient; we also need to consider its shape.

In order to be able to reduce spectral redundancy through decimation, we need to solve a number of problems, the first of which is finding all possible sublattices of an assigned *order* (density reduction ratio). Given a sampling lattice, the number of possible distinct decimation grids of a given *order* is finite, but it increases very rapidly with the order. As a consequence, in order to make an appropriate selection of the subgrid according to the spectral properties of the signal, not only do we need to define and estimate the spectral extension of the signal, but we need some criteria for limiting the search space of candidate sublattices by ruling out those that do not meet some specific conditions on the spectral extension.

Another problem is the definition of spectral extension according to which both decimation grid and anti-aliasing filter are designed. As a matter of fact, although the spectral energy of non-synthetic 2D signals is usually rather irregularly scattered [6], what we need is a spectral extension model that simplifies the design of the decimation prefilter. As a general rule, it is convenient to select a class of prefilters of interest and choose the spectral extension model accordingly.

If we restrict the class of prefilters we are interested in to those that have convex pass-band, we do not need a detailed description of the spectral extension shape, as we just need to determine the direction around which the spectral energy is maximally concentrated and a measure of the energy dispersion about it. This way of quantifying the anisotropy of the spectral distribution corresponds to approximating the spectral extension with an ellipse whose shape is decided by the ratio between the inertia moments of the power spectrum, while its size is chosen according to the severity of the spectral truncation we are willing to apply.

Adopting a second-order model for the spectral extension seems quite a reasonable choice, as it allows us to detect the dominant direction over which high-frequency portions of the 2D signal are concentrated the most, and adopt decimation grids that exhibit a higher density of samples in that specific direction. According to our experience, any

further extension of the class of prefilters we use would only modestly improve the quality of the signal reconstructed after decimation, and would not justify the heavy complications that would arise from it. Another important problem concerns the design and the implementation of the anti-aliasing filter. The design problems, in fact, crticially depend on the shape of the prefilter's passband.

In this article we propose and test a method for optimally determining the decimation sublattice and designing the relative prefilter. The method, starting from a second-order spectral extension model of the signal, jointly determines the decimation grid and designs a zero-phase 2D FIR prefilter, in a fully automatic fashion [22].

Starting from an estimate of the principal axes of the power spectrum, an upper bound k_{MAX} for the decimation order is determined as the ratio between the area of one period of the spectrum of the original signal and that of the spectral ellipse (i.e. an ellipse whose axes are proportional to the inertia axes). Among all possible decimation grids of order $k \leqslant k_{MAX}$, those that are *compatible* with the elliptical spectral extension are selected. In other words, we keep only the decimation grids that do not give rise to aliasing between *elliptic* truncations of the spectral extension of the signal. If no compatible decimation grids of a certain order exist, then we must repeat the search among grids of lower order until we find some compatible ones. For each compatible decimation grid we generate a hexagonal ideal prefilter by following a fully geometric approach. The best prefilter-grid pair is finally chosen as the one whose prefilter's passband best fits the spectral extension. The geometrical parameters of sublattice and hexagonal prefilter, can now be used for automatically designing a zero-phase FIR hexagonal prefilter that best approximates the ideal one. In order to do so, a new method, which extends the McClellan's frequency transformation technique [16–18], is proposed and tested. With this method we automatically design a frequency transformation that fits the prefilter's hexagonal passband through an entirely geometrical procedure. This way, any hexagonal prefilter can be synthesized from the same 1D zero-phase FIR filter.

The decimation strategy proposed in this article has been implemented in a fully automatic computer procedure and applied to a variety of 2D test signals. In order to make the examples of application as intuitive and visual as possible, we used non-synthetic images with different spectral properties.

In Section 2 of this article we have included all the basic information that is needed for the comprehension of the concepts discussed in the next sections. In particular, a brief summary of the properties of sampling lattices and sublattices is included together with a brief review of concepts of Fourier analysis on lattice structures. The readers who are already familiar with the theory of sampling lattices may ignore this section, while those who would like a more complete introduction to lattice structures, may refer to [1, 2, 8, 9, 15, 24]. The method for the determination of decimation subgrid and prefilter is described in Section 3. The method for designing zero-phase FIR 2D filters with hexagonal pass-band through frequency transformation is proposed and illustrated in detail in Section 4. Finally, some examples of application of the proposed technique are reported in Section 5. Though the notation adopted in this article is the same used in [8, 9], in order to facilitate the reading, we have also included a glossary in Appendix A.

2. Mathematical preliminaries

The decimation grids considered in this article are regular point sets with enough algebraic structure to allow us to define a Fourier transform on them, so that spectral analysis can be performed. Such sampling structure are well-known as sampling *lattices* [8–10, 24], and are characterized by the property of being invariant with respect to translation and reflection [24].

2.1. Lattices and sublattices

The M-dimensional lattice Λ generated by a non-singular matrix $\boldsymbol{A} \in R^M$ is defined as

$$\Lambda = \mathrm{LAT}(\boldsymbol{A}) = \{\boldsymbol{x} \in R^M \mid \boldsymbol{x} = \boldsymbol{A}\boldsymbol{n}, \boldsymbol{n} \in Z^M\},$$

which is the set of all possible linear combinations, with integer coefficients, of the M linearly independent vectors (*basis* of the lattice) represented by the

columns of $\boldsymbol{A}$. Each basis generates a unique lattice, while a lattice may have several bases.

Given a basis $\boldsymbol{A}$, it is possible to derive all other bases of the lattice $\Lambda = \mathrm{LAT}(\boldsymbol{A})$, by using the fact that the generic basis $\boldsymbol{A}'$ can always be written as $\boldsymbol{A}' = \boldsymbol{AU}$, where $\boldsymbol{U} \in Z_M$ is a *unimodular* matrix, i.e. an integer matrix whose determinant has absolute value equal to one (see Appendix B). As a consequence, all bases of a lattice must have the same *modulus* (absolute value of its determinant). The modulus $d(\Lambda)$ of a lattice Λ has a meaningful geometric interpretation, which is related to the definition of its *fundamental cell*.

A fundamental cell $\mathscr{S}$ of an M-dimensional lattice Λ is a (non-necessarily connected) closed region of R^M such that the collection $S_\Lambda = \{\mathscr{S} + \boldsymbol{a}, \boldsymbol{a} \in \Lambda\}$ of all shifted version of $\mathscr{S}$ on all points of the lattice *tiles* the whole space R^M without overlapping between distinct translations. It is not difficult to verify that there exist infinite fundamental cells for a single lattice, but their hypervolume is always the same. Since the hyper-parallelogram $\mathscr{P}$ corresponding to the vectors of the basis $\boldsymbol{A}$ is a fundamental cell of the lattice $\Lambda = \mathrm{LAT}(\boldsymbol{A})$ and its volume is equal to $d(\Lambda)$, we conclude that $d(\Lambda)$ is the volume of all fundamental cells of Λ, therefore $1/d(\Lambda)$ can be interpreted as a measure of the lattice *density*. The fundamental cell is very important in problems of multidimensional sampling and interpolation, and some of its properties will be extensively used in the next sections. Unfortunately, there exist no general geometric classifications of all possible fundamental cells of a given lattice. The only results that are available in the literature concern convex cells [14] and are particularly simple in the two-dimensional case, in which case the only convex regions that tile R^2 are hexagons with central symmetry.[1]

The concept of *sublattice* is particularly important for the decimation problem. The decimation grids we are considering are, in fact, sublattices of the original sampling grid Λ, i.e. subsets of Λ that have a lattice structure. In order to be able to select a suitable sampling grid for a given digital signal, it is of paramount importance to be able to classify and generate all possible sublattices of given order, i.e. those that have a pre-assigned decimation ratio.

Given an M-dimensional lattice $\Lambda = \mathrm{LAT}(\boldsymbol{A})$, the grid $\Gamma = \mathrm{LAT}(\boldsymbol{B})$, is a sublattice of Λ if and only if there exists a non-singular integer matrix $\boldsymbol{H}$ such that $\boldsymbol{B} = \boldsymbol{AH}$. This result [4] is particularly important as it provides us with a method for analytically generating all sublattice bases. The integer number $|\det(\boldsymbol{H})| = |\det(\boldsymbol{B})|/|\det(\boldsymbol{A})|$ represents the ratio between the densities of Λ and Γ, and is often referred to as the *index* of Γ in Λ. This number, which is also written as $(\Lambda : \Gamma)$, corresponds to what we have already called the *decimation ratio* from Λ to Γ. Notice, however, that a lattice may have several bases, therefore the above result is not enough for classification purposes, The problem of automatically generating all k-th order sublattices of a given M-dimensional lattice $\Lambda = \mathrm{LAT}(\boldsymbol{A})$, however, has already been solved [11] and corresponds to finding a partition of the set $\mathscr{W} = \{\boldsymbol{W} \in R_M : \boldsymbol{W} = \boldsymbol{AK}, \boldsymbol{K} \in Z_{M,k}\}$ into distinct classes $\mathscr{W}_i$, each of which identifies a unique sublattice. In order to do so, we can look for a partition of the set $Z_{M,k}$ into distinct classes $\mathscr{K}_i$ of right-equivalent matrices. More specifically, we would like each class $\mathscr{K}_i$ to univocally determine a sublattice $\Gamma_i = \mathrm{LAT}(\boldsymbol{W}_i)$ of Λ through the relationship $\boldsymbol{W}_i = \boldsymbol{AK}, \boldsymbol{K} \in \mathscr{K}_i$. According to the Hermite normal form theorem (see Appendix B), each class of right-equivalent matrices of $Z_{M,n}$ contains one and only one matrix in Hermite normal form. Consequently all kth order sublattice of $\Lambda = \mathrm{LAT}(\boldsymbol{A})$ will be given by $\Gamma_i = \mathrm{LAT}(\boldsymbol{W}_i)$, where $\boldsymbol{W}_i = \boldsymbol{AH}_i$, $\boldsymbol{H}_i \in \mathscr{H}_{M,k}$, where $\mathscr{H}_{M,k}$ is the set of integer matrices in Hermite normal form whose determinant is equal to k.

The number of distinct kth order sublattices of a given lattice Λ corresponds to the number of distinct matrices in Hermite normal form with determinant k, i.e. on the number of possible integer factorization of k. Such a number increases very rapidly with the decimation index k (and the dimension of Λ). For example, there exist 91 distinct sixth-order sublattices of a 3D lattice while, in the 2D case, the number drops down to 12.

[1] Such hexagons can become parallelograms when two opposite sides have zero length.

