

A Wreath Product Group Approach to Signal and Image Processing: Part I — Multiresolution Analysis

R. Foote*
G. Mirchandani†
D. Rockmore‡
D. Healy§
T. Olson¶

August 9, 1999

Abstract

We propose the use of spectral analysis on certain noncommutative finite groups in digital signal processing, and in particular, image processing. We pay significant attention to groups constructed as wreath products of cyclic groups. Within this large class of groups our approach recovers the DFT, Haar wavelet transform, various multichannel pyramid filter banks and other aspects of multiresolution analysis as special cases of a more general phenomenon. In addition, the group structure provides a rich algebraic structure which can be exploited for the analysis and manipulation of signals. Our approach relies on a synthesis of ideas found in the early work of Holmes, Karpovsky, Trachtenberg and others on noncommutative filtering, as well as Diaconis's spectral analysis approach to understanding data.

EDICS numbers: SP 2-HIER and SP 4-MTHE

*Department of Mathematics and Statistics, University of Vermont, Burlington VT 05405. Supported by NASA-Vermont Space Grant

†Department of Electrical and Computer Engineering, University of Vermont, Burlington VT 05405. Supported by NASA\JOVE

‡Department of Mathematics, Dartmouth College, Hanover, NH 03755. Supported by NSF, ONR, ARPA, National Center of Atmospheric Research

§Department of Mathematics, Dartmouth College, Hanover NH 03755. Supported by ARPA

¶Department of Mathematics, Univ. of Florida, Gainesville, FL. Supported by ONR

1 INTRODUCTION

1.1 Motivation

This paper is intended to expand upon the ideas of [32] which introduced the possibility of using finite noncommutative groups obtained by iterated wreath products of cyclic groups for certain signal processing tasks. The hierarchical structure of these groups translates into a natural group invariant multiresolution analysis for the corresponding Fourier analysis, thereby linking the spectral analysis approach for data on finite groups with the world of wavelets. In so doing we realize the DFT, Haar wavelet transform, and certain multichannel pyramid filter banks as special cases of a more general approach. Moreover, by working with an underlying group, we are able to exploit a richer and more general algebraic structure.

The use of finite groups in digital signal processing is not new. The most important instance of this is the formulation of the discrete Fourier transform (DFT) in terms of the representation theory of finite cyclic groups as well as subsequent representation theoretic interpretations of various versions of the Fast Fourier Transform (FFT) which includes the famous Cooley-Tukey FFT [16]. The references [1, 2, 12, 38, 67] give a few examples of this approach.

Indeed, to date, the finite groups that have found significant applications in signal processing are almost exclusively commutative—which is to say groups that are direct products of cyclic groups. This is easily attributed to the broad applicability of Fourier analysis on \mathbb{R}^n , or on tori. Actual computation of a Fourier transform on these continuous groups requires sampling, quantization and possibly truncation of the data. This reduces the original data to a finite set of real or complex machine-representable numbers and allows the required Fourier transform to be rewritten as (or approximated by) a DFT which is then usually computed via some FFT algorithm.

Nevertheless, there have been sporadic appearances of finite noncommutative groups in the signal processing literature. In particular, we appreciatively cite our debt to the fundamental work of R. Holmes whose papers [35, 36] lay the foundation for “group-theoretic signal processing” and continue to be an inspiration. Primarily, noncommutative finite groups have been suggested for applications to filtering and pattern matching. Under the name of “group filters,” Fourier transforms on noncommutative groups have appeared as possible approximations to Wiener filters (cf., [35, 36, 39]). Related to this is the work of Hannan [31] and Willsky [71] which investigates the idea of stationarity for Markov processes on noncommutative state spaces. In fact, Willsky’s paper marks the first appearance of an FFT for a noncommutative finite group. A current survey of such algorithms can be found in [48].

Convolution over dihedral groups [24, 43, 44] and the discrete motion group [41] has been explored in image processing for pattern matching with rotational symmetry; convolution over the affine group, in the context of wideband correlation processing was considered in [70]. Diaconis has suggested a spectral analysis approach for the analysis of data with natural symmetry groups [7, 18, 19]. This generalizes the usual Fourier analysis approach applied to the study of time series. These and other applications are surveyed more fully in [26, 56].

Our point of view is something of a hybrid of the above. Following the basic philosophy of Diaconis’s spectral analysis, we hope to find appropriate symmetry groups for certain types of signals, and then study the signals via their decomposition into orthogonal projections onto group-invariant subspaces. Furthermore, in the

spirit of the work of Holmes, Karpovsky and Trachtenberg, we use the underlying group structure to make use of the associated convolution and correlation operations for signal manipulation.

Of particular interest is the setting in which wreath product groups act on the space of discrete signals of a fixed finite length. This action is realized by indexing the samples as leaves in a (graph theoretic) tree. Wreath products are the symmetry groups of regularly branching rooted trees, and consist in those relabelings that only permute nodes at a fixed distance from the root. We shall see that these wreath product groups arise in a natural way in pyramid structured multichannel DFT filter banks (see [38] for a group-theoretic approach to multidimensional filter banks). The analysis stage of these filter banks in polyphase form turns out to be the same as the Fourier transform of a tree-encoded discrete signal with respect to the action of a wreath product group. The presence of an underlying group, however, now provides additional structure that has heretofore been unexplored. In particular, there is a group-based convolution which gives rise to a wealth of new group-invariant filters for investigating classical problems such as filtering, pattern recognition, data compression, and noise reduction.

In this first part of a two-part sequence of papers we introduce a general theory of finite group-based signal processing: multiresolution analysis and group-based Fourier decompositions of discrete signals. We also describe the basic structure of (iterated) wreath product groups, focusing specifically on wreath products of cyclic groups with the application to two-dimensional images in mind. We show how the irreducible decomposition of any signal provides a natural multiresolution analysis that, in the case of wreath product groups, generalizes the Haar wavelet decomposition. We describe a *quadtree* scanning scheme for converting a discrete $2^N \times 2^N$ signal into a row vector in such a way that these wreath product groups act on nested grids of subimages, and we describe this action in detail in the tractable yet realistic case of 512×512 images. We also give a quadtree scheme for displaying the multiresolution spectrum of such images, and we examine the effect of group operations on these quadtree spectra for 512×512 images.

In Part II of this work, [50], we introduce the general theory of finite group-based convolution and correlation. For wreath products of cyclic groups we show how these may be effectively computed by passing to the frequency (spectral) domain using the quadtree schemes. Building on the applications described in Part I, we give a “geometric” description of convolution. Finally, we apply convolution and correlation to the problems of recognizing perceptually similar or transformed images, and we see that, empirically, the wreath product group-based results appear to have certain advantages over other image recognition methods.

1.2 Other related work

The particular focus of our development which uses mainly the representation theory of finite groups over the complex numbers for filtering and spectral analysis is but one way in which group theory appears in signal processing. Here are a few other significant applications.

1.2.1 Coding theory

The DFT (over finite fields) is an important tool in the subject of error correcting codes (cf., [10]). Many well-known codes have natural descriptions in terms of the group algebra of a cyclic group which translate into

easily stated properties of their Fourier transforms over finite fields. MacWilliams [45] appears to have been the first to investigate the extension of this to a noncommutative group (the dihedral group). More recently, Lafferty and Rockmore [42] have been able to formulate some of the expander codes of Sipser and Spielman [62] in terms of the representation theory of the special linear groups of degree two over a finite field.

1.2.2 Continuous groups

Although to this point, the use of finite noncommutative groups and their representation theory has been rare, continuous noncommutative groups play a significant role in signal processing and data analysis.

Many classes of orthogonal functions can be realized as the special functions of certain non-abelian Lie groups (see e.g. [68]). The papers [21, 23] provide some new general approaches for fast transforms using these functions. Of particular importance are the use of spherical harmonics for the analysis of functions defined on the 2-sphere. Computational aspects of these functions have long been of interest for their use in the numerical solution of PDEs in spherical geometry, applicable to problems in the atmospheric sciences (cf., [20, 33, 34]). Extensions of these ideas to the general framework for sampling and computing on compact groups can be found in [47].

In the noncompact setting, Fourier analysis on motion groups is now finding applications in robotics [15, 25] as well as automatic target recognition [63, 64].

1.2.3 Wavelets and representation theory

One origin of wavelet theory can be traced back to the mathematical physics community and the study of cyclic representations of the affine group [30]. The work of Flornes, *et al.*, [29] develops a finite analog of this by working with the affine group mod p . The papers [8, 9, 51] pursue a different aspect of the construction, focusing instead on the condition of generating a basis for a representation space by the set of (group) translates of a fixed vector.

1.3 Organization

In order to put things into a familiar framework, we begin in Section 2 by first recalling the representation-theoretic formulation of the DFT, paying specific attention to those aspects which will motivate our general approach. Sections 3 and 4 of this paper lay the foundations of some theoretical aspects of group-based signal processing; in each of these sections the theory is followed by applications to the specific case of wreath product groups. In this way the paper serves both as an in-depth illustration of the general methodology of group-based signal processing for future research, and an investigation of the intrinsically interesting application of wreath products to signal processing.

In Section 3 we give a general group-theoretic formulation of signal processing. We pay special attention to the example of wreath products of cyclic groups. Section 4 describes the general theory of finite group based Fourier analysis, and relates it to multiresolution analysis and wavelet decompositions. The specific decompositions obtained from wreath products of cyclic groups are described. Section 5 applies the theory to

discrete $2^N \times 2^N$ images, describing specific scanning methods and how wreath product groups act on images and their spectra. Section 6 places this theory in some larger contexts. It describes multiresolution schemes that extend this work and which may lead to areas of future research.

2 The DFT, FFT, AND GROUP THEORY

The connection between the DFT and group theory is immediate from its definition. Given an input $f = (f(0), \dots, f(N-1)) \in \mathbb{C}^N$, its DFT is defined as the collection of sums (Fourier coefficients)

$$\hat{f}(k) = \sum_{n=0}^{N-1} f(n)W^{nk} \quad (2.1)$$

where $W = e^{2\pi j/N}$ for $j = \sqrt{-1}$. The N^{th} roots of unity W^{jk} are the values of the irreducible characters (or in this case, equivalently, the irreducible matrix elements) of the cyclic group $\mathbb{Z}/N\mathbb{Z}$. If computed for all $k \in \{0, \dots, N-1\}$, the expression (2.1) computes the change of basis for the function $f : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ from the basis of delta functions

$$\delta_n(k) = \begin{cases} 1 & \text{if } n = k \\ 0 & \text{otherwise} \end{cases}$$

to the basis of irreducible matrix elements

$$\chi_n(k) = W^{nk}.$$

The new expression for f is then given by the Fourier inversion formula, or inverse discrete Fourier transform

$$f(n) = \frac{1}{N} \sum_{k=0}^{N-1} \hat{f}(k)W^{-nk}. \quad (2.2)$$

The choice of the functions χ_n is motivated by their invariance under the group of cyclic shift. Given any $m \in \mathbb{Z}/N\mathbb{Z}$ we can consider the translation of a function $f : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ by m , denoted (mf) and defined as

$$(mf)(n) = f(n - m).$$

Thus, the characters χ_n are simultaneous eigenfunctions for all translation operators,

$$(m\chi_n) = W^{-nm}\chi_n.$$

Consequently, we derive the invariance of the Fourier coefficients (up to a phase factor) under the entire group of translations, or (cyclic) shifts of the origin. This motivates the interpretation of these roots of unity as character values for the cyclic group. The approach of understanding the input according to these Fourier coefficients is usually known as the *spectral analysis* of the input. The square of the magnitudes of the Fourier coefficients give the power spectrum.