2.2. Fourier transform

The Fourier transform pair is defined on a lattice $\Lambda = \mathrm{LAT}(\boldsymbol{A})$ as

$$U(\boldsymbol{f}) = \sum_{\boldsymbol{x}\in\Lambda} u(\boldsymbol{x})\mathrm{e}^{-\mathrm{j}2\pi\boldsymbol{f}^{\mathrm{T}}\boldsymbol{x}} = \sum_{\boldsymbol{n}\in Z^M} u(\boldsymbol{An})\mathrm{e}^{-\mathrm{j}2\pi\boldsymbol{f}^{\mathrm{T}}\boldsymbol{An}},$$

$$\boldsymbol{f}\in R^M, \tag{2.1}$$

$$u(\boldsymbol{An}) = |\det(\boldsymbol{A})|\int_{\mathscr{P}} U(\boldsymbol{f})\mathrm{e}^{\mathrm{j}2\pi\boldsymbol{f}^{\mathrm{T}}\boldsymbol{An}}\,\mathrm{d}\boldsymbol{f}, \quad \boldsymbol{n}\in Z^M.$$

The Fourier transform $U(\boldsymbol{f})$ is periodic and its periodicity centers are given by the reciprocal lattice [4, 8] $\Lambda^* = \mathrm{LAT}(\Lambda^{-\mathrm{T}}) = \{y\in R^M : y^{\mathrm{T}}x\in Z, x\in\Lambda\}$. The Fourier transform $U(\boldsymbol{f})$ is thus completely specified by its values in any fundamental cell $\mathscr{P}$ of Λ^*.

2.3. Decimation and interpolation

In the multidimensional case, several different subgrids having the same decimation ratio are available. This fact is a source of complications with respect to the one-dimensional case, but it also offers greater freedom in the decimation setup. Additional complications are caused by the fact that for a single decimation grid, there exist a non-numerable multitude of antialiasing filters, which can always be chosen in such a way to favor some spectral regions instead of others. In order to understand the reason of this increased freedom, we need two operators that will be extensively used in the following sections: the ideal interpolator and the decimator [7].

The interpolation can be though of as the cascade of an elementary interpolator (zero interleaving) and a filter. More specifically, the elementary interpolation of a signal $s(\boldsymbol{x})$, from $\Gamma_1 = \mathrm{LAT}(\boldsymbol{A}_1)$ to $\Gamma_2 = \mathrm{LAT}(\boldsymbol{A}_2)$, Γ_1 being a sublattice of Γ_2, is defined as $w(\boldsymbol{x}) = s(\boldsymbol{x})$ for $\boldsymbol{x}\in\Gamma_1$, and $w(\boldsymbol{x}) = 0$ for $\boldsymbol{x}\in\Gamma_2$ but $\boldsymbol{x}\notin\Gamma_1$. It is not difficult to verify [8] that $W(\boldsymbol{f}) = S(\boldsymbol{f})$ which means that the ideal interpolation does not affect the Fourier transform but just the periodicity lattice, which is now Γ_2^*.

Decimating u from $\Gamma_1 = \mathrm{LAT}(\boldsymbol{A}_1)$ to $\Gamma_2 = \mathrm{LAT}(\boldsymbol{A}_2)$, whew Γ_2 is now a sub-lattice of Γ_1 (i.e. $\Gamma_2\subset\Gamma_1$, and $\boldsymbol{A}_1^{-1}\boldsymbol{A}_2 = \boldsymbol{M}$ unimodular), returns the signal $v = (\boldsymbol{x}) = u(\boldsymbol{x}), \boldsymbol{x}\in\Gamma_2$. In this case the relationship between Fourier transform [8] turns out to be

$$V(\boldsymbol{f}) = \frac{1}{(\Gamma_1:\Gamma_2)}\sum_{\boldsymbol{a}\in\mathscr{I}} U(\boldsymbol{f}+\boldsymbol{a}), \tag{2.2}$$

$\mathscr{I}$ being any Γ_1^* period of Γ_2^*, and $(\Gamma_1:\Gamma_2) = |\det(\boldsymbol{M})|$ the decimation factor.

In order to be able to perfectly reconstruct a signal u defined on Γ_1 from its decimated version v on Γ_2, i.e. in order for the decimation to invertible, it is necessary for the support of $U(\boldsymbol{f})$ to be confined inside some fundamental cell $\mathscr{P}_2$ of Γ_2^*. In this case the reconstruction can be done by using the cascade of an ideal interpolator from Γ_2 to Γ_1 and an ideal filter with frequency response

$$H(\boldsymbol{f}) = \begin{cases}(\Gamma_1:\Gamma_2), & \boldsymbol{f}\in\mathscr{P}_2,\\ 0, & \boldsymbol{f}\in\mathscr{P}_1,\ \boldsymbol{f}\notin\mathscr{P}_2.\end{cases} \tag{2.3}$$

As the impulse response of the filter is defined on Γ_1, the frequency response $H(\boldsymbol{f})$ is a Γ_1^*-periodic of R^2, therefore it is sufficient to specify it in a Γ_1^*-period $\mathscr{P}_2$ of R^2.

3. Decimation approach

As seen in Section 2, decimating a signal on a kth order sublattice causes its spectrum to replicate $k-1$ times, in correspondence to the points of the decimation lattice's reciprocal. In order for the prefilter to prevent aliasing, its passband must be completely contained in a fundamental cell of the decimation grid's reciprocal.

What we need now is a strategy for approaching decimation problem in its globality. As we are interested in reducing spectral redundancy, not only do we need a systematic strategy for modeling the spectral extension of a generic two-dimensional digital signal, but we also need to be able to interpret this information in order to make an optimal selection of the decimation grid and, at the same time, to design the anti-aliasing filter that best fits the spectral extension. Such two operations cannot be considered as separate and independent, therefore they but must be performed jointly.

In general, designing prefilter and sublattice jointly is quite a difficult task whose solution can be kept reasonably simple only under some simplifying assumptions. The design method proposed in this article is based on a strategy presented in a previous work [22], which approaches the problem by using information on the spectral extension for drastically reducing the number of candidate subgrids to consider and geometrically constructing a prefilter for each one of them.

The first step of the procedure consists of estimating the second-order spectral extension of the signal through the evaluation of the inertia axes of its power spectrum. We can then determine an upper bound for the index of the sublattices to choose among, from the spectral occupancy of the estimated spectral ellipse. The maximum decimation index k_0 is the order from which to start looking for suitable decimation grids. Given a decimation order k, all distinct k-th order sublattices are generated and, among them, all those that are *non-compatible* with the spectral extension, i.e. those that cause spectral ellipses to overlap, are ruled out. If no compatible candidate subgrids can be found, then the search will be repeated among sublattices of an inferior order. For each compatible sublattice, we geometrically generate the fundamental cell that best fits the elliptical spectral extension. The sublattice-cell pair that best fits the spectral extension can be finally selected among the remaining candidates.

3.1. Spectral extension estimation

In order to decimate the signal with a minimum loss of information we need to determine the spectral extension of the signal to be decimated. In other words, we need to estimate a spectral region that contains the most significant information on the original signal. Such a region will be used later for the selection of a decimation grid and the design of an antialiasing filter that does not cause overlapping between spectral replicas.

The spectral extension of a 2D discrete signal $u(\boldsymbol{x}), \boldsymbol{x} \in \Lambda$, can be simply defined as the portion $\mathscr{E}$ of one Λ^*-period $\mathscr{P}$ of R^2, where the magnitude of the spectrum $U(\boldsymbol{f})$ of $u(\boldsymbol{x})$ exceeds an assigned threshold S. The *bandwidth*, or *spectral occupancy* of u can thus be defined as $B = \text{area}(\mathscr{E})$, and its normalized value $e = \text{area}(\mathscr{E})/\text{area}(\mathscr{P})$, $0 < e < 1$, can be used as an index of the spectral efficiency.

The determination of decimation grid and fundamental cell that best fit a spectral extension obtained through simple amplitude thresholding is quite a formidable task. In fact, when dealing with non-synthetic signals such as 2D images, that definition of spectral extension usually produces sets $\mathscr{E}$ that are rather complex in shape [6]. The spectral energy of non-synthetic 2D signals, however, is normally concentrated near the origin, and its distribution is often limited to a connected region of the frequency plane. It is thus reasonable to restrict the class of prefilters we are interested in to those that have a *compact* and *convex* passband region. As explained in Section 2.1, the only convex regions that tile R_2 are parallelograms and hexagons with central symmetry, therefore the above assumption greatly simplifies the problem.

In order to determine a compact and convex model of the signal's spectral extension, we just need to determine the direction around which the spectral energy is maximally concentrated (principal axis) and a measure of the energy dispersion about that axis. Quantifying the anisotropy of the spectral distribution through this second-order model corresponds to approximating the spectral extension with an ellipse whose shape is decided by the inertia moments of the power spectrum, while its size is chosen according to the severity of the spectral truncation that we are willing to apply.

Any non-synthetic discrete 2D signal $u(\boldsymbol{x}), \boldsymbol{x} \in \Lambda$, has a limited region of support $Q \subset R^2$, therefore we may consider its Ψ-periodized version $u_\Psi(\boldsymbol{x})$, Ψ being some sublattice of Λ, so that the spectrum $U_\Psi(\boldsymbol{f})$ of $u_\Psi(\boldsymbol{x})$ results as being discrete over Ψ^* and periodical over Λ^*. It should be quite clear that, if the Ψ-periodization is correct, then the frequency sampling associated to it does not cause any loss of information. It is thus possible to derive all spectral information about the signal $u(\boldsymbol{x})$, for example, through DFT computation.

In order to estimate the second-order model of the power spectrum $S(\boldsymbol{f}) = |U(\boldsymbol{f})|^2$, we need to determine the principal inertia axes d_1 and d_2 and the relative radii of gyration ρ_1 and ρ_2 [25]. The

axes d_1 and d_2 are determined from the eigenvectors of the inertia matrix (matrix of the inertia moments), while the radii $\rho_1 = \sqrt{I_1/M}$ and $\rho_2 = \sqrt{I_2/M}$ depend on the principal inertia moments I_1 and I_2 (eigenvalues of the inertia matrix) and the total energy M of the 2D signal.

With respect to the reference frame (d_1, d_2), the inertia ellipse is described by the equation $\rho_1^2 d_1^2 + \rho_2^2 d_2^2 = 1$ therefore the second-order model of the spectral extension is a set of the form $\rho_1^2 d_1^2 + \rho_2^2 d_2^2 \leqslant r^2$, r being a scaling factor for the radii of gyration, to be chosen in order for such a spectral ellipse to contain the most significant portion of the power spectrum (e.g. by selecting a threshold for the signal energy included in the ellipse). The spectral occupancy of the signal, i.e. the area $A = \pi\rho_1\rho_2/r^2$ of the spectral ellipse, provides us with an upper bound for the index k of the sublattices that could be used for decimating the signal. In fact, we must have $k \leqslant k_0$, where k_0 is the largest integer which is not greater than area $(\mathscr{P})/A$ and $\mathscr{P}$ is any fundamental cell of the reciprocal sampling lattice.

3.2. Determination of prefilter and decimation grid

As already seen in Section 2, the problem of determining all kth order sublattices of a lattice Λ, with $k \leqslant k_0$, can be solved by generating all matrices in Hermite normal form [6] that have determinant equal to k. For each one of such subgrids we need to determine the *compatible* ones, i.e. those that do not cause the above spectral ellipses to overlap, and the relative ideal prefilters [22].

3.2.1. Compatible sublattices

The fact that the spectral extension model is elliptical makes the compatibility check particularly simple. In fact, given an ellipse, the curve described by the centers of all of its tangent replicas is itself an ellipse whose radii of gyration are twice the ellipse's radii. This *threshold ellipse* γ_{th}, described by the equation $(2\rho_1)^2 d_1^2 + (2\rho_2)^2 d_2^2 = r^2$, represents the limit region for the replication centers beyond which no overlapping occurs, therefore we just need to make sure that all points of Γ^* lie outside γ_{th}.

Letting $(\bar{d}_1, \bar{d}_2)$ be the coordinates of the generic point of Γ^*, referred to the principal axes of the ellipse (see Fig. 1), the compatibility check becomes

$$(2\rho_1)\bar{d}_1^2 + (2\rho_2)\bar{d}_2^2 > r^2.$$

3.2.2. Prefilter's geometry

The set of all compatible subgrids is much more limited than the set of all possible sublattices of the order $k \leqslant k_0$. All such grids are good candidates for decimation as none of them causes the spectral ellipse's replica to overlap. The choice will then be made according to the prefilter that will be associated to each one of them. In order to do so, we need a fast method for generating a fundamental cell that 'well fits' the spectral extension, for each compatible sublattice. More specifically, the compact and convex fundamental cells that we are looking for must enclose the elliptical spectral extension entirely, and the principal directions of the prefiltered spectrum must be as close as possible to those of the original spectrum

According to the results of Section 2, all convex fundamental cells of two-dimensional lattice are hexagons with central symmetry. A method for determining a hexagonal fundamental cell of a given *compatible* sublattice which entirely encircles the elliptical spectral extension, is described in Fig. 1. We begin by determining the six pairwise opposite points a, b, c, a', b', c' of Γ^* that lie closest to the

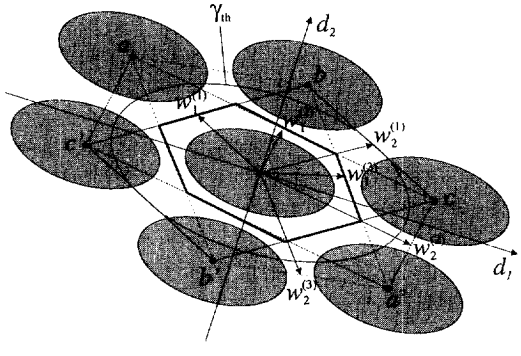

Fig. 1. The hexagonal passband is given by the intersection of the three parallelograms $\boldsymbol{W}_1 = [\boldsymbol{w}_1^{(1)} \quad \boldsymbol{w}_2^{(1)}] = \{b, c, b'\, c'\}$, $\boldsymbol{W}_2 = [\boldsymbol{w}_1^{(2)} \quad \boldsymbol{w}_2^{(2)}] = \{a, c, a'\, c'\}$, $\boldsymbol{W}_3 = [\boldsymbol{w}_1^{(3)} \quad \boldsymbol{w}_2^{(3)}] = \{a, b, a'\, b'\}$, built on the six points of Λ^* that lie the closest to the threshold ellipse γ_{th}.

threshold ellipse. We then build three parallelograms on the three quadruples of points that can be obtained by excluding an opposite pair of points, as shown in Fig. 1. For example, the parallelogram $\{b, c, b' c'\}$ is obtained by ruling out the pair $\{a, a'\}$ of opposite points. This polygon can be geometrically described by the vectors $\boldsymbol{w}_1^{(1)}$ and $\boldsymbol{w}_2^{(2)}$ that connect the origin with the middle points of two non-opposite parallelogram's sides. In matrix form, the three parallelograms are represented by the 2×2 matrices

$$W_i = [\boldsymbol{w}_1^{(i)} \quad \boldsymbol{w}_2^{(i)}], \quad i = 1, 2, 3. \tag{3.1}$$

As we can see from Fig. 1, the intersection of the three parallelogram $\boldsymbol{W}_i$, $i = 1, 2, 3$, is a fundamental cell that well fits the spectral ellipse and can be used as a reference for constructing both the anti-aliasing and the reconstruction filters. With this choice, the shape of the refilter's passband tends to capture the spectral anisotropy in a natural fashion.

3.2.3. Final selection of the decimation set-up

Last step of the decimation procedure consists of choosing the best decimation setup among the compatible subgrid/refilter pairs. Although the available candidates are all suitable for decimation, we need a criterion of optimality that takes into account how well the fundamental cell fits the spectral extension. As the computational complexity of the prefilter is expected to be kept reasonably modest, it would be desirable to have a transition band which is as uniform as possible. An operative way of applying the above criterion consists of selecting the sublattice whose prefilter has minimum impact on the principal axis of the power spectral distribution. More precisely, we can choose the prefilter that minimizes the angle between the principal axis of the non-prefiltered spectrum and that of the prefiltered spectrum.

4. Prefilter's design

As already explained in Section 3, the prefilter's passband is determined as an intersection of three parallelograms whose matrix representation (Eq. (3.1)) constitute a complete geometrical description of the whole decimation setup (subgrid and reference refilter) and, in fact, represent the only output of the procedure described in Section 3.

In order to design a low-pass anti-aliasing filter with the above hexagonal passband, several strategies can be adopted. It would be highly desirable, however, to define and adopt a design strategy that, starting from just the geometrical description of the prefilter, generates the desired filter in a totally unsupervised and automatic fashion.

The simplest solution to the prefilter's design problem is represented by spectral windowing based on FFT computation [23]. Considering the fact that the 2D test signals we selected are images, however, a few considerations are in order. The ideal prefilter determined in Section 3 exhibits sharp transitions at the border of the passband region. As a consequence, abrupt spectral truncation may give rise to an undesired 'ringing' effect, which is highly visible as it results in artifacts that are parallel to the image edges. It is well-known, in fact, that our visual system reacts to noise and disturbance in highly anisotropic way [12]. For example, the disturbance components that are parallel to edges are much more visible than those that are perpendicular to it. In principle, one should design a prefilter taking into account the impulse response as well as the frequency response [3, 20, 26]. A simple solution consists of adopting a smoothed version of the ideal hexagonal prefilter, although smoothing the filters's transitions is not a straightforward procedure.

It is also important to keep in mind that computing the 2D FFT of a limited-support signal defined on the lattice Γ, requires signal periodization over a lattice Ψ which is a sublattice of the decimation grid Λ and, consequently, of Γ. Failing to perform zero-padding prior the computation of the FFT may result in the generation of undesirable artifacts.

The major drawback of spectral windowing is the fact that it forces us to work in the frequency domain. This lack of flexibility motivates us in searching for a design solution which allows us to choose between an implementation in the image domain and one in the frequency domain.

The design strategy that we propose in this section can be seen of as an extension of the McClellan's frequency transformation method [13, 16, 17]. As we will see, our technique uses just one pre-designed 1D zero-phase FIR filter and

geometrically transforms it into a 2D zero-phase FIR filter that meets all the requirements specified in Section 3. In order to do so, the frequency transformation is automatically designed through a fully geometrical procedure in such a way to generate the desired hexagonal passband.

4.1. Frequency transformation

It is well-known that a frequency transformation of the form

$$\cos(2\pi f\Delta) = T(\boldsymbol{f}), \tag{4.1}$$

Δ being the 1D sampling step, can be used for mapping the frequency response $H_{1D}(f)$ of a 1D zero-phase FIR filter onto a 2D zero-phase FIR filter as follows:

$$H(\boldsymbol{f}) = \sum_{i=0}^{N} a_i \cos^i(2\pi f\Delta)\Big|_{\cos(2\pi f\Delta)=T(\boldsymbol{f})} = \sum_{i=0}^{N} a_i T^i(\boldsymbol{f}). \tag{4.2}$$

According to Eq. (4.2), Eq. (4.1) maps the 1D frequency value $\bar{f}$ onto an iso-level curve described by the implicit equation $T(\boldsymbol{f}) = \bar{c}$, where $\bar{c} = \cos(2\pi\bar{f}\Delta)$. The idea behind the method we propose and the goal of this section is to determine a frequency transformation function $T(\boldsymbol{f})$ that, starting from the matrix description $\boldsymbol{W}_i$, $i = 1, 2, 3$, of the three parallelograms that constitute the hexagon, maps $\bar{f} \in [0, 1/(2\Delta)[$ onto a closed curve that lies entirely inside the hexagonal passband region, while th Nyquist frequency $f_N = 1/(2\Delta)$ is mapped onto the whole stopband region. With this choice, a 1D zero-phase FIR filter having a zero at the Nyquist frequency would be mapped onto a 2D filter which prevents aliasing from occurring in the decimation process. Furthermore, all geometrical parameters of the sublattice/prefilter pair would be totally embedded into the transformation map, therefore a single 1D filter could be mapped onto any of the 2D prefilters of the type defined in Section 3.

Notice that, because of the flat response in the stopband region, such an ideal transformation function cannot be found in FIR form, therefore we will have to look for an approximate FIR solution.

In order to simplify the description of the hexagonal prefilter's design method, we will first consider the synthesis of prefilters whose passband is shaped like a parallelogram and then will proceed with the hexagonal case.

4.1.1. Transformation map for parallelograms

Let us consider the classical McClellan's transformation map

$$T_m(\boldsymbol{\xi}) = -1 + \tfrac{1}{2}(1 + \cos(2\pi\xi_1\Delta_1)) \times (1 + \cos(2\pi\xi_2\Delta_2)), \tag{4.3}$$

of Fig. 2(b) and 2(c). Apart from a scale factor and a range shift, the map (Eq. (4.3)) is separable, therefore it can be written as a product of the form

$$2(1 + T_m(\boldsymbol{\xi})) = F_1(\xi_1)F_2(\xi_2),$$

where $F_1(\xi_1) = 1 + \cos(2\pi\xi_1\Delta_1)$ and $F_2(\xi_2) = 1 + \cos(2\pi\xi_2\Delta_2)$. Equivalently, it can also be written as a periodic convolution of the form

$$2(1 + T_m(\boldsymbol{\xi})) = G_1(\boldsymbol{\xi}) \circledast G_2(\boldsymbol{\xi}), \tag{4.4}$$

where

$$G_1(\boldsymbol{\xi}) = F_1(\xi_1)\sum_k \delta\left(\xi_2 - \frac{k}{\Delta_2}\right),$$

$$G_2(\boldsymbol{\xi}) = F_2(\xi_2)\sum_k \delta\left(\xi_1 - \frac{k}{\Delta_1}\right).$$

As the extension of the convolution does not exceed the Γ^*-period, we can limit our analysis to one square period. In order to do so, we define the aperiodic *blade* function $\bar{G}_1(\boldsymbol{\xi}) = \bar{F}_1(\xi_1)\delta(\xi_2)$ and $\bar{G}_2(\boldsymbol{\xi}) = \bar{F}_2(\xi_2)\delta(\xi_1)$, which are oriented along ξ_1 and ξ_2, respectively, and whose profile is shaped as

$$\bar{F}_j(\xi_j) = \begin{cases} F_j(\xi_j), & -\frac{1}{2\Delta_j} \leqslant \xi_j < \frac{1}{2\Delta_j}, \\ 0, & \text{elsewhere}, \end{cases} \qquad j = 1,2,\ldots,$$

as shown in Fig. 2(c). The linear convolution of such blade functions results in

$$\bar{T}_m(\boldsymbol{\xi}) = -1 + \tfrac{1}{2}\bar{G}_1(\boldsymbol{\xi}) * \bar{G}_2(\boldsymbol{\xi}),$$

which needs to be Γ^*-periodized in order to return $T_m(\boldsymbol{\xi})$.