Representation of a function in the delta function basis is often referred to as a *spatial-domain* expression, while the basis of sampled exponentials give the *frequency-domain* expression (see (2.2)).

The N -dimensional vector space of functions $L^2(\mathbb{Z}/N\mathbb{Z}) = \{f : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}\}$ is not only a linear space, but also carries a natural \mathbb{C} -linear multiplication, better known as *convolution*, which extends the group multiplication. For $f, g \in L^2(\mathbb{Z}/N\mathbb{Z})$, this is defined as

$$f \star g(n) = \sum_{m \in \mathbb{Z}/N\mathbb{Z}} f(m)g(n-m). \quad (2.3)$$

Considered with this additional structure, the vector space $L^2(\mathbb{Z}/N\mathbb{Z})$ is called the (complex) group algebra of $\mathbb{Z}/N\mathbb{Z}$. Any fixed function h gives rise to an associated *convolution operator* T_h or *filter* defined by $T_h(f) = f \star h$. Application of convolution often carries the interpretation of smoothing (in the spatial domain). Notice that in this instance convolution is commutative.

A well-known (and often-used) property of the basis of characters, is that the Fourier coefficients of the convolution of two functions multiply:

$$(\widehat{f \star g})(k) = \widehat{f}(k)\widehat{g}(k) \quad (2.4)$$

making possible the direct specification of filters in terms of their effect in the frequency domain. *Correlation* is a relative of convolution and is computed as $f \star \tilde{h}$ where $\tilde{h}(n) = h(-n)$.

The diagonalization of convolution operators by the Fourier basis (2.4) is equivalent to the fact that any circulant matrix is diagonalized by the Fourier basis. This is useful in statistical signal processing since it is further equivalent to the fact that the Fourier transform decorrelates noise from a Markov-1 process. This also underlies the use of the DFT as an approximate Wiener filter for stationary signals as well as its use for compression in this setting. At the heart of this is the asymptotic equivalence of the DFT and the Karhunen-Loeve transform in this setting [66].

Group theory also contributes to an understanding of the efficient algorithms for computing the DFT (cf., [48]). In particular the Cooley-Tukey FFT has a beautiful interpretation in terms of subgroups and cosets.

For this, let $N = N_1 N_2$, $N_i > 1$. Then

$$\begin{aligned} \widehat{f}(k) &= \sum_{n=0}^{N-1} f(n)W^{nk} \\ &= \sum_{n_2=0}^{N_2-1} \sum_{n_1=0}^{N_1-1} f(n_1 N_2 + n_2)W^{(n_1 N_2 + n_2)k} \\ &= \sum_{n_2=0}^{N_2-1} W^{n_2 k} \sum_{n_1=0}^{N_1-1} f(n_1 N_2 + n_2)W^{(n_1 N_2)k} \\ &= \sum_{n_2=0}^{N_2-1} W^{n_2 k} \sum_{n_1=0}^{N_1-1} f^{n_2}(n_1) e^{2\pi j n_1 k / N_1} \end{aligned} \quad (2.5)$$

where $f^{n_2}(n_1) = f(n_1 N_2 + n_2)$. Notice that the inner summation of the last equation has the form of a DFT of length N_1 . In group-theoretic terms, we have used the decomposition of the group $\mathbb{Z}/N\mathbb{Z}$ into cosets by the subgroup $\mathbb{Z}/N_1\mathbb{Z}$ to rewrite the original DFT in terms of a DFT on the subgroup $\mathbb{Z}/N_1\mathbb{Z} < \mathbb{Z}/N\mathbb{Z}$. An efficient computation of the DFT of f comes about by first computing the (smaller) DFTs of each of the

N_2 functions $\widehat{f^{n_2}}$ and then combining them as per (2.5). As a result of this reorganization, instead of requiring $N^2 = (N_1 N_2)^2$ operations to compute the DFT, we now have an algorithm requiring $N_1 N_2 (N_1 + N_2)$ operations.

The goal of the remainder of the paper is to explain how all of the above makes sense for any finite group G and to give some evidence for the applicability of some particular noncommutative groups. That is, given a function $f : G \rightarrow \mathbb{C}$, (or equivalently, an element of the complex group algebra of G defined in terms of the basis of delta functions) we can rewrite it in terms of a collection of irreducible matrix elements of G . This is the *spectral analysis* of the input f . The result is a collection of Fourier coefficients which exhibit some sort of G -invariance, and which transform nicely with respect to convolution over G , or multiplication in the group algebra. This makes possible the description of particular noncommutative filters. Fourier analysis can be accomplished efficiently using algorithms which are completely analogous to the Cooley-Tukey FFT [48, 49].

3 SIGNALS ON FINITE GROUPS —A GROUP ACTION APPROACH

At the heart of this theory is the observation (cf., [35, 36]) that a discrete signal of fixed finite length N , $f = (f(x_0), \dots, f(x_{N-1}))$ can be viewed as a function on any set $X = \{x_0, \dots, x_{N-1}\}$ (with a fixed ordering) on which a group G acts. The group action determines a decomposition of the *signal space* (the complex vector space \mathbb{C}^N , now identified with the space of possible discrete signals) into group-invariant pieces. Most of classical signal processing assumes that both G and X are the sets $\mathbb{Z}/N\mathbb{Z}$ with the association $x_n \leftrightarrow n \in \mathbb{Z}/N\mathbb{Z}$, implicitly ordered by $n < n + 1$. In higher dimensions we would have $N = m_1 m_2 \cdots m_r$ and $G = X = \mathbb{Z}/m_1\mathbb{Z} \times \cdots \times \mathbb{Z}/m_r\mathbb{Z}$. Nevertheless, other choices are possible and may be preferable. In this section we lay down the basic theory of this general approach which relies on the notion of a permutation representation corresponding to a group action.

3.1 Basic theory

Let X be a finite set. We think of X as an index set for a signal. In the most familiar examples, X will be the set of points on a sampling grid. A *signal* is just a complex-valued function on X , $f : X \rightarrow \mathbb{C}$, although in some instances it may be useful to allow the range of f to be a finite field (cf., [29, 37, 52, 65]). By concentrating on the complex or characteristic zero case, we make use of the “ordinary representation theory” of the designated symmetry group [59]. Much of the following goes through in the finite field case as well.¹

Let $L(X)$ denote the complex vector space of signals on X . When $|X| = N$, $L(X)$ is N -dimensional with *spatial-domain basis* given by the impulse functions $\{\delta_x\}_{x \in X}$ defined by

$$\delta_x(y) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{otherwise.} \end{cases} \quad (3.1)$$

¹As long as the characteristic of the field is prime to the order of the group, and we work in a sufficiently large extension field, then the representation theory is the same.

In this basis a signal f has the expansion

$$f = \sum_{x \in X} f(x) \delta_x.$$

Fixing an order for the elements of X is equivalent to identifying $L(X)$ with the complex vectors of length N , in which case the basis $\{\delta_x\}_{x \in X}$ is identified with the standard basis. Also, since X is a finite set, all functions in $L(X)$ are continuous in the discrete topology on X and so as a vector space $L(X)$ is the discrete space version of the L^p -space $L^p(X)$ for every p , although they are of course distinct as metric spaces. Our applications assume the use of the L^2 -norm, although recent work [69] indicates that other values of p may be better-suited for some applications.

Let G be any finite group acting on X as permutations. This simply means that G can be realized as a set of bijections of X with itself, closed under composition and inverse. For any $\alpha \in G$ and $x \in X$ we write αx for the image of x under the bijection α and say that αx is the *translate* of x by α .

For each $\alpha \in G$ and each $f \in L(X)$ we define a new function $(\alpha f) : X \rightarrow \mathbb{C}$ by

$$(\alpha f)(x) = f(\alpha^{-1}x), \quad \text{for all } x \in X. \quad (3.2)$$

The function αf is the *translate* of f by α . Notice that this makes $\alpha \in G$ a linear transformation of $L(X)$, i.e., for $f_1, f_2 \in L(X)$, $\alpha \in G$ and $c \in \mathbb{C}$,

$$\begin{aligned} \alpha(f_1 + f_2) &= (\alpha f_1) + (\alpha f_2) \\ \alpha(cf) &= c(\alpha f). \end{aligned} \quad (3.3)$$

Furthermore, the composition of these linear transformations is compatible with the composition of elements in the group G :

$$(\alpha\beta)f = \alpha(\beta f), \quad \text{for all } \alpha, \beta \in G. \quad (3.4)$$

The inverse in (3.2) is necessary so that (3.4) is satisfied for non-commutative groups.

Thus G acts as a group of linear transformations of $L(X)$, and (3.3) and (3.4) are the conditions which define the action as a (linear) *representation* of G on $L(X)$. The particular representation defined by (3.2) is called the *permutation representation associated to the action of G on X* . We now give some examples.

Example 0. Let $X = G$. Then G acts on itself by left (right) translation. The associated action of G on $L(G)$ is called the *left (right) regular representation* of G .²

Example 1. Let $G = Z_N = \langle \sigma \rangle$ be the cyclic group of order N with generator σ , and let $X = \{0, \dots, N-1\} = \mathbb{Z}/N\mathbb{Z}$. Then G acts on X by translation,

$$\sigma^m n = \overline{n + m}$$

where the bar indicates that the addition is performed modulo N . With the association $f \in L(\mathbb{Z}/N\mathbb{Z}) \leftrightarrow (f(0), \dots, f(N-1))$ the action of σ^m on $L(X)$ effects a cyclic shift (to the right) of f by m positions.

$$\sigma f = ((\sigma f)(0), \dots, (\sigma f)(N-1)) = (f(N-1), f(0), \dots, f(N-2)). \quad (3.5)$$

²In the case of right translation the action is defined by $(\alpha f)(\beta) = f(\beta\alpha)$.