Let us now consider the linear transformation $\boldsymbol{f} = \boldsymbol{W}\boldsymbol{\xi}$ of the frequency plane, where $\boldsymbol{W}$ is a 2×2

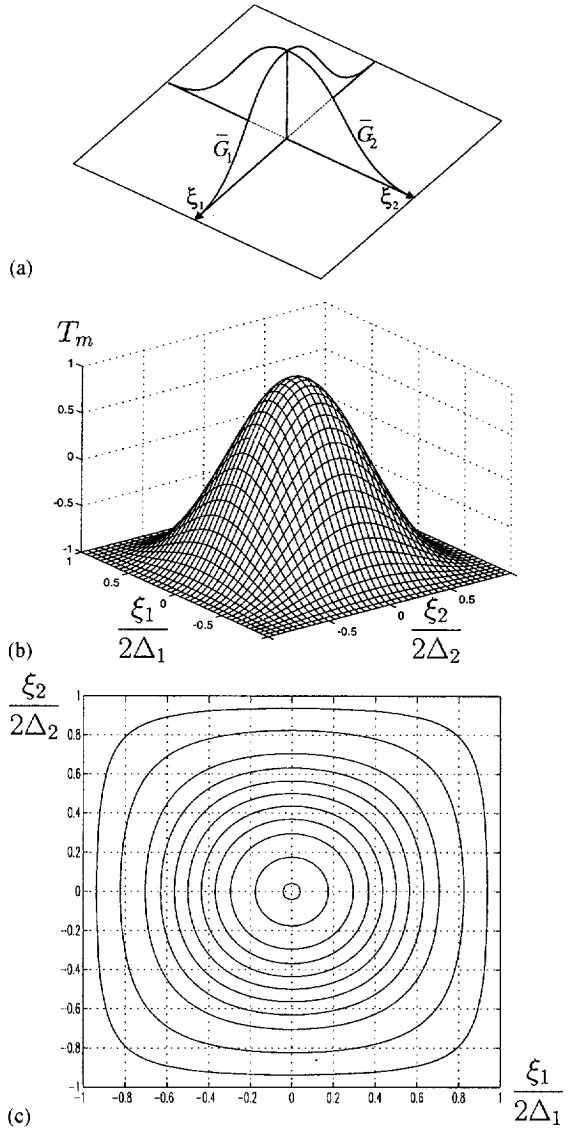

Fig. 2. McClellan's frequency transformation map. The *blade* functions (a) are used for generating the 2D profile (b) with the iso-level curves, (c) through convolution.

matrix such that $|\det(\boldsymbol{W})| < (\Delta_1\Delta_2)^{-1}$. The matrix $\boldsymbol{W}$ characterizes a parallelogram (its columns represent the vectors that connect the origin with the centers of two non-parallel sides of the parallelogram) which we assume to be contained in the Γ^*-period. The transformation map $\boldsymbol{f} = \boldsymbol{W}\boldsymbol{\xi}$ stretches the McClellan's iso-level curves in such a way to fit inside this parallelogram.

It is not difficult to show that, if $\bar{G}(\boldsymbol{\xi}) = \bar{G}_1(\boldsymbol{\xi}) * \bar{G}_2(\boldsymbol{\xi})$, then

$$\bar{G}(\boldsymbol{W}_i^{-1}\boldsymbol{f}) = \frac{1}{\det(\boldsymbol{W}_i)}\,\bar{G}_1(\boldsymbol{W}_i^{-1}\boldsymbol{f}) * \bar{G}_2(\boldsymbol{W}_i^{-1}\boldsymbol{f}),$$

therefore the transformation map $T_i(\boldsymbol{f})$ that fits the parallelogram $\boldsymbol{W}_i$, $i = 1, 2, 3$, can be obtained through Γ^*-periodization of

$$\bar{T}_i(\boldsymbol{f}) = -1 + \frac{1}{2\det(\boldsymbol{W}_i)}\,\bar{G}_1(\boldsymbol{W}_i^{-1}\boldsymbol{f}) * \bar{G}_2(\boldsymbol{W}_i^{-1}\boldsymbol{f}),$$

$$i = 1, 2, 3.$$

In conclusion, apart from a scaling and a range shift, a transformation map that fits the parallelogram $\boldsymbol{W}$ can be expressed as a convolution of two blade functions shaped like raised cosines and nonzero along the lines that connect the centers of opposite sides.

The above representation of a parallelogram-shaped transformation map is not very useful from the operative standpoint. In fact, the above method implies frequency de-periodization (space interpolation on $\mathscr{R}^2$) followed by frequency (space) warping and, finally, frequency periodization (space resampling over Γ). Furthermore, the presence of a flat region in the Γ^*-period (outside the parallelogram) prevents the transformation mask from having finite extension. The above discussion, however, will become useful later for defining a hexagonal transformation map.

In order to define an operative procedure for constructing a parallelogram-shaped zero-phase FIR transformation map, let us apply the linear transformation $\boldsymbol{f} = \boldsymbol{W}\boldsymbol{\xi}$ to the axes of the periodic McClellan's transformation map of Fig. 3(a). The result of this change of reference frame is shown in Fig. 3(c), where it is quite apparent that the iso-level curves of Fig. 3(a) are now all mapped onto the inside of the parallelogram $\boldsymbol{W}$. Notice that, in order for this result to be correct, we need to make sure that Γ^* is a sublattice of $\Upsilon = \mathrm{LAT}(\boldsymbol{W})$, so that the new transformation mat $T_w(\boldsymbol{f}) = T_m(\boldsymbol{W}^{-1}\boldsymbol{f})$ is also Γ^*-periodical.

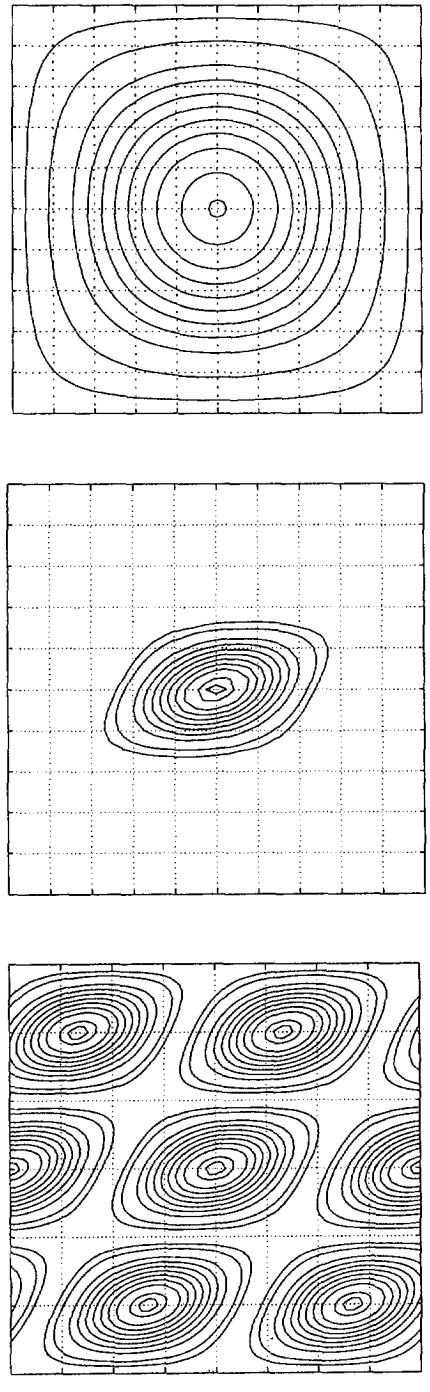

Fig. 3. The McClellan's frequency transformation map (a) can be used for deriving map (b) by removing all undesired replicas from map (c). Map (c) can be directly obtained from (a) through linear transformation of the reference frame.

If $t_m(\boldsymbol{y}), \boldsymbol{y} \in \Gamma$, is the transformation mask associated to $T_m(\boldsymbol{\xi})$ (see Fig. 3(c)), then the transformation mask associated to $T_w(\boldsymbol{f}) = T_m(\boldsymbol{W}^{-1}\boldsymbol{f})$ is

$$t_w(x) = \begin{cases} t_m(\boldsymbol{W}^{-\mathrm{T}}\boldsymbol{x}), & \boldsymbol{W}^{-\mathrm{T}}\boldsymbol{x} \in \Lambda, \\ 0, & \boldsymbol{x} \in \Gamma,\ \boldsymbol{x} \notin \Lambda, \end{cases}$$

which is obtained from $t_m(\boldsymbol{y})$ by redistributing its samples over the sublattice Λ while leaving all other samples of Λ to zero.

The transformation map that we are looking for can be derived from $T_w(\boldsymbol{f})$ by eliminating the undesired replicas of the parallelogram $\boldsymbol{W}$, as shown in Fig. 3(b). This operation corresponds to ideally interpolating $t_w(\boldsymbol{x})$ on Γ. The resulting transformation function $T(\boldsymbol{f})$ maps $\bar{f} \in [0, f_N[$ onto one of a family of closed curves that lie entirely inside the parallelogram $\boldsymbol{W}$ and tend to approximate its border as $\bar{f}$ approaches f_N. The Nyquist frequency f_N, instead, is mapped onto both the border and the outside of this parallelogram, as required.

It is important to notice that, in order for our design method to be applicable, the final transformation map $t(\boldsymbol{x})$ must have limited support. A mask obtained through ideal interpolation of $t_w(\boldsymbol{x})$, however, would not be FIR therefore a non-ideal FIR interpolation would be in order.

In alternative to interpolation, we propose to derive $t(\boldsymbol{x})$ from a properly chosen analog window $t_a(\boldsymbol{x}) = h(\boldsymbol{W}^{-\mathrm{T}}\boldsymbol{x})$, where $h(\boldsymbol{x})$ is any analog window such that $h(\boldsymbol{x}) = t_m(\boldsymbol{x})$ for $\boldsymbol{x} \in \Gamma, \boldsymbol{x} \neq 0$. A possible choice for $h(\boldsymbol{x})$ is a separable analog Hanning window, whose amplitude range and whose sample at the origin are modified in such a way that its Fourier transform ranges from -1 to 1, as required by Eq. (4.1). One apparent advantage in constructing the transformation map through a sampling of an analog function is the fact that it allows us to remove the hypothesis according to which $\Gamma^* \subset \Upsilon = \mathrm{LAT}(\boldsymbol{W})$, thus giving us the freedom of arbitrarily scaling and/or modifying the shape of the parallelogram.

4.1.2. Transformation map for hexagons

Our goal is to build a transformation map that decreases monotonically from 1 to -1 as it approaches the border of a regular hexagon with central symmetry. Similarly to what seen in the

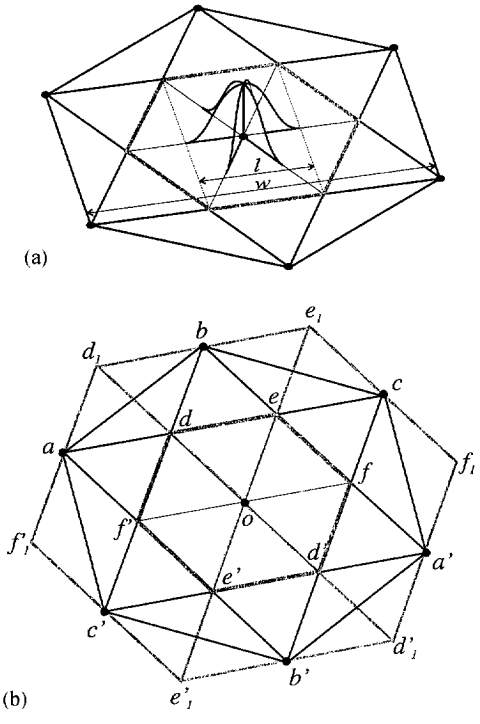

Fig. 4. (a) Construction of the hexagonal frequency transformation map. The three *blade* functions are aligned with the vertex lines of the hexagon. The extension l of any blade function is half the length of the vertex line which, in turn, is 2/3 the length w of the corresponding parallelogram. (b) Convolving two blade functions whose extension are the vertex lines $\{d, d'\}$ and $\{e, e'\}$ results in a transformation map whose extension is the parallelogram $\{b, f, b', f'\}$. Convolving this map with a blade function whose support is the vertex line $\{f, f'\}$ results in the external hexagon $\{d_1, e_1, f_1, d'_1, e'_1, f_1\}$. The vertex lines of this hexagon are twice those of the desired one $\{d, e, f, d', e', f'\}$.

previous section, this function can be constructed through the appropriate convolution of *blade* functions. In particular, with reference to Fig. 4(a), a proper hexagonal transformation map can be obtained by convolving together three blade functions oriented along the three lines that connect opposite vertices of the hexagon (vertex lines) and whose length are half such lines. Their shape is, once again, given by raised cosine functions. In order to verify that this procedure gives the correct result, it might be easier to determine the result of the convolution of blade functions whose extension are the whole vertex lines of the hexagon. The result will be a hexagon whose vertex lines are twice the desired ones, as shown in Fig. 4(b).

In order to construct the desired hexagonal transformation map we thus need to construct three FIR transformation masks which approximate the Γ^*-periodized blade functions in the frequency domain. In fact, convolving together their Fourier transform corresponds to multiplying them together in the space domain.

With reference to Fig. 4, we notice that each one of the three parallelograms

$$\boldsymbol{W}_1 = [\boldsymbol{w}_1^{(1)} \quad \boldsymbol{w}_2^{(1)}] = \{b, c, b', c'\},$$
$$\boldsymbol{W}_2 = [\boldsymbol{w}_1^{(2)} \quad \boldsymbol{w}_2^{(2)}] = \{a, c, a', c'\},$$
$$\boldsymbol{W}_3 = [\boldsymbol{w}_1^{(3)} \quad \boldsymbol{w}_2^{(3)}] = \{a, b, a', b'\},$$

which intersect in the desired hexagon, exhibit a dominant direction of extension along a hexagon's vertex line. This fact is not incidental and can be used for determining a FIR approximation of the blade functions.

Let $\hat{T}_i(\boldsymbol{f}), i = 1, 2, 3$, be the parallelogram-shaped frequency transformation associated to the scaled parallelograms $\hat{\boldsymbol{W}}_i = \alpha \boldsymbol{W}_i$, $i = 1, 2, 3$. With reference to Fig. 4, it should be quite clear that the extension of any of the three blade functions is 2/3 the longest one of the two vectors that generate the relative parallelogram. As a consequence, $\hat{T}_i(\boldsymbol{f}), i = 1, 2, 3$, represent the desired approximations of the blade functions when $\alpha = 2/3$, and can be derived by following the procedure of Section 4.1.1. It is important to emphasize the fact that the scaling factor $\alpha = 2/3$ is fixed, and does not depend on the geometry of the sublattice/prefilter pair.

Notice that convolving together three parallelogram-shaped maps, rather than blade functions, has the desirable effect of reducing the residual ripple of the final map, caused by imperfect interpolation of the warped McClellan mask. Furthermore, convolving together the three frequency maps corresponds to multiplying together the three corresponding masks, with the result of limiting the extent of the final transformation mask.

4.2. *One-dimensional filter design*

One problem that still needs to be solved is how to choose the specifications of the zero-phase 1D

FIR filter to be transformed. As already explained above, the suppression of undesired replicas depends on how well the frequency transformation map flattens at -1 in the stopband region. Because of the resampling of the analog window on Γ, besides depending on the choice of window, stopband suppression depends also on the alignment of the hexagon's sides with the principal directions (lines of points) of Γ^*. In general, however, the stopband's residual ripple is never greater than the sidelobes of the Fourier transform of the adopted analog window t_a. The iso-level curves associated to the frequency transformation will not invade the stopband region as long as the level is above the residual ripple.

Let f_{th} be the threshold level beyond which the iso-level curves invade the stopband region. In order to guarantee the unwanted replicas to be suppressed, we need to make sure that the stopband region of the 1D filter includes the interval $[f_{\text{th}}, f_N[$. In the rest of the frequency band, we would like the 1D filter to have a response which is as flat as possible around 1.

In general, the 1D filter needs to be designed according to the specific type of 2D signal that we are dealing with. In the case of images, for example, it is preferable for its frequency response not to exhibit sharp transitions, in order to limit consequent ringing phenomena.

4.3. *Some practical considerations*

The specifications determined above are quite restrictive for a FIR filter. In order to smoothen them down a little, we can introduce a *transition band* before f_{th}, which requires further scaling of the matrices $\bar{W}_i$ in order to obtain the desired passband's size. Such scaling, however, needs to be balanced with the loss of accuracy with which the passband's shape will approximate a hexagon. It is very important to remind, however, that the above additional scaling factor, as well as α (see Section 4.1.2), are fixed, in the sense that they do not depend on the geometry of the sublattice/prefilter pair. This means that, changing the sublattice will result in a change of just the transformation mask, with no tuning to perform.

As far as the computational complexity is concerned, spectral windowing benefits from the availability of fast algorithms for the computation of the DFT. Conversely, direct convolution is quite a costly operation. However, the fact that the prefilter is immediately followed by an ideal sampler suggests us that the only the samples that lie on Λ should be computed, with the result that the convolution's complexity is reduced of a factor $[\Gamma : \Lambda]$. In conclusion, the complexity of the two strategies is comparable, while the advantage of the one based on the frequency transformation method over the FFT-based strategy are quite apparent. The former is fully geometrical, and requires the synthesis of just one 1D FIR filter as all geometrical properties of the prefilter/sublattice pair are incorporated into the transformation mask. In fact, if the sublattice/prefilter pair changes, then we just need to replace the mask while leaving the 1D filter unchanged.

5. Example of application

As already mentioned in the introduction, the search for a decimation grid and the generation of a prefilter according to the method proposed in this article has been implemented into a completely automatic computer procedure and tested over a series of real images.

A first example of application is represented by the test image *Train*, shown in Fig. 5(a). The spectrum of this image, shown in Fig. 5(b), exhibits a certain anisotropy, as confirmed by the orientation and the elongation of the spectral ellipse. This fact can mainly be attributed to the prevalence of detail aligned with the body of the train. The principal axes of the elliptical spectral extension have been chosen to be 3.5 times the inertia axes of the power spectrum samples. The maximum order of decimation in which some compatible sublattices can be found is $k = 12$. Among all hexagonal fundamental cells (built by using the method of Section 3.2.2) that are associated to the compatible 12th order sublattices, the one whose principal axes are closest to those of the elliptical spectral extension is chosen to define the passband region of the image prefilter (see Fig. 5(c)),

while the relative sublattice, whose basis (in Hermite normal form) is

$$V = \begin{bmatrix} 12 & 8 \\ 0 & 1 \end{bmatrix},$$

is the corresponding decimation grid.

Decimating the image over the selected subgrid causes the truncated power spectrum to replicate like in Fig. 5(d), where it is quite apparent how the elongation of the power spectrum due to the prevalence of some edges aligned along a specific direction, causes subgrid and prefilter to preserve the spectrum in that direction.