Example 2. Let $G = Z_n \times Z_m = \langle \sigma \rangle \times \langle \tau \rangle$, the direct product of cyclic groups, and $X = \{0, \dots, n-1\} \times \{0, \dots, m-1\} = \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$. We think of X as corresponding to the usual grid points on the $n \times m$ lattice. In this way, to any function $f \in L(X)$ we associate a matrix, given by the corresponding function values. Then $(\sigma^i, \tau^j)(r, s) = (\overline{i+r}, \overline{j+s})$ with the bar indicating the appropriate modular arithmetic. Thus the group action performs independent cyclic shifts of left and right indices.

$$\begin{aligned}
f &= \begin{pmatrix} f(0,0) & \dots & f(0,m-1) \\ \vdots & \vdots & \vdots \\ f(n-1,0) & \dots & f(n-1,m-1) \end{pmatrix} \\
(\sigma, 1)f &= \begin{pmatrix} f(n-1,0) & \dots & f(n-1,m-1) \\ f(0,0) & \dots & f(0,m-1) \\ \vdots & \vdots & \vdots \\ f(n-2,0) & \dots & f(n-2,m-1) \end{pmatrix} \\
(1, \tau)f &= \begin{pmatrix} f(0,m-1) & f(0,0) & \dots & f(0,m-2) \\ \vdots & \vdots & \vdots & \vdots \\ f(n-1,m-1) & f(n-1,0) & \dots & f(n-1,m-2) \end{pmatrix} \\
(\sigma, \tau)f &= \begin{pmatrix} f(n-1,m-1) & f(n-1,0) & \dots & f(n-1,m-2) \\ f(0,m-1) & f(0,0) & \dots & f(0,m-2) \\ \vdots & \vdots & \vdots & \vdots \\ f(n-2,m-1) & f(n-2,0) & \dots & f(n-2,m-2) \end{pmatrix}
\end{aligned}$$

Examples 1 and 2 are the canonical examples of digital signal processing, giving rise to the application of the 1-D and 2-D DFT respectively.

Example 3. Let $G = S_N$, the symmetric group on N symbols and $X = \{0, \dots, N-1\}$ (i.e., G is the group of all permutations of the set X). Then any permutation $\pi \in S_N$ acts on X by sending i to $\pi(i)$. All of the preceding examples are really subexamples of this, in the sense that G is a subgroup of an appropriate S_N .

Given a group action G on X and corresponding representation of G on the signal space $L(X)$, we are interested in group-invariant decompositions of the signal space under the group action. A subspace W is G -invariant if given any $f \in W$, then $\alpha f \in W$ for all $\alpha \in G$. Thus, a one-dimensional G -invariant subspace is simply a common eigenspace for the group of linear operators defined by the action of G on $L(X)$. In this sense, invariant subspaces are generalized eigenspaces, and a decomposition of a signal (function $f \in L(X)$) according to projections onto these spaces is a generalized Fourier expansion of the signal.

As a first step in this direction, notice that if the set of nodes X decomposes as the disjoint union $X_1 \cup X_2$, with each X_i being G -invariant (so that G only maps elements of X_1 among themselves, and similarly for X_2) then $L(X)$ decomposes as a sum of G -invariant subspaces, $L(X) = L(X_1) \oplus L(X_2)$. The theory (developed in Section 4) may then be applied individually to each $L(X_i)$. Thus, we may reduce to the case in which X has no nonempty, proper G -invariant subsets, i.e., *we may assume G acts transitively on X* . Thus, for each $x, y \in X$ there is some $\alpha \in G$ such that $\alpha x = y$. In this case we also say that X is a *homogeneous space* for G .

Each of the following Examples 0–3 are transitive. In Example 1 the cyclic group of order N acts transitively on $X = \mathbb{Z}/N\mathbb{Z}$. Moreover, the action gives us a way to identify the sets G and X by associating to each $i \in X$ the unique group element $\sigma_i \in G$ that maps 0 to i . In other words, if G is generated by the N -cycle σ that cyclically permutes $0, 1, \dots, N-1$, then i may be identified with σ^i .

The above describes a *regular* group action, one in which each element is left fixed only by the identity. More generally, $X = \{x_0, \dots, x_{N-1}\}$ may be identified with a set of left cosets in G as follows. For a distinguished “basepoint” $x_0 \in X$, let

$$G_0 = \{\alpha \in G \mid \alpha(x_0) = x_0\}. \quad (3.6)$$

This defines G_0 as the *stabilizer subgroup of x_0 in G* . Under the assumption of transitivity, for each i we may choose an $\alpha_i \in G$ with $\alpha_i(x_0) = x_i$. We choose α_0 to be the identity element of G . It follows easily that the left coset $\alpha_i G_0$ is the set of all elements of G that send x_0 to x_i . Furthermore, the association

$$x_i \longleftrightarrow \alpha_i G_0 \quad (3.7)$$

is a one-to-one correspondence that identifies X with the set G/G_0 of left cosets of G_0 in G .

Examples 0, 1 and 2. These are regular actions, where in each case G_0 is the identity subgroup, so each left coset contains a single element of G , i.e., the left cosets of G_0 are the elements of G .

Example 3. We fix a basepoint $N-1$. Then the stabilizer is the subgroup of permutations that permute only $0, \dots, N-2$ among themselves, so is isomorphic to S_{N-1} .

This identification of X with left cosets is compatible with the group action on each of the two sets X and G/G_0 as follows. The group G acts by left multiplication on the set of left cosets of G_0 ; namely each $\alpha \in G$ maps $\alpha_i G_0$ to the left coset $(\alpha\alpha_i)G_0$, and this left multiplication by α is a permutation on the set G/G_0 . Under the identification (3.7) each α in G induces the same permutation on X and G/G_0 :

$$\alpha(x_i) = x_k \quad \text{if and only if} \quad (\alpha\alpha_i)G_0 = \alpha_k G_0.$$

Thus we may identify X and G/G_0 as sets admitting an action by G . We use this to embed $L(X)$ into the (possibly larger) space $L(G)$ as follows.

Observe that the set of functions that are constant on the left cosets of any subgroup forms a subspace of $L(G)$. In particular, for G_0 , the stabilizer of x_0 , let $L(G/G_0)$ denote this subspace:

$$L(G/G_0) = \{f \in L(G) \mid f(\alpha\beta) = f(\alpha), \text{ for all } \alpha \in G, \beta \in G_0\}. \quad (3.8)$$

It is immediate from (3.2) that $L(G/G_0)$ is also a G -invariant subspace. Namely, if f is constant on all left cosets of G_0 , then so is αf for any $\alpha \in G$.

Under the identification in (3.7) of X with G/G_0 , the complex-valued functions on G/G_0 may be identified with those on X . As we shall see in [50], in order for convolution multiplication to be compatible between $L(G)$ and $L(X)$ we choose a specific way of making this identification (which amounts to a specific choice of

normalizing factor): If $f \in L(G/G_0)$, identify f with the function in $L(X)$ whose value on the point x_n is $|G_0|f(\alpha_n)$. This identification is the restriction to the subspace $L(G/G_0)$ of the map

$$\pi_X : L(G) \longrightarrow L(X)$$

defined by

$$\pi_X(f)(x_n) = \sum_{\beta \in \alpha_n G_0} f(\beta), \tag{3.9}$$

for all $f \in L(G)$. This maps any function in $L(G)$ to $|G_0|$ times its average value on each coset $\alpha_n G_0$; the factor of $|G_0|$ is retained so that f and $\pi_X(f)$ have the same total sum, i.e., $\sum_{\alpha \in G} f(\alpha) = \sum_{x_n \in X} \pi_X(f)(x_n)$. Thus, if δ_{α_n} is the basis function in $L(G)$ defined as in (3.1), then $\pi_X(\delta_{\alpha_n})$ is the unit impulse delta function in $L(X)$ supported at x_n .

Finally, note that the actions of G on $L(X)$ and $L(G)$ are compatible via π_X :

$$\pi_X(\alpha f) = \alpha(\pi_X(f)),$$

for all $\alpha \in G, f \in L(G)$ i.e., π_X is a G -equivariant linear transformation or G -homomorphism.

Remark. The permutation representation of G on $L(G/G_0)$ is equivalent to the *induction of the trivial representation from G_0 to G* . We shall elaborate a bit more on this point of view in Sections 4 and 6.

3.2 Trees and wreath product group actions

Our primary application of the general theory of group actions on signal spaces will be to wreath product groups acting on trees. We describe this situation and indicate its relevance to signal and image processing. We consider the general case of a spherically homogeneous rooted tree of many levels, with symmetries given by independent cyclic shifts at each level, ultimately realizing a group of symmetries on the set of leaves of the tree. Our applications arise by identifying the leaves of the tree with sample points, with particular attention paid to the case of two-dimensional sampling grids and the associated permutation representation on the vector space of images.

3.2.1 Generalities

The applications discussed here depend on understanding the action of certain wreath product groups on a particular class of directed graphs, spherically homogeneous rooted trees (SHRTs) (see [6]). Recall that a *tree* is a graph with no cycles. A *rooted tree* has a distinguished vertex, the *root*, and all other vertices are judged in relation to their distance from the root, measured in terms of the length of the shortest path to the root. A vertex of distance k from the root is said to be at *level k* . A *child* of a vertex v is a vertex at level $k+1$ connected to v (of level k). A SHRT is a rooted tree such that all vertices at a fixed distance from the root have the same number of children. If m_k is the number of children of any vertex at distance k , then the SHRT is said to be of *type* (m_0, m_1, \dots, m_l) where $l+1$ is the *height* (or *depth*) of the tree. The sequence $\mathbf{m} = (m_0, \dots, m_l)$

determines the SHRT which is denoted as $\mathcal{T}_{\mathbf{m}}$.³ Notice that a SHRT with all $m_k = 2$ is a complete binary rooted tree. Figure 1 gives a few examples of SHRTs.

The vertices of maximal distance from the root are called the *leaves* of the tree (i.e., those at level $l + 1$). We denote the leaves of the tree $\mathcal{T}_{m_0 \dots m_l}$ by $X_{(m_0, \dots, m_l)}$. Notice that $|X_{\mathbf{m}}| = m_0 \cdots m_l$, and that the leaves are naturally indexed by sequences (i_0, \dots, i_l) with $0 \leq i_k < m_k$. Pictorially, with the “usual” picture of an SHRT (see Figure 1) this would indicate the leaf obtained by choosing the i_k^{th} branch (going from left to right) at step k as we make our way down from the root to the leaf. Figure 2 illustrates this.

SHRTs have recursive structure in a variety of ways. Each vertex of level one of a SHRT of type m_0, \dots, m_l may be viewed as the root of a SHRT of type m_1, \dots, m_l . These are the *subtrees* of level one. More generally, any vertex at level k can be viewed as the root of a SHRT of type m_k, \dots, m_l , which make up the *subtrees of level k* . On the other hand, by collapsing all subtrees of level k to a point, the SHRT of type m_0, \dots, m_l is converted to a SHRT of type m_0, \dots, m_{k-1} . We call this the k^{th} *truncation* of $\mathcal{T}_{(m_0, \dots, m_l)}$, as illustrated in Figure 3.

Wreath product groups may be realized geometrically as adjacency preserving permutations of SHRTs that fix the root. This implies that only vertices within a given level are permuted. In particular, symmetries of a SHRT give rise to structured permutations of the leaves. The full group of symmetries of $T_{\mathbf{m}}$ is described in a hierarchical manner: For each i between 0 and k , and each vertex at level i , choose a permutation in S_{m_i} . The permutations at level i define a reordering of the children of the vertex at level i . Apply this reordering while leaving fixed the relative ordering within the subtrees determined by each of the children. Thus, any symmetry of $T_{\mathbf{m}}$ can be written as a sequence $(\sigma_l, \dots, \sigma_0)$ where σ_k is a function on the leaves of the k^{th} truncation of $T_{\mathbf{m}}$, with values in the group S_{m_k} . This may be pictured as the tree $\mathcal{T}_{(m_0, \dots, m_l)}$ with all nodes at level k labeled by permutations in S_{m_k} .

The above describes the full automorphism group which is a wreath product of symmetric groups. More generally, certain subgroups, also wreath products, can be obtained by choosing for each level i , a specific subgroup H_i of S_{m_i} then choosing the permutations at level i from H_i . This defines the (iterated) wreath product of the groups H_i , denoted by

$$G = H_l \wr H_{l-1} \wr \cdots \wr H_0. \quad (3.10)$$

We note that wreath product groups are usually defined algebraically, without reference to trees (see, for example, [22] or [58]), however we adopt this “geometric” approach to defining them with a view to applications in signal and image processing. In this context there are natural tree structures on our spaces of discrete signals which help to illuminate the abstract theory. Also, (iterated) wreath products arise naturally as the automorphism groups of nested designs [4]. They also occur in chemistry as the symmetry groups of certain regularly branching non-rigid molecules [5, 72].

Given a choice of G as in (3.10) we can describe the symmetries of $T_{\mathbf{m}}$ in the following recursive manner: (1) Choose an element of H_0 , and permute the m_0 subtrees of the root among themselves, leaving fixed the relative

³We only concern ourselves with finite SHRTS. The infinite theory is discussed in some detail in [6].

order within each subtree. (2) Independently choose m_0 elements of the iterated wreath product $H_l \wr H_{l-1} \wr \cdots \wr H_1$ to act on each of the subtrees of the root, considered as separate SHRTs of the type (m_1, \dots, m_l) .

Thus, if we let $G_{(r)} = H_l \wr H_{l-1} \wr \cdots \wr H_r$, then

$$G = G_{(0)} \cong G_{(1)} \wr H_0.$$

We can write an element of G as (v, σ) with $v \in G_{(1)}^{m_0}$ and $\sigma \in H_0$, where for any group A the direct product of A with itself m times is denoted by A^m . Multiplication of group elements (v, σ) and (w, τ) is then performed by

$$(v, \sigma)(w, \tau) = (v \cdot (\sigma w), \sigma \tau)$$

where

$$(\sigma w)(k) = w(\sigma^{-1}(k))$$

and

$$v \cdot (\sigma w)(k) = v(k)w(\sigma^{-1}(k)).$$

Implicitly, the iterated wreath product $G = H_l \wr H_{l-1} \wr \cdots \wr H_0$ has a recursive structure. If, as above, we define $G_{(r)} = H_l \wr H_{l-1} \wr \cdots \wr H_r$ and $G^{(s)} = H_s \wr H_{s-1} \wr \cdots \wr H_0$ then note that $G_{(r)}$ is an automorphism group of $\mathcal{T}_{(m_r, \dots, m_l)}$ and $G^{(s)}$ is an automorphism group of $\mathcal{T}_{(m_0, \dots, m_s)}$ and $G = G_{(r)} \wr G^{(r-1)}$. This is equivalent to viewing any symmetry as being given by a symmetry of the truncated tree of depth r followed by independent symmetries of the $m_0 \cdots m_{r-1}$ subtrees at level r .

This description reveals the semidirect product structure of the wreath product groups. In particular we see that any symmetry can be written uniquely as a product of an element of $G^{(r-1)}$ and $m_0 \cdots m_{r-1}$ independent elements of $G_{(r)}$, or more succinctly, a single element of $(H_l \wr H_{l-1} \wr \cdots \wr H_r)^{m_0 \cdots m_{r-1}}$. The latter makes up the subgroup of G which leaves invariant all nodes of $\mathcal{T}_{(m_0, \dots, m_l)}$ in levels $0, 1, \dots, r$. As such, this is a normal subgroup of G (it is the kernel of the action on the nodes at level r). Thus, if we let B_r denote this subgroup, then B_r is a normal subgroup of G (denoted $B_r \trianglelefteq G$). Moreover, a particular subgroup complement to B_r in G consists of the set of all elements of G that send each leaf $(i_0, \dots, i_{r-1}, i_r, \dots, i_l)$ in the indexing scheme described earlier to another leaf $(i'_0, \dots, i'_{r-1}, i_r, \dots, i_l)$ having the same $l - r + 1$ last indices. Denoting this subgroup by $G^{(r-1)}$ we have $B_r \cap G^{(r-1)} = \{1\}$ and $B_r G^{(r-1)} = G$; or equivalently, G is a semidirect product of B_r and $G^{(r-1)}$. We collect this in the following theorem. (The verification of all the details of the proof is a straightforward exercise; alternatively, these details may easily be checked from the algebraic, semidirect product definition of wreath product groups.)

Theorem 3.1 *Let $G = H_l \wr H_{l-1} \wr \cdots \wr H_0$. With all notation as above, the subgroups B_r form a normal series*

$$\{1\} = B_{l+1} \trianglelefteq B_l \trianglelefteq B_{l-1} \trianglelefteq \cdots \trianglelefteq B_0 = G \tag{3.11}$$

where B_0/B_1 is a copy of H_0 , and for $1 \leq r \leq l$

$$B_r/B_{r+1} \text{ is the direct product of } m_0 m_1 \cdots m_{r-1} \text{ copies of } H_r. \tag{3.12}$$

Moreover, G/B_r acts as permutations of the truncated tree $\mathcal{T}_{(m_0, \dots, m_{r-1})}$ obtained from $\mathcal{T}_{\mathbf{m}}$ by deleting all nodes at levels greater than r and

$$G^{(r-1)} = G/B_r \text{ is isomorphic to } H_{r-1} \wr H_{r-2} \wr \dots \wr H_0. \quad (3.13)$$

Furthermore, B_r has a subgroup complement in G isomorphic to $G^{(r-1)}$ and G is a semidirect product of B_r with $G^{(r-1)}$ for each r .

The order of any specific wreath product group may be determined from (3.12) and (3.13) recursively.

In Sections 4 and 6 we shall see that the subgroup structure of wreath products given in Theorem 3.1 gives the “multiresolution” filtration of the spatial domain $L(X)$ into a chain of group invariant subspaces.

In this paper we shall focus on the smallest wreath product groups that are transitive on the nodes at each level i . For each $i \in \{0, 1, \dots, l\}$ let $H_i = Z_{m_i}$ be the cyclic subgroup of S_{m_i} generated by the m_i -cycle $(1\ 2 \dots m_i)$. Define

$$Z_{\mathbf{m}} = Z_{m_l} \wr Z_{m_{l-1}} \wr \dots \wr Z_{m_0}. \quad (3.14)$$

We shall call the groups $Z_{\mathbf{m}}$ *wreath product cyclic groups*, abbreviated *WPC groups*. WPC groups are non-abelian whenever $l \geq 1$. The group $Z_{\mathbf{m}}$ acts on $\mathcal{T}_{\mathbf{m}}$ by independent cyclic shifts at each level.

For example, consider $Z_{(3,4)}$ which acts on the tree $\mathcal{T}_{(3,4)}$ depicted in Figure 4. We specify any element as $(i_1, i_2, i_3; k)$ where $0 \leq k \leq 2$, and $0 \leq i_m \leq 3$ indicating the various powers of cyclic shift of three and four elements respectively. In Figure 4 we show the effect of two different group elements.

3.2.2 Wreath products and nested grid decompositions

We now describe how the particular case of (3.14) when each $m_i = 4^k$ arises naturally in the study of nested grid decompositions of $2^N \times 2^N$ discrete images. In this special case, for any n -tuple $(4^k, 4^k, \dots, 4^k)$, we streamline notation by denoting the tree $\mathcal{T}_{(4^k, 4^k, \dots, 4^k)}$ by $Q(k, n)$, and referring to it as a *quadtrees*. In the notation of Section 3.2.1 we have $n = l + 1$, so $Q(k, n)$ has n nonzero levels (depth n) and 4^{kn} leaves.

A $2^{kn} \times 2^{kn}$ array of pixel positions may be given a quadtree structure (see e.g. [28, 57]) in the following nested grid fashion: The entire array is the root node. The array is divided into a $2^k \times 2^k$ grid of subframes, each of dimension $2^{k(n-1)} \times 2^{k(n-1)}$; these are the nodes of the tree at level one. Each of these subframes is then subdivided into another $2^k \times 2^k$ grid of subframes, each of dimension $2^{k(n-2)} \times 2^{k(n-2)}$, and they are the nodes at level two descending from the node at level one containing them. This nested grid decomposition process continues until the individual pixels at level n are reached; these are the leaves of the tree, i.e., the set $X_{(4^k, \dots, 4^k)}$. Thus the tree structure $Q(k, n)$ on a $2^{kn} \times 2^{kn}$ pixel set has the structure of a nested grid with successive nested block sizes scaled by a factor of $\frac{1}{2^k}$ in dimension (cf., Figure 5).

Different choices of n and k give different tree structures. For example, a 4×4 image may be given the distinct structures of $Q(2, 1)$ or $Q(1, 2)$. Figure 6 shows the different trees given by these choices.

For the $Q(k, n)$ tree structure on a $2^{kn} \times 2^{kn}$ array we adopt a *quadtree scanning* for converting each image to a row vector. At each node the 4^k nodes descending from it are indexed by $0, 1, \dots, 4^k - 1$ in a clockwise spiral fashion starting from the upper lefthand corner.

Example 1. A 4×4 array with the quadtree structure $Q(2, 1)$ is scanned as:

$$Q(2, 1) \text{ quadtree scanning sequence: } \begin{pmatrix} 0 & 1 & 2 & 3 \\ 11 & 12 & 13 & 4 \\ 10 & 15 & 14 & 5 \\ 9 & 8 & 7 & 6 \end{pmatrix}. \quad (3.15)$$

Once the 4^k nodes descending from each fixed node in the tree have been associated to submatrices within the nested grid with the above ordering, the leaves of the tree (i.e., the individual pixels) have a lexicographic order described in Section 3.2.1. Each image is scanned into a 1×4^{kn} vector in this sequence.

Example 2. When a 4×4 array is given the structure of the tree $Q(1, 2)$, the individual pixels are scanned in the following order:

$$Q(1, 2) \text{ quadtree scanning sequence: } \begin{pmatrix} 0 & 1 & 4 & 5 \\ 3 & 2 & 7 & 6 \\ 12 & 13 & 8 & 9 \\ 15 & 14 & 11 & 10 \end{pmatrix}. \quad (3.16)$$

Here the upper lefthand 2×2 submatrix is the zeroth node at level 1, the upper right hand 2×2 submatrix is the first node at level 1, etc. These nodes at level 1 are then each spirally scanned, resulting in the displayed sequencing.