The geometrical description of the reference hexagonal filter of Fig. 5(c) is represented by the three parallelograms

$$W_1 = \frac{1}{24}\begin{bmatrix} 3 & -1 \\ 0 & 8 \end{bmatrix}, \quad W_2 = \frac{1}{24}\begin{bmatrix} 1 & -5 \\ 4 & 4 \end{bmatrix},$$

$$W_3 = \frac{1}{24}\begin{bmatrix} 4 & -2 \\ 4 & 4 \end{bmatrix}.$$

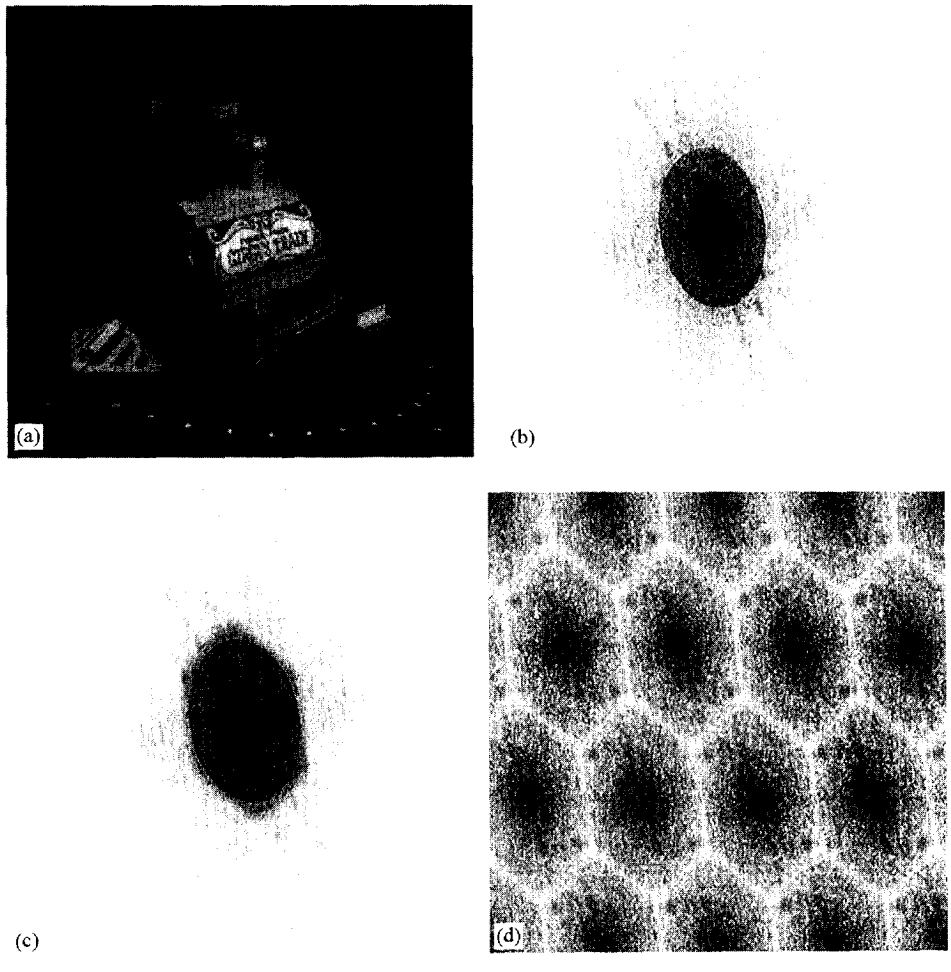

Fig. 5. Geometric characterization of prefilter and decimation grid for the test image *train*. (a) Original image, (b) spectrum and spectral ellipse, (c) hexagonal passband, (d) tiled prefiltered spectrum.

The blade functions that characterize the hexagonal transformation map are approximated by the frequency responses of three zero-phase FIR filters, obtained with reference to the above parallelograms through the method explained in Section 4.1.1. The iso-level curves of the Fourier transforms of such approximated maps are shown in Fig. 6(a)–6(c). Such maps can be convolved together in order to obtain the global transformation map of Fig. 6(d) and 6(e).

By using the transformation map of Fig. 6, the 1D FIR filter whose frequency response is shown in Fig. 7 is mapped onto the 2D FIR filter of Fig. 8. Notice from Fig. 8(c) that the passband's replicas of the resulting filter (obtained by decimating its impulse response over Λ) tile the frequency plane

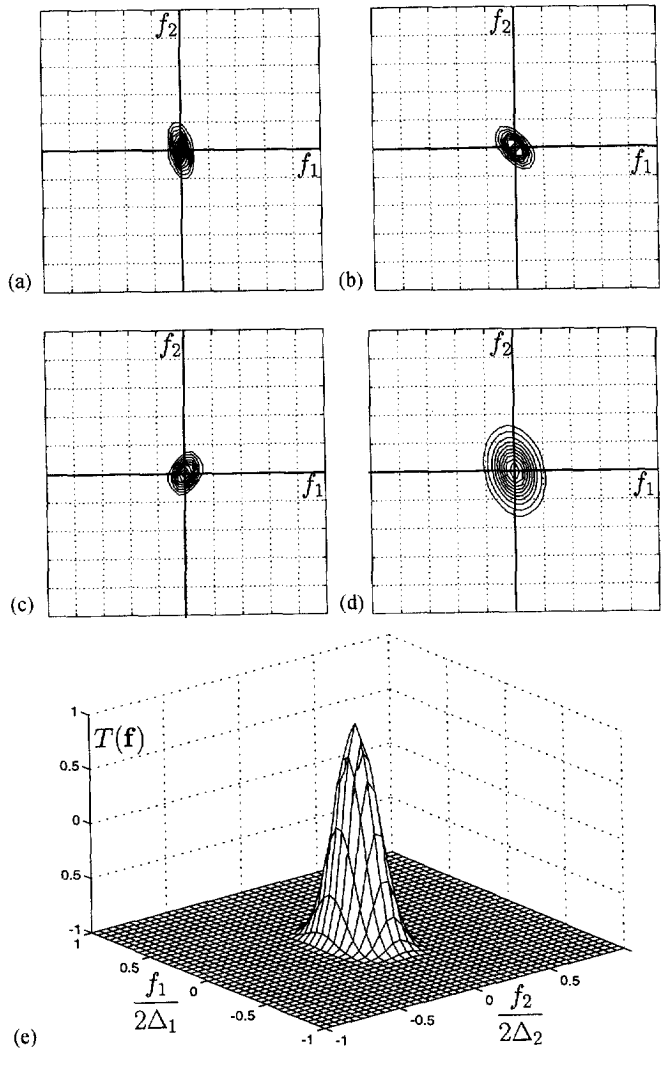

Fig. 6. Construction of the frequency transformation map for the test image *train*. The Fourier transforms of the three FIR approximations of the blade functions (a), (b) and (c) are convolved together to obtain the global transformation map (d) and (e).

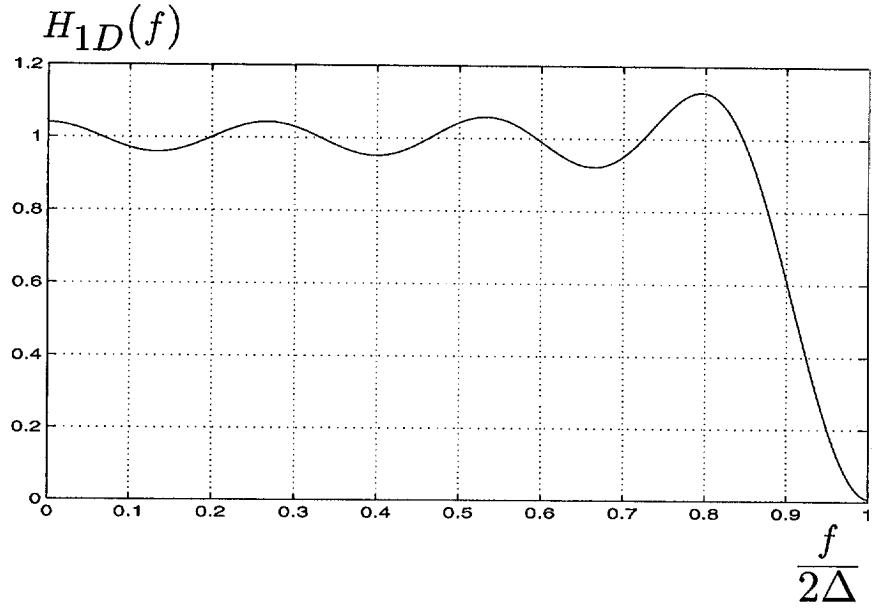

Fig. 7. Frequency response of the zero-phase 1D FIR filter used for synthesizing the desired 2D prefilters.

perfectly. The anti-aliasing filter can also be used for reconstructing the original image from the decimated one.

The prefiltered image is shown in Fig. 9(a). Such results are shown on the original sampling lattice only because of obvious visualization limits. A comparison between corresponding zoomed-in details of original and prefiltered image are shown in Fig. 9(b) and 9(c). As we can see, the blurring due to the low-pass anti-aliasing filtering is still acceptable, considering the reduction of 12 times in the amount of samples that can actually be used for describing the image itself. The residual ringing can be further reduced by selecting a 1D filter specifically designed for that purpose.

Another reconstruction example is shown in Figs. 10–12. In this case the spectral extension exhibits a different orientation in the spectral anisotropy and the decimation factor is $k = 15$. Notice that the prefilter of Fig. 11 has been obtained from the same 1D filter used for the previous example. As already emphasized in Section 4, the anti-aliasing filter can also be used for reconstructing the original image from the decimated one. In both cases, however, the reconstructed image is identical to the prefiltered one as no significant aliasing is introduced by the decimation process.

6. Conclusions

In this article we presented a new technique for decimating discrete 2D signals, which is capable of considerably reducing the spectral redundancy while suppressing the least amount of spectral energy. Spectral characteristics of the signal, such as spectral extension shape and spectral occupancy, are taken into account for determining both decimation grid and anti-alias filter. In particular, the spectral extension has been modeled as a second-order energy distribution through its principal axes and inertia moments of the power spectrum. This choice corresponds to approximating the spectral energy distribution with an ellipse, whose principal axes correspond to the radii of gyration of the spectrum, and allows us to considerably simplify