The group $Z_{(4^k, 4^k, \dots, 4^k)}$ (n factors) acts on the $2^{kn} \times 2^{kn}$ array of pixel positions having tree structure $Q(k, n)$, scanned in the quadtree fashion. For simplicity we denote this wreath product group by $Z(k, n)$. By Theorem 3.1

$$Z(k, n)/B_i \cong Z(k, i) \quad \text{and} \quad Z(k, n)/B_i \text{ acts on the truncated tree } Q(k, i). \quad (3.17)$$

It follows easily from this by recursion that

$$|Z(k, n)| = 4^{k(4^n - 1)/3}. \quad (3.18)$$

Geometrically, the elements of $Z(k, n)$ perform local spiral shifts on images. Each $2^k \times 2^k$ submatrix corresponding to a node at level $n - 1$ has been scanned spirally, and these 4^k points are cyclically permuted by the elements of $Z(k, n)$ fixing that node. For $k \geq 2$, a cyclic permutation of this nature is not a planar symmetry of the $2^k \times 2^k$ submatrix. For example a cyclic permutation of the entries in the 4×4 matrix (3.15) cannot be achieved by a rigid motion of the plane. For $k = 1$, however, the elements of $Z(1, n)$ perform local rotations of $0, 90, 180$ or 270 degrees on submatrices within the nested grid structure. In particular, any 2×2 submatrix whose upper lefthand pixel $a_{p,q}$ has both p and q odd may be rotated by some group element (even by one that fixes all other pixels). For example, the matrix in (3.16) may be rotated clockwise 90 degrees by first rotating each 2×2 submatrix in the nested grid by 90 degrees, and then by permuting the four 2×2 blocks

as a 90 degree rotation of these level one nodes (leaving rigid the ordering within each 2×2 block). Both of these operations are group symmetries of the quadtree.

Rotations of the full image through multiples of 90 degrees are also group symmetries. Most elements of this group, however, do not perform planar symmetries on the image matrix as a whole. Thus, a rotation by 90 degrees of the upper left 2×2 block of the matrix in (3.16), but keeping all other entries fixed, is a group symmetry of the quadtree, but not a rigid motion of the 4×4 image as a whole.

3.3 The WPC group $Z(1, 9)$ acting on 512×512 images

We now apply these concepts to the realistic yet tractable example of the WPC group $Z(1, 9)$ acting on the 9-level tree $Q(1, 9)$ shown in Figure 7; we shall build upon this example throughout this paper and its successor.

Levels of the tree $Q(1, 9)$ are numbered from 0 to 9, with the root node at level 0 and the leaves at level 9. At the i^{th} level the 2^{2i} nodes are numbered from left to right, beginning with node 0. We associate the nodes of the tree to subimages of the image in the following nested grid manner, as depicted in Figure 8: Let Q^0 denote the whole image, and associate to it the root node of the tree at level 0. Label the four 256×256 quadrants in a clockwise fashion as

$$\begin{pmatrix} Q_0^1 & Q_1^1 \\ Q_3^1 & Q_2^1 \end{pmatrix} \quad (3.19)$$

and associate Q_i^1 to the i^{th} node at level 1. Proceeding recursively, suppose Q_i^k is the $2^{9-k} \times 2^{9-k}$ subimage attached to ν_i^k , the i^{th} node at level k . Divide Q_i^k into its four quadrants, ordering them clockwise beginning with the upper left quadrant. These subimages are then associated to the four nodes descending from ν_i^k , and are labeled $Q_{i_0}^{k+1}, \dots, Q_{i_3}^{k+1}$, where the subscripts match the indices of their respective nodes at level $k+1$. Thus for example, the nodes at level 8 of the tree correspond to 2×2 blocks (labeled Q_i^8) within the image, and the leaves of the tree to the individual pixels, Q_i^9 , $i = 0, \dots, 4^9 - 1$.

In this scheme any path in the tree from the root node to a specific leaf represents a sequence of nested arrays, scaled successively by a factor of $1/4$, from the whole image down to the pixel corresponding to the given leaf. By writing the subscript i of any block Q_i^k in base four, the “digits” (with leading zeros included) read from left to right specify the successive branches down the tree of the unique path that starts from the root node and ends at the node corresponding to Q_i^k at level k . Thus, block Q_i^k is seen to lie inside the larger block Q_p^q , where $q \leq k$, precisely when the base four expression for p is given by truncating the right-most $k - q$ digits of the base four expression for i .

Referring again to Figure 7, we next describe the action of the WPC group $Z(1, 9)$ on the tree $Q(1, 9)$ and on 512×512 images. For $k = 0, \dots, 8$ let $\alpha_i^{(k)}$ be the group element that cyclically shifts the four subtrees descending from the i^{th} node at level k , and fixes the remainder of the tree. Equivalently, $\alpha_i^{(k)}$ cyclically and rigidly permutes the four quadrants within the i^{th} subimage (of dimension $2^{9-k} \times 2^{9-k}$), and fixes the remaining portion of the image.

For example, at level 0 the group element $\alpha_0^{(0)}$ cyclically permutes the four nodes at level 1 in a clockwise direction, and rigidly shifts the trees descending from them. Equivalently, $\alpha_0^{(0)}$ acts on images associated to $Q(1, 9)$ by cyclically permuting the four quadrants Q_0^1, \dots, Q_3^1 , but keeping intact these four subimages (and not rotating them). Also, for any i , $0 \leq i \leq 4^8 - 1$, the element $\alpha_i^{(8)}$ acts as a 4-cycle on the leaves of the tree or the pixels of an image by rotating the appropriate 2×2 subimage and fixing all other pixels. Every element of $Z(1, 9)$ may be written as a product of these permutations (where we allow repeated applications, or powers, of a given group element). Equivalently, $Z(1, 9)$ is generated by these elements. In fact, it is easy to see that:

$$Z(1, 9) \text{ is generated by } \alpha_0^{(0)}, \alpha_0^{(1)}, \dots, \alpha_0^{(8)} \quad (3.20)$$

and these are a minimal generating set (the minimality follows from Burnside's Basis Theorem, cf., Exercise 26, Section 6.1 in [22]).

4 FOURIER ANALYSIS AND MULTIREOLUTION

In this section we shall describe the decomposition of the space $L(X_{\mathbf{m}})$ of discrete signals indexed by the leaves of the tree $\mathcal{T}_{\mathbf{m}}$, into a sum of irreducible subspaces under the action of the WPC group $Z_{\mathbf{m}}$. From this we shall obtain a wreath product based discrete Fourier series, and a multiresolution-like decomposition of the space. The natural level structure of the tree—as reflected by the sequence of normal subgroups (3.11) of the group—will imply that the decomposition may be implemented by a multichannel pyramid DFT filter bank (in polyphase form). In order to describe this precisely we first continue the brief outline of the theoretical underpinnings of finite group-based signal processing begun in Section 3.1.

4.1 Invariant decompositions

Continuing the development in Section 3.1, let G denote an arbitrary finite group acting transitively on the set X . We describe how the left (or right) regular representation of G on $L(G)$ decomposes, and then use this decomposition to infer the spectral decomposition of G acting on the subspace $L(G/G_0) = L(X)$ of all discrete signals of length N . When there is a natural hierarchy to the subspaces, this decomposition can give a finite, multiresolution analysis of each discrete signal $f \in L(X)$ (as discussed in [46]). This is the case for the iterated wreath products, and here the decomposition is a finite analog of the multiresolution analysis usually associated to a wavelet decomposition. The proofs of results cited in this section may be found in standard references on representation theory such as [59] or [22].

The decomposition of the regular representation of G described in Example 0 of Section 3.1 is given by the classical Wedderburn's Theorem:

Theorem 4.1 *Let G be a finite group acting on $L(G)$ by right (or left) translation. Then $L(G)$ has an orthogonal G -invariant decomposition as*

$$L(G) = M_0 \oplus M_1 \oplus \dots \oplus M_{r-1} \quad (4.1)$$

where each M_k is a uniquely determined G -invariant subspace of dimension d_k^2 . Each M_k further decomposes as a direct sum of d_k irreducible G -invariant subspaces. If I_k is any one of the irreducible subspaces in M_k , then the representation of G on M_k is equivalent to (i.e., the same up to a change of basis) the representation of G on the direct sum of d_k copies of I_k . We denote this by

$$M_k \cong I_k \oplus I_k \oplus \cdots \oplus I_k \quad (d_k \text{ factors}). \quad (4.2)$$

The subspaces M_k are called the *isotypic components* of $L(G)$. If G is abelian, then each d_k is one, $M_k = I_k$, and the M_k 's are the simultaneous eigenspaces of all elements of G . If G is non-abelian, however, some of the d_k are at least two, the corresponding M_k have dimension at least four, and these M_k are necessarily reducible (by (4.2)).

Every irreducible representation of G is equivalent to one of I_0, I_1, \dots, I_{r-1} , and no pair of these are equivalent. The number, r , of inequivalent irreducible representations of G equals the number of conjugacy classes of G (which equals the order of G when G is abelian).

For arbitrary finite groups the projection map π_k of $L(G)$ onto the (not necessarily irreducible nor one-dimensional) isotypic component M_k is described as follows. Each irreducible subspace I_k within M_k has an associated irreducible character χ_k . The character χ_k is a complex-valued function on G whose value on the group element α is the trace of any matrix representing the linear transformation of α acting on I_k (which is a G -invariant subspace of $L(G)$).

Theorem 4.2 (See e.g. [59], Theorem 8 (ii), Section 2.6) For any $f \in L(G)$ the projection of f onto M_k is then given by

$$\pi_k(f)(\alpha) = \frac{d_k}{|G|} \sum_{\beta \in G} \overline{\chi_k(\alpha\beta^{-1})} f(\beta), \quad \text{for all } \alpha \in G \quad (4.3)$$

where the overbar denotes complex conjugation.

Example 0. Consider the action of G on $L(G)$ defined in Example 0 in Section 3.1 (the right or left regular representation). With a change of variables (4.3) may be rewritten in operator form as a generalized Fourier transform:

$$\pi_k(f) = \frac{d_k}{|G|} \sum_{\beta \in G} \overline{\chi_k(\beta)} \beta f. \quad (4.4)$$

The (generalized) *Fourier series* or *spectral decomposition* of f with respect to the group G is

$$f = \sum_{k=0}^{r-1} \pi_k(f). \quad (4.5)$$

Example 1. Let G be the cyclic group of order N generated by σ , acting on a set X of N points by cyclic shift. Then G and X may be identified with each other, and similarly, $L(G)$ with $L(X)$. Let $f(n)$ denote $f(\sigma^n) = f(x_n)$. In this case the decomposition in (4.4) reduces to the N -point DFT decomposition: Let W

be the N^{th} root of unity $e^{2\pi j/N}$ where $j = \sqrt{-1}$. Then the sampled exponential function $E^k(n) = W^{nk}$ is an eigenfunction for σ , and is a basis for the one-dimensional space M_k . The k^{th} coefficient of a signal $f(n) \in L(X)$ expressed in this basis is given by its k^{th} discrete Fourier coefficient:

$$\hat{f}_k = \sum_{n=0}^{N-1} f(n) \overline{E^k(n)}, \quad 0 \leq k \leq N-1 \quad (4.6)$$

where the overbar denotes complex conjugation. Notice that this is simply the inner product of f with E_k , which is N times the projection of f onto E_k . The decomposition (4.1) is the representation of each f with respect to the basis of eigenfunctions, i.e., the discrete Fourier series for f :

$$f(n) = \frac{1}{N} \sum_{k=0}^{N-1} \hat{f}_k E^k(n). \quad (4.7)$$

The function $\frac{1}{N} \hat{f}_k E^k$ is the projection, $\pi_k(f)$, of f onto M_k , the one-dimensional eigenspace for G .

Example 2. For the two-dimensional DFT, we also proceed as in Example 2 of Section 3.1. In this case the Fourier basis is the matrix of sampled exponential products, $((E_m(i)E_n(k))_{i,k})$. This corresponds to taking a tensor product of the sampled exponentials as the basis. Projection onto this basis is the usual two-dimensional DFT.

Examples 1 and 2 are abelian, while Example 0 describes the general case. Section 4.2 is devoted to examining the particular non-abelian case of the WPC groups $Z_{\mathbf{m}}$ in some detail.

Theorem 4.1 describes the decomposition of the regular representation, when G acts on the space $L(G)$ of signals indexed by G itself. From it we may infer the decomposition of G acting transitively on any set X (eg. $Z_{\mathbf{m}}$ acting on the set of leaves of a tree):

Theorem 4.3 *The decomposition of any G -invariant subspace U of $L(G)$ into its isotypic components, $U \cap M_k$, can be computed by applying to U each of the projection maps π_k to obtain the spectral decomposition of U . In particular, when $U = L(G/G_0) = L(X)$ this gives the spectral decomposition of $L(X)$ and the Fourier series of each $f \in L(X)$ with respect to G .*

As noted previously, the isotypic components M_k for the regular representation need not be irreducible, and in general there need not be a natural choice of basis functions for M_k . It may still be the case, however, that the k^{th} isotypic component of the G -invariant subspace U is irreducible. We shall see in this situation that there is a generator of this component of U (where the irreducible component is spanned by G -translates of this function), and this generator plays a critical role in the convolution of functions (discussed in [50]). We say that the representation of G on the invariant subspace U is *multiplicity-free* if every nonzero isotypic component of U is irreducible. The family of multiplicity-free representations includes many groups acting in various natural ways as permutations on sets X , i.e., for many signal spaces $L(X)$. In particular, this will be the case for the wreath product groups $Z_{\mathbf{m}}$ acting on $L(X_{\mathbf{m}})$.

When the isotypic or irreducible components have a natural order or hierarchy, then decomposing a discrete signal $f \in L(X)$ into its isotypic or irreducible components can give a natural (finite) multiresolution decomposition of f . Each of these components (“harmonics”) of f contains some of the information of f , and one of the challenges of this group-based approach is to determine which groups and which isotypic (or irreducible) subspaces for some action of G on the spatial domain $L(X)$ carry concentrations of specific kinds of information. We shall see for wreath product groups that the multiresolution analysis has additional structure analogous to a finite wavelet decomposition.

Remark. For many groups there exist efficient algorithms for computing the spectral decomposition of any $f \in L(X)$. See the papers [48, 49] for recent progress in this area.

4.2 A Multiresolution approach for Fourier analysis on wreath products

Continuing in the notation of Section 3.2 let $\mathbf{m} = (m_0, \dots, m_{n-1})$ be an n -tuple of integers, all at least two. For the moment let G be any wreath product group (3.10) acting on the tree $\mathcal{T}_{\mathbf{m}}$, and let $L(X_{\mathbf{m}})$ denote the space of all complex valued functions on the set $X_{\mathbf{m}}$ of leaves of $\mathcal{T}_{\mathbf{m}}$. Representations of wreath product groups have been studied in [40].

We first describe a G -invariant filtration of the spatial domain $L(X_{\mathbf{m}})$.

Theorem 4.4 *Set $V_n = L(X_{\mathbf{m}})$ and for $0 \leq i \leq n - 1$ let V_i be the space of all functions that are constant on each block of $m_i m_{i+1} \cdots m_{n-1}$ leaves that descend from a common node at level i . (Thus V_0 is the one-dimensional space of functions that are constant on all leaves.) These subspaces form a complete filtration of $L(X_{\mathbf{m}})$,*

$$0 \subset V_0 \subset V_1 \subset V_2 \subset \cdots \subset V_n = L(X_{\mathbf{m}}). \quad (4.8)$$

Moreover, the spaces V_i are all G -invariant, and B_i is the subgroup of G that acts trivially on V_i .

Example 1. For $\mathbf{m} = (3, 4)$ ($n = 2$) as depicted in Figure 4, then V_1 is the set of functions f such that $f(x_0) = f(x_1) = f(x_2) = f(x_3)$, and $f(x_4) = \cdots = f(x_7)$, and $f(x_8) = \cdots = f(x_{11})$.

Example 2. Let $\mathcal{T}_{\mathbf{m}}$ be the quadtree $Q(k, n)$ structure on the spatial domain of $2^{kn} \times 2^{kn}$ images scanned in the quadtree manner. Then V_i consists of images which have constant color/intensity on each $2^{k(n-i)} \times 2^{k(n-i)}$ subimage within the nested grid.

In general, if ν_i^j is a node of the tree at level j let $e_{j,i}$ be the function that is 1 on all leaves below $\nu_{j,i}$ and zero on all other nodes. Then the unit “step functions” $e_{j,i}$ form a basis of V_j that is orthogonal (disjoint support) under the usual inner product. The collection $\{e_{n,i} \mid i = 0, \dots, m_0 \cdots m_{n-1} - 1\}$ is the basis of $L(X_{\mathbf{m}})$ of delta functions on the set of leaves. These bases are “shift invariant” under the group action:

Lemma 4.5 *For each j the group G permutes the basis $\{e_{j,0}, e_{j,1}, \dots, e_{j,m_0 \cdots m_{i-1}-1}\}$ of V_j .*

There is a natural notion of “scaling” among these basis functions given by the level of the tree from which they derive. Each step function $e_{j,i}$ is the sum of all step functions $e_{j+1,i'}$ corresponding to the nodes ν_i^{j+1} at level $j+1$ descending from ν_i^j . The number of leaves in the support of $e_{j,i}$ is m_j times the support size of $e_{j+1,i'}$. We shall see that for certain groups G there are similar relations between other choices of bases at different scales i.e., they are “generalized Haar wavelets”.

Next we exhibit a G -invariant decomposition⁴ of V_{j+1} as $V_j \oplus W_j$. For each j define the j^{th} (*normalized Radon transform* or augmentation map as follows. For each leaf x in the tree let $\nu_j(x)$ be the unique node at level j above x , and let $A_j(x)$ be the set of all leaves that lie below $\nu_j(x)$ in the tree. For example, if $X_{\mathbf{m}}$ indexes the space of $2^{kn} \times 2^{kn}$ discrete images with quadtree structure $Q(k, n)$, then each leaf x represents a pixel position; $\nu_j(x)$ is the particular $2^{k(n-j)} \times 2^{k(n-j)}$ subframe in the nested grid that contains pixel x and $A_j(x)$ is the set of all pixels within that subframe. The j^{th} Radon transform is a map,

$$\mathcal{R}_j : L(X_{\mathbf{m}}) \longrightarrow V_j \quad \text{defined by} \quad \mathcal{R}_j(f)(x) = \frac{1}{|A_j(x)|} \sum_{x_n \in A_j(x)} f(x_n). \quad (4.9)$$

In other words, $\mathcal{R}_j(f)$ is the function that is constant on each $A_j(x)$ with value at x equal to the average value of f on $A_j(x)$. For fixed j all the blocks $A_j(x)$ have the same size, $m_j m_{j+1} \cdots m_{n-1}$, so the averaging constant depends only on the level j . Note that each Radon transform is a linear transformation, and is G -equivariant. The latter means that,

$$\alpha \mathcal{R}_j(f) = \mathcal{R}_j(\alpha f), \quad \text{for all } \alpha \in G. \quad (4.10)$$

Evidently if $f \in V_j$ then $\mathcal{R}_j(f) = f$, so \mathcal{R}_j is the projection of $L(X_{\mathbf{m}})$ onto V_j . Now \mathcal{R}_j also projects V_{j+1} to V_j . Let W_j be the nullspace of the Radon transform from V_{j+1} to V_j , so $W_j = \{f \in V_{j+1} \mid \mathcal{R}_j(f) = 0\}$. Thus W_j consists of those functions with zero average value on each block $A_j(x)$. Using (4.10) it is not difficult to show the following.

Theorem 4.6 *The Radon transform gives a G -invariant decomposition of V_{j+1} as*

$$V_{j+1} = V_j \oplus W_j, \quad \text{for all } j, \quad (4.11)$$

where W_j is the nullspace of \mathcal{R}_j .

This gives a general description of a multiresolution decomposition of $L(X_{\mathbf{m}})$ for any wreath product group G . The further decomposition of $L(X_{\mathbf{m}})$ into irreducible G -invariant subspaces depends on the specific factors H_j in (3.10) that comprise G .

Theorem 4.6 provides a recursive decomposition of V_{j+1} which we use to determine its irreducible components. Now V_j may be considered as the space of all functions on the leaves of the truncated tree $\mathcal{T}_{(m_0, \dots, m_{j-1})}$ obtained from $\mathcal{T}_{\mathbf{m}}$ by deleting all nodes at levels greater than j ; a function f in V_j is deemed to have the value on node ν at level j equal to the sum of all its (equal) values on leaves beneath ν . The normal subgroup B_j of

⁴The existence of a G -invariant complement W_j to V_j is guaranteed by Maschke's Theorem from representation theory, however we explicitly construct it.

G acts trivially on V_j , so by Theorem 3.1 the quotient group G/B_j is a wreath product group acting on this truncated tree and in a corresponding way on its spatial domain V_j . Thus the decomposition of V_j reduces to the problem of decomposing a spatial domain with a tree structure and wreath product group action, but of smaller sizes. In this way there is a natural recursive methodology for decomposing the components of V_j . We now focus on finding the irreducible decomposition of subspaces W_j in the special case of a wreath product of cyclic groups. At the conclusion we shall recursively assemble the pieces and describe the overall decomposition in this case. We begin by revisiting the usual one-dimensional DFT, interpreted as the tree of one level with cyclic symmetry group.