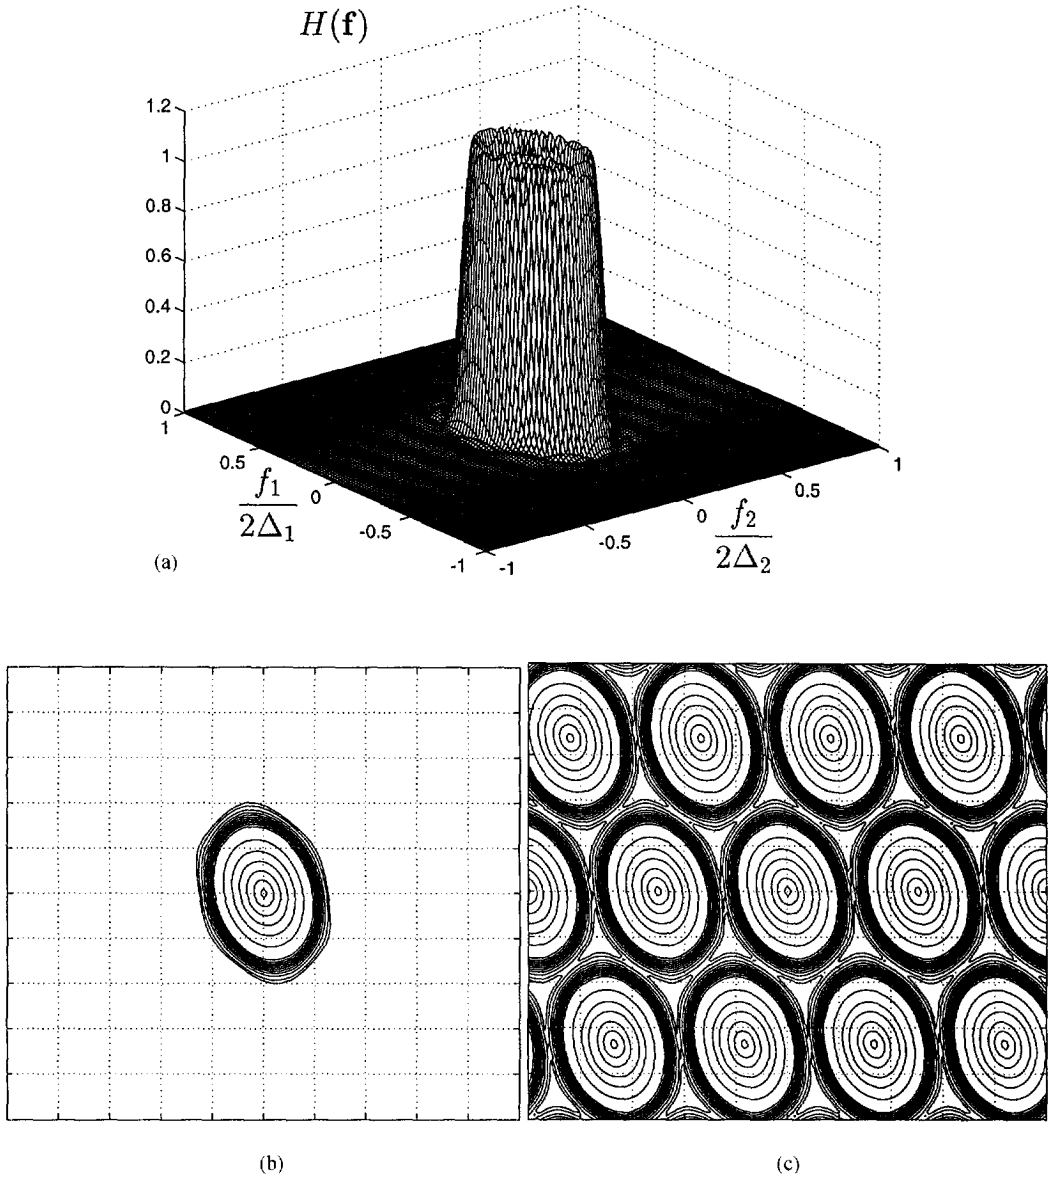

Fig. 8. 2D FIR prefilter for the test image *train*, constructed through frequency transformation. Frequency response (a) and corresponding iso-level curves (b). Fourier transform (c) of the impulse response, after decimation over Λ.

the structure of the prefilter. As a consequence, the procedure for jointly designing subgrid and prefilter becomes purely geometrical and of immediate application.

We have implemented out decimation technique into a fully-automated computer procedure which, after having analyzed the spectral content of a discrete 2D signal, generates all sublattices that are

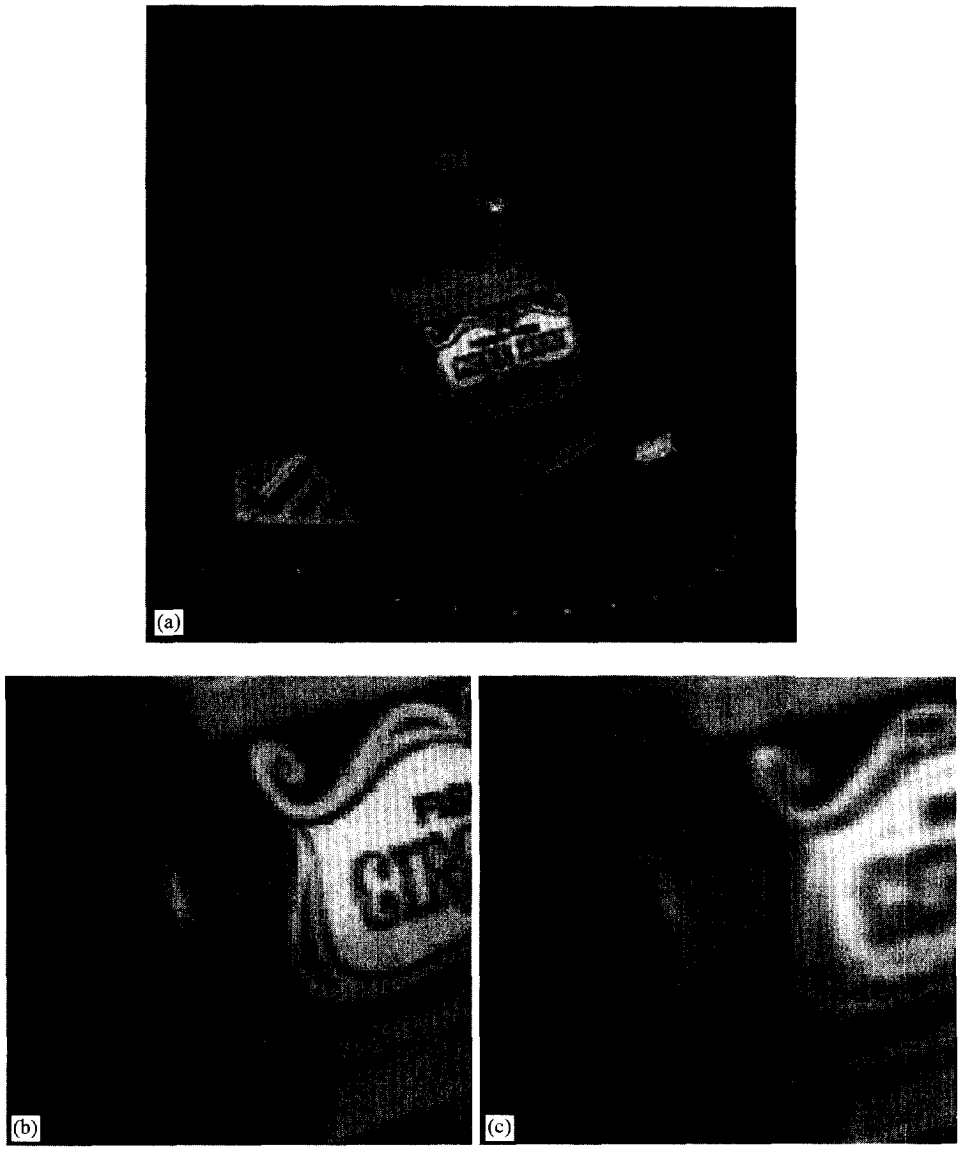

Fig. 9. Prefiltered version of the test image *train*. (a) Corresponding zoomed-in details in the original (b) and prefiltered (c) images.

compatible with its spectral extension, finds the corresponding fundamental cells, and selects the prefilter/subgrid pair that best fits the estimated spectral extension. In addition, the procedure synthesizes a 2D zero-phase FIR filter that corresponds to the previously determined ideal prefilter, by means of a special application of the frequency transformation method. This last methods is entirely geometrical and uses all geometrical information about subgrid/prefilter for constructing a proper transformation mask.

The proposed method has been proven to be quite effective for reducing the spectral redundancy of a 2D signal. In fact, after having applied it on a series of test images with various spectral content, we have shown that we can typically reach

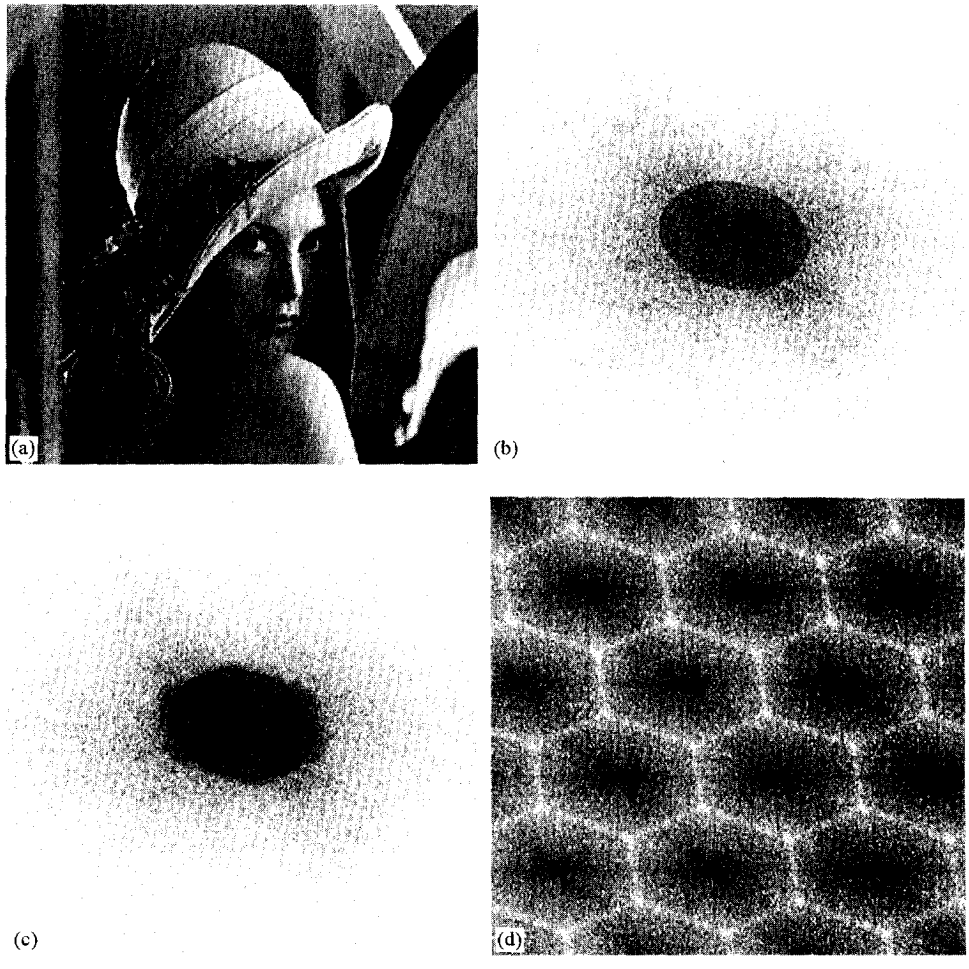

Fig. 10. Geometric characterization of prefilter and decimation grid for the test image *Lenna*. (a) Original image, (b) spectrum and spectral ellipse, (c) hexagonal passband, (d) tiled prefiltered spectrum.

decimation ratios that range from ten to twenty with an acceptable loss of quality after reconstruction. It should be quite clear, however, that the 2D signals that are appropriate the most for the proposed decimation technique are those whose spectral extension is well described by a second-order model. Signals with more complex spectral content might give rise to a more modest results in the reconstruction quality. For example, the method might not perform at its best with signals having a non-compact spectrum. The ideas presented in this article, however, could be adapted to different spectral geometries by modifying the spectral extension model and the geometrical procedure for synthesizing the relative prefilter.

We are currently investigating the possibility of estimating the spectral extension through the analysis of the signal rather than its DFT. In particular, we are investigating the possibility of using the projection-slice theorem for estimating the spectral extension along a limited number of properly chosen 1D FFT's.