Theorem 4.7 *Assume G is the wreath product of cyclic groups $Z_{\mathbf{m}}$ (cf., (3.14)). When the tree has only one level (i.e., $n = 1$) G is the cyclic group Z_{m_0} cyclically permuting the m_0 leaves of the tree, and the decomposition of $L(X_{\mathbf{m}})$ is given by the usual cyclic DFT expansion of a function as its discrete Fourier series (cf., the remarks preceding (4.7)):*

$$L(X_{(m_0)}) = V_0 \oplus W_{0,1} \oplus W_{0,2} \oplus \cdots \oplus W_{0,m_0-1}. \quad (4.12)$$

where each subspace is one-dimensional and G -invariant. The projection of f onto $W_{0,k}$ is given by the k^{th} Fourier coefficient of f for the m_0 -point DFT; the projection of f onto V_0 is given by the 0^{th} Fourier coefficient, which is the sum of the values of f .

The decomposition of $L(X_{\mathbf{m}})$ for $n > 1$ is thus seen to be a natural (non-abelian) generalization of the discrete Fourier transform.

Theorem 4.8 *For $G = Z_{\mathbf{m}}$, the vector space $L(X_{\mathbf{m}})$ decomposes into G -invariant subspaces as*

$$\begin{aligned} L(X_{\mathbf{m}}) &= V_n = V_{n-1} \oplus W_{n-1} \\ &= V_{n-1} \oplus W_{n-1,1} \oplus W_{n-1,2} \oplus \cdots \oplus W_{n-1,m_{n-1}-1} \end{aligned} \quad (4.13)$$

where V_{n-1} and W_{n-1} are as in Theorem 4.6. Furthermore, $W_{n-1,k}$ is irreducible and

$$\dim W_{n-1,k} = s = m_0 m_1 \cdots m_{n-2}, \text{ for every } k. \quad (4.14)$$

In (4.14) s equals the number of nodes of the tree at level $n - 1$. Also, $W_{n-1,k}$ and $W_{n-1,k'}$ are inequivalent representations whenever $k \neq k'$.

As noted earlier, the decomposition of V_{n-1} into irreducibles may be arrived at recursively.

Theorem 4.9 *In the notation of Theorem 4.8 for $\mathbf{m} = (m_0, \dots, m_{n-1})$, let $\mathbf{m}' = (m_0, m_1, \dots, m_{n-2})$. Then G/B_{n-1} is isomorphic to the wreath product of cyclic groups $Z_{\mathbf{m}'}$ acting on the truncated tree $\mathcal{T}_{\mathbf{m}'}$ with leaves $X_{\mathbf{m}'}$. Under this identification the Radon transform \mathcal{R}_{n-1} gives that $V_{n-1} = L(X_{\mathbf{m}'})$, and hence, this space may be decomposed recursively according to Theorems 4.7 and 4.8.*

Corollary 4.10 *The representation of $Z_{\mathbf{m}}$ on $L(X_{\mathbf{m}})$ is multiplicity-free.*

This follows, since by Theorem 4.8 the irreducible spaces $W_{j,k}$ are all inequivalent. For fixed j and different k this was noted above; and for different values of j they have different dimensions.

By Theorem 4.9 it suffices to describe bases of each $W_{n-1,k}$ defined in Theorem 4.8, and to give the coefficients of the expansion of any $f \in L(X_{\mathbf{m}})$ with respect to these basis functions.

Let W be the $(m_{n-1})^{\text{th}}$ root of unity $e^{2\pi j/m_{n-1}}$ where $j = \sqrt{-1}$. Then for each $k \geq 1$ the signal

$$h_k(x_p) = \begin{cases} W^{pk}, & 0 \leq p \leq m_{n-1} - 1 \\ 0, & p \geq m_{n-1} \end{cases} \quad (4.15)$$

is a basis vector whose independent G -translates form a basis of $W_{n-1,k}$. Note that h_k is supported on the initial m_{n-1} leaves of the tree, i.e., on the leaves descending from the zeroth node of the tree at level $n-1$. If $h_{k,q}$ is the corresponding impulse response filter supported on the ordered set of leaves descending from the q^{th} node at level $n-1$, then $h_k = h_{k,0}$ and $h_{k,0}, h_{k,1}, \dots, h_{k,s-1}$ is a basis of $W_{n-1,k}$ (where again s is the number of nodes at level $n-1$ as in (4.14)). (The proof that these functions form a basis is outlined in a remark at the end of this subsection.)

We now observe that the wreath product decomposition also generalizes the one-dimensional Haar decomposition for signals of length 2^n .

Theorem 4.11 *In the special case when $\mathcal{T}_{\mathbf{m}}$ is the binary tree ($\mathbf{m} = (2, 2, \dots, 2)$), in which case G is the wreath product of cyclic groups of order 2, the basis functions $h_{n-1,q}$ are precisely the discrete Haar wavelets at scale (level) n . The function $h_{n-1} = h_{n-1,0}$ is the mother wavelet, and the other wavelets at level n (i.e., at the highest level of detail) are its G -translates. The wavelets at scale j are obtained by applying the Radon transform \mathcal{R}_j to these.*

For a general WPC group, the projection of a signal f onto $W_{n-1,k}$ is obtained from the m_{n-1} -point DFTs of f on each of the s sets of leaves descending from the same nodes at level $n-1$. For the q^{th} node at level $n-1$ let $\hat{f}_{k,q}$ be the k^{th} coefficient of the m_{n-1} -point DFT of f on these (ordered) leaves, i.e.,

$$\hat{f}_{k,q} = \sum_{i=0}^{m_{n-1}-1} f(x_{qm_{n-1}+i})W^{-ik}. \quad (4.16)$$

Then the projection of f into $W_{n-1,k}$ is given by

$$\pi_{n-1,k}(f) = \sum_{q=0}^{s-1} \hat{f}_{k,q} h_{k,q}. \quad (4.17)$$

For $k \geq 1$ the coefficients $\hat{f}_{k,q}$ are the k^{th} wreath product Fourier coefficients of f at level $n-1$. The collection of all wreath product Fourier coefficients is called the *wreath product spectrum*, or *WPC spectrum*, of f .

The wreath product Fourier coefficients of f at all levels can now be obtained recursively, and in this way the WPC spectrum of f is computed. We describe how this may be computed effectively through multichannel filter banks.

As usual, let $f \in L(X_{\mathbf{m}})$ be a discrete signal of length $m_0 m_1 \cdots m_{n-1}$, let $s = m_0 m_1 \cdots m_{n-2}$. Using a block transform DFT filter bank (in polyphase form) the signal is split into s disjoint blocks, each of size m_{n-1} (where each block corresponds to a set of leaves descending from a common node at level $n-1$). An m_{n-1} -point DFT is applied to each block. For each $k \in \{0, 1, \dots, m_{n-1} - 1\}$ the k^{th} output coefficient is stored in a vector $w_{n-1,k}$ of length s . For $k \geq 1$ the highpass coefficient vectors, $w_{n-1,k}$, are the projections of f onto $W_{n-1,k}$ and the coefficients in these vectors are the wreath product Fourier coefficients. The lowpass coefficient vector $w_{n-1,0}$ is obtained from the 0^{th} DFT coefficients. This lowpass vector is now fed to an m_{n-2} -point block DFT filter bank by splitting it into $m_0 \cdots m_{n-3}$ blocks, each of size m_{n-2} . This pyramid process is continued, at each stage storing the highpass vectors $w_{i,k}$ for $k \geq 1$ and continuing to filter the lowpass vector $w_{i,0}$. The pyramid terminates after n stages, when the single block $w_{1,0}$ of length m_0 is passed through a m_0 -point DFT filter bank. The 0^{th} (lowpass) coefficient of $w_{1,0}$ is seen to be the sum of all values of f . We shall refer to this algorithm as an **m-channel DFT filter bank**. In this language we restate Theorem 4.11 as

Theorem 4.12 *In the case of a binary tree, the $(2, 2, \dots, 2)$ -channel DFT filter bank is the Haar transform of a discrete signal.*

The inverse wreath product Fourier transform or synthesis algorithm is also easily described. The synthesis stage of the **m-channel DFT filter bank** algorithm gives perfect reconstruction. This algorithm is completely analogous to the analysis algorithm, but uses inverse DFTs (which involve appropriate normalizing factors)—the details are omitted.

Remark. As mentioned at the end of Section 3.1, in representation theoretic terms, the representation of G on $L(X_{\mathbf{m}}) = L(G/G_0)$ is the principal (trivial) representation, I , of the subgroup G_0 induced to G , denoted as $Ind_{G_0}^G(I)$. This may be computed by first inducing I from G_0 to the larger subgroup $G_1 = G_0 B_{n-1}$, which is the stabilizer in G of the zeroth node at level $n-1$, to obtain:

$$Ind_{G_0}^{G_1}(I) = J_0 + J_1 + \cdots + J_{m_{n-1}-1} \quad (4.18)$$

where each J_k is a one-dimensional representation and J_0 is the trivial representation of G_1 . It then follows that $V_{n-1} = Ind_{G_1}^G(J_0)$ and $W_{n-1,k} = Ind_{G_1}^G(J_k)$ for $k \geq 1$. From the formulas for induced representations the characters of these representations may be explicitly computed and hence the projections of each signal f onto the subspaces can be computed as well.

4.3 An explicit realization of the one-dimensional spectral transform

While the WPC spectrum of a signal f in $L(X_{\mathbf{m}})$ can be determined and evaluated by the fast algorithm described above, it may also be seen in terms of a one-dimensional unitary block transform which we shall call the *wreath product transform* or WPT. This perspective provides additional insight into the multiresolution structure of the decomposition.

As an illustration with reference to Figure 9(a) and (b), we consider the WPT on the vector space $L(1, 2)$ of signals indexed by the tree $\mathcal{T}_{1,2}$. Here signals are of length $M = 16$. Elements in $L(1, 2)$ may also be considered as 4×4 images scanned in the $k = 1, n = 2$ quadtree fashion, as displayed explicitly in (3.16).