A final consideration concerns the fact that only uniform decimation grids are considered. In fact, the proposed method analyzes the signal as if it were stationary, an assumption that is generally not

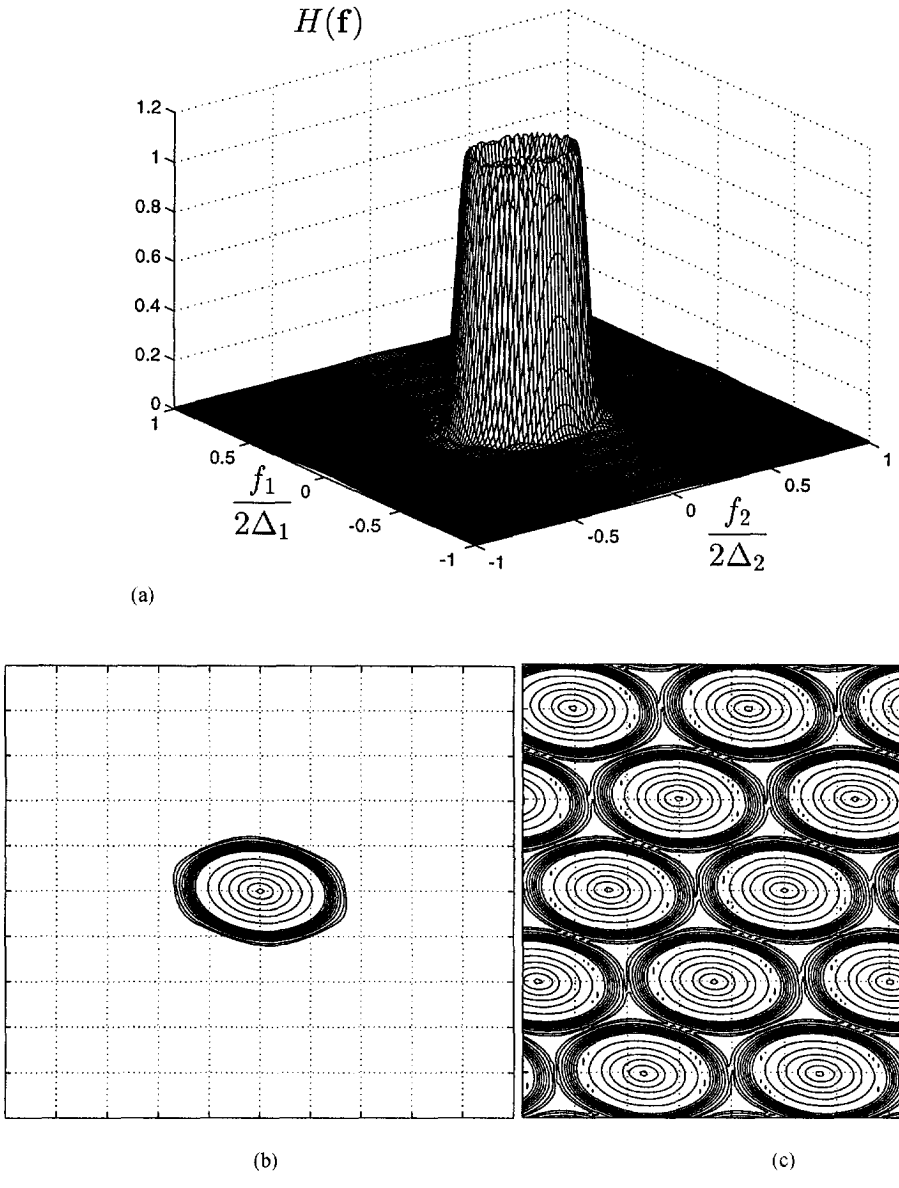

Fig. 11. 2D FIR prefilter for the test image *Lenna*, constructed through frequency transformation. (a) Frequency response, (b) corresponding iso-level curves, (c) Fourier transform of the impulse response, after decimation over Λ.

correct for images, in which case it would be better to assume region-wise stationarity. As a matter of fact, it is reasonable to expect a region-wise adaptive implementation of the algorithm to outperform the proposed one. The derivation of a region-wise implementation, however, would involve a number of problems, such as those that filtering operations would cause in the proximity of region borders, or those connected with the estimation of the spectral extension parameters in irregular regions of small size.

Fig. 12. Prefiltered version of the test image *Lenna* (a); corresponding zoomed-in details in the original (b) and prefiltered (c) images.

Appendix A. Glossary

Z, Q, R	sets of integer, rational and real elements, respectively
Z_M, Q_M, R_M	$M \times M$ matrices with integer, rational and real elements, respectively
$A_{i,j}$	(i,j)th element of the matrix $\boldsymbol{A}$
$A^{-\mathrm{T}} = (A^{-1})^{\mathrm{T}}$	transposed of the inverse of the matrix $\boldsymbol{A}$
$Z_{M,n}$	$\{A \in Z_M, \lvert\det(\boldsymbol{A})\rvert = n\}$
$\mathscr{H}_{M,n}$	set of matrices of $Z_{M,n}$ in Hermite normal form.

Appendix B. Hermite normal form

A non-singular matrix $A \in Z_M$ is said to be *unimodular* if $A^{-1} \in Z_M$. Necessary and sufficient condition for A to be unimodular is $|\det(A)| = 1$. Two matrices $A, B \in Z_M$ are said to be *right-equivalent* if there exists a unimodular matrix $V \in Z_M$ such that $B = AV$. Left-equivalence is similarly defined. Two matrices $A, B \in Z_M$ are said to be *equivalent* if there exist two unimodular matrices $U, V \in Z_M$ such that $B = UAV$.

Two equivalent (or just right/left-equivalent) matrices always have the same *modulus* (absolute value of their determinant). As a consequence, given a positive integer n, it is always possible to uniquely subdivide the set $Z_{M,n} = \{A \in Z_M, |\det(A)| = n\}$ into equivalence classes. The set of equivalence classes represents a partition of $Z_{M,n}$.

Definition B.1. *A matrix $A \in Z_{M,n}$ is said to be in hermite normal form* [27] *if:*
(1) *A is upper triangular;*
(2) $A_{ij} \geqslant 0$;
(3) $A_{ij} < A_{ii}$, $1 \leqslant i < j \leqslant M$, if $a_{ii} \neq 0$;
(4) $A_{ij} = 0$ *if* $A_{ii} = 0$.

The number of distinct matrices of $Z_{M,n}$ in Hermite normal form, is thus given by the sum of all possible distinct integer factorizations of n.

Theorem B.2 (Hermite normal form theorem). *Every non-singular matrix of Z_M is the right-equivalent of one and only one matrix in Hermite normal form.*

Such result [19] is of fundamental importance as it implies that each class of right-equivalents in $Z_{M,n}$ contains a unique matrix in Hermite normal form.

References

[1] R. Bernardini, R. Manduchi, On the reduction of multidimensional DFT to separable DFT by Smith normal form theorem, European Transactions on Telecommunications and Related Technologies 5 (3) (May–June 1994) 377–380.

[2] G. Cariolaro, La Teoria Unificata dei Segnali, UTET, Turin, Italy, 1991.

[3] P. Carrai, G.M. Cortelazzo, G.A. Mian, Characteristics of minimax filters for video interpolation/decimation, IEEE Trans. Circuits Syst. Video Technol. 4 (October 1994) 453–467.

[4] J.W.S. Cassels, An Introduction to the Geometry of Numbers, Springer, Berlin, 1959.

[5] K.F. Cheung, R.J. Marks II, Image sampling below the Nyquist density without aliasing, Journal of the Optical Society of America – A 7 (1) (January 1990) 92–105.

[6] G.M. Cortelazzo, R. Manduchi, On the determination of all the sublattices of preassigned index and its application to multidimensional subsampling, IEEE Transactions on Circuits and Systems for Video Technology, 1993.

[7] P.E. Crochiere, L. Rabiner, Interpolation and decimation of digital signals – a tutorial review, Proceedings of the IEEE 69 (March 1981) 300–331.

[8] E. Dubois, The sampling and reconstruction of time-varying imagery with applications in video systems, Proceedings of the IEEE 73 (April 1985) 502–522.

[9] D.E. Dudgeon, R.M. Mersereau, Multidimensional Digital Signal Processing, Prentice-Hall, Englewood Cliffs, NJ, 1984.

[10] A. Guessam, R.M. Mersereau, Fast algorithms for the multidimensional discrete fourier Transform, IEEE Trans. Acoust. Speech and Signal Process. 34 (4) (August 1986) 937–943.

[11] A. Knoll, Filter design for the interpolation of highly subsampled pictures, Signal Processing: Image Communications 3 (1991) 239–248.

[12] H.E. Knutsson, R. Wilson, G.H. Granlund, Anisotropic nonstationary image estimation and its applications: Part I – Restoration of noisy images, IEEE Transactions on Communications 31 (3) (March 1983) 388–397.

[13] J.S. Lim, Two-Dimensional Signal and Image Processing, Prentice Hall, Englewood Cliffs, NJ, 1990.

[14] P. MacMullen, Convex bodies which tile space by translations, Mathematika 27 (1980) 113–121.

[15] R. Manduchi, Problemi di discretizzazione di immagini in movimento, Ph.D. thesis, University of Padua, Italy, 1992.

[16] J.H. McClellan, The design of two-dimensional digital filters by transformation, Proc. 7th Annu. Princeton Conf. Inform. Sci. Syst., 1973, pp. 247–251.

[17] R.M. Merserau, The design of arbitrary 2D zero-phase FIR filters using transformations, IEEE Trans. Circuits Systems CAS-27 (February 1980) 142–144.

[18] R.M. Mersereau, T.E. Speake, The processing of periodically sampled multidimensional signals, IEEE Trans. Acoust. Speech Signal Process. 31 (February 1983) 188–194.

[19] M. Newman, Integral Matrices, Academic Press, New York, 1972.

[20] V. Ouvrard, P. Siohan, Design of two-dimensional video filters with spatial constraints, in: Proc. of EUSIPCO '92, Brussels, August 1992.

[21] Özdoğan Yilmaz, Seismic data processing, Chapter 4. Series: Investigation in Geophysics, Vol. 2. Society of Exploration Geophysicists, 1987.
[22] F. Pedersini, A. Sarti, S. Tubaro, Joint sublattice selection and prefilter design for the optimal decimation of 2D digital signals, IEE Proceedings – Vision, Image and Signal Processing, 1997.
[23] D.P. Petersen, Discrete fast Fourier transform on N-dimensional lattices, Proceedings of the IEEE 58 (8) (August 1970) 1286–1288.
[24] D.P. Petersen, D. Middleton, Sampling and reconstruction of wave-number-limited functions in N-dimensional Euclidean spaces, Information and Control 5 (1962) 279–323.
[25] F. Scheck, Mechanics, Springer, Berlin, 1990.
[26] D. Taubman, A. Zakhor, Orientation adaptive subband coding of images, IEEE Trans. on Image Processing 3 (4) (July 1994) 421–437.
[27] E. Viscito, J.P. Allebach, The analysis and design of multidimensional FIR perfect reconstruction filter banks for arbitrary sampling lattices, IEEE Transactions on Circuits and Systems 38 (January 1991) 29–41.