The WPC group $Z(1, 2)$ decomposes the 4^2 -dimensional vector space $L(1, 2)$ into irreducible, group invariant subspaces,

$$L(1, 2) = V_0 \oplus W_{0,1} \oplus W_{0,2} \oplus W_{0,3} \oplus W_{1,1} \oplus W_{1,2} \oplus W_{1,3} \quad (4.19)$$

where $V_0, W_{0,1}, W_{0,2}, W_{0,3}$ are each one-dimensional, and $W_{1,1}, W_{1,2}, W_{1,3}$ are each four-dimensional. Explicitly, each $W_{1,k}$ is obtained from the k^{th} highpass coefficients of a 4-point DFT block transform; and each $W_{0,k}$ is then the k^{th} coefficient of a 4-point DFT on the four DC terms obtained from the four blocks. Thus the projections onto $W_{1,1}, W_{1,2}, W_{1,3}$ are decompositions at the detail level, while those onto $W_{0,1}, W_{0,2}, W_{0,3}$ are decompositions at the next coarsest detail level; projection onto V_0 provides the coarsest approximation, giving the 16 times the average value of the signal f . This spectral (or generalized Fourier) decomposition can be shown in terms of a block transform, \mathbf{A} , on signals $\mathbf{f} = (f(0), \dots, f(15))$:

$$\mathbf{f} = \mathbf{A}\hat{\mathbf{f}} \implies f(p) = \sum_{q=0}^{M-1} a(p, q)\hat{f}(q), \quad 0 \leq p \leq M-1 \quad (4.20)$$

where,

$$\mathbf{A} = \left[\begin{array}{cccc|cccc|cccc|cccc} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -i & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & -i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & i & 0 & 0 & 0 \\ 1 & i & -1 & -i & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & i & -1 & -i & 0 & i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -i & 0 & 0 \\ 1 & i & -1 & -i & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ 1 & i & -1 & -i & 0 & -i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & i & 0 & 0 \\ 1 & -1 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & -1 & 1 & -1 & 0 & 0 & i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -i & 0 \\ 1 & -1 & 1 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 \\ 1 & -1 & 1 & -1 & 0 & 0 & -i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -i & 0 \\ 1 & -i & -1 & i & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & -i & -1 & i & 0 & 0 & 0 & i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -i \\ 1 & -i & -1 & i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 \\ 1 & -i & -1 & i & 0 & 0 & 0 & -i & 0 & 0 & 0 & -1 & 0 & 0 & 0 & i \end{array} \right]$$

$$= [\mathbf{a}_0, \mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_{15}] \quad (4.21)$$

and $\hat{f}(p)$ is the p^{th} generalized Fourier coefficient of \mathbf{f} . Here \mathbf{a}_j denotes the $j+1^{\text{st}}$ column of \mathbf{A} . Thus \mathbf{A} is the synthesis matrix of a wavelet pyramid algorithm for finite sized sequences, and its inverse represents the analysis stage. In order to carry out convolution products (presented in [50]) the column vectors \mathbf{a}_i are not normalized. However, they are orthogonal (of varying lengths). Thus $\mathbf{A}\mathbf{A}^{*\text{T}}$ is the diagonal matrix \mathbf{D} whose i, i entry is $\|\mathbf{a}_i\|^2$. Thus the analysis matrix $\mathbf{A}^{-1} = \mathbf{A}^{*\text{T}}\mathbf{D}^{-1}$ is easily computed, and it gives the spectral transformation:

$$\hat{\mathbf{f}} = \mathbf{A}^{*\text{T}}\mathbf{D}^{-1}\mathbf{f} \implies \hat{f}(p) = \sum_{q=0}^{M-1} b^*(p, q)f(q), \quad \text{for } 0 \leq p \leq M-1. \quad (4.22)$$

Columns of \mathbf{A} constitute an orthogonal basis for $L(1, 2)$: \mathbf{a}_0 is a basis for the one-dimensional space of constants V_0 ; columns $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are bases for the respective one-dimensional subspaces $W_{0,1}, W_{0,2}$ and $W_{0,3}$. Scaled and translated versions of the prototype \mathbf{a}_k span respectively the (detail level) space $W_{1,k}$ for $k = 1, 2, 3$. More specifically, \mathbf{a}_4 to \mathbf{a}_7 span $W_{1,1}$, \mathbf{a}_8 to \mathbf{a}_{11} span $W_{1,2}$, and \mathbf{a}_{12} to \mathbf{a}_{15} span $W_{1,3}$. Basis functions $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are progressively scaled by a factor of $\frac{1}{4}$ and translated by increments equal to their support, so that the set of all basis functions at any scale completely cover the interval. No “lapped” effect is present. Note that the basis functions are structured and that the transform, being based on cyclic groups, is nonexpansive.

As we saw in the discussion preceding Theorem 4.11, for general k and n the associated transform matrix \mathbf{A} has basis functions (its orthogonal columns) containing powers of the root of unity $e^{2\pi j/4^k}$, and these functions are progressively scaled by 4^{-k} . Again, each of the scaled basis functions are translated by increments equal to their support, and the complete set of basis functions at any scale cover the interval. For normalized \mathbf{A} , we obtain a unitary transform.

5 WPC GROUP BASED IMAGE PROCESSING

In this section we consider the space $L(k, n)$ of discrete signals indexed by the leaves of the tree $\mathcal{T}_{(4^k, \dots, 4^k)}$ as this space decomposes under the action of the WPC group $Z(k, n)$, and we apply the theory to the special case where the signals represent $2^{kn} \times 2^{kn}$ discrete images with the quadtree structure $Q(k, n)$, scanned in the quadtree fashion. We first show that the WPC spectrum of such images may also be stored in a “quadtree” fashion that preserves locality: the so-called *quadtree spectrum*. We then examine, by way of the space of 512×512 images discussed in Section 4, the action of individual elements of the WPC group on images and their corresponding action on the quadtree spectra of images. Finally, we again relate these observations to the one- and two-dimensional Haar spectra.

5.1 The quadtree spectrum of an image

The spectral coefficients may be grouped and stored in various ways to illuminate relations between the original signal and its spectrum. We describe one such scheme for a $2^{kn} \times 2^{kn}$ image f scanned into a one-dimensional signal in $L(k, n)$ with the quadtree indexing scheme. This quadtree spectrum of f is denoted by $\mathcal{Q}(f)$. In this scheme the WPC spectrum of f is displayed as a $2^{kn} \times 2^{kn}$ matrix in such a way that the projections of f onto the irreducible, WPC group-invariant components (i.e., its multiresolution components) appear as a nested grid of subimages within this spectral matrix. Moreover, the WPC Fourier coefficients of f from a given irreducible subspace $W_{j,i}$ are stored together in a $2^{kj} \times 2^{kj}$ subimage so that locality of the original image is preserved: (1) each pixel value within the spectrum depends only on a block within the original image at a corresponding location, (2) different pixels in the same spectral subimage depend on disjoint blocks, and (3) adjacent pixels depend on adjacent blocks. Furthermore, the projections of the image into the multiresolution subspaces V_j in Theorem 4.8 (i.e., the spectral coefficient vectors $w_{j,0}$ described after Theorem 4.11) appear as nested increasing subimages of size $2^{kj} \times 2^{kj}$ in the upper lefthand corner of this (quadtree) spectral matrix.

This is accomplished as follows: First store each of the vectors $w_{n-1,p}$ in a $2^{k(n-1)} \times 2^{k(n-1)}$ array A_p in a clockwise spiral manner, beginning in the upper lefthand corner, for $p = 1, 2, \dots, 4^k - 1$. Then position the arrays A_p as submatrices of a $2^{kn} \times 2^{kn}$ matrix $\mathcal{Q}(f)$ in clockwise spiral manner, but leaving the upper lefthand corner block empty and starting with the upper secondmost lefthand block position. The upper lefthand $2^{k(n-1)} \times 2^{k(n-1)}$ submatrix of $\mathcal{Q}(f)$ is now recursively filled in with the Fourier coefficient vectors for the lower scales, inserted in quadtree manner. The final upper lefthand $2^k \times 2^k$ submatrix then contains $w_{1,0}$; and the 1,1 entry of $\mathcal{Q}(f)$ is the sum of all values of f . We shall refer to $\mathcal{Q}(f)$ as the *quadtree Fourier series* or *quadtree spectrum* of f .

Example 1. Let $f = (a_{i,j})$ be a 4×4 matrix as depicted in (3.16). Here $n = 2$ and $\mathbf{m} = (4, 4)$. The first step in the $(4, 4)$ -channel DFT filter bank algorithm divides f into four blocks of size 4; in the quadtree algorithm these are a 2×2 grid of 2×2 submatrices:

$$f = \begin{pmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{pmatrix}.$$

The quadtree spectrum of f is then

$$Q(f) = \begin{pmatrix} Q_0 & Q_1 \\ Q_3 & Q_2 \end{pmatrix}.$$

For $i = 1, 2, 3$, Q_i is the 2×2 matrix whose p, q entry is the i^{th} coefficient of the 4-point DFT applied to $A_{p,q}$. Thus letting $W = e^{2\pi j/4}$, $j = \sqrt{-1}$ and overbar denoting complex conjugation

$$Q_1 = \begin{pmatrix} a_{1,1} + a_{1,2}\overline{W} + a_{2,2}\overline{W}^2 + a_{2,1}\overline{W}^3 & a_{1,3} + a_{1,4}\overline{W} + a_{2,4}\overline{W}^2 + a_{2,3}\overline{W}^3 \\ a_{3,1} + a_{3,2}\overline{W} + a_{4,2}\overline{W}^2 + a_{4,1}\overline{W}^3 & a_{3,3} + a_{3,4}\overline{W} + a_{4,4}\overline{W}^2 + a_{4,3}\overline{W}^3 \end{pmatrix}.$$

The matrices Q_2 and Q_3 are computed using the same formulas for the entries, but with W replaced by W^2 and W^3 , respectively. Likewise let Q'_0 be the 2×2 matrix whose p, q entry, $s_{p,q}$, is the 0^{th} coefficient of the 4-point DFT of $A_{p,q}$, i.e., the sum of the entries of $A_{p,q}$. The recursive step in the quadtree Fourier algorithm then gives the entries of Q_0 as the coefficients of the 4-point DFT applied to Q'_0 :

$$Q_0 = \begin{pmatrix} s_{1,1} + s_{1,2}\overline{W} + s_{2,2}\overline{W}^2 + s_{2,1}\overline{W}^3 & s_{1,3} + s_{1,4}\overline{W} + s_{2,4}\overline{W}^2 + s_{2,3}\overline{W}^3 \\ s_{3,1} + s_{3,2}\overline{W} + s_{4,2}\overline{W}^2 + s_{4,1}\overline{W}^3 & s_{3,3} + s_{3,4}\overline{W} + s_{4,4}\overline{W}^2 + s_{4,3}\overline{W}^3 \end{pmatrix}.$$

This gives the complete formula for the $Q(1, 2)$ quadtree Fourier series algorithm.

For general $2^{kn} \times 2^{kn}$ images calculation of the quadtree spectrum is especially amenable to a recursive algorithm which we now describe. Note also that (ignoring delays in sequential signals) the blocks may be processed independently, so that the algorithm is naturally parallelizable.

For clarity we describe the algorithm only for $k = 1$; extending the algorithm to larger k is transparent. Figure 11 shows an example of such a spectrum. Suppose that the image f is the $2^n \times 2^n$ matrix $(a_{p,q})$ with complex entries. The entries of the $2^n \times 2^n$ matrix $(b_{p,q}) = Q(f)$ are determined as follows: f is partitioned into blocks of size 4, namely its 2×2 submatrices, each of whose upper lefthand entry has both indices odd. Scan each of these 2×2 submatrices $\begin{pmatrix} a_{2s-1,2t-1} & a_{2s-1,2t} \\ a_{2s,2t-1} & a_{2s,2t} \end{pmatrix}$ in a clockwise spiral fashion beginning with the upper left corner to form the vector

$$v_{s,t} = (a_{2s-1,2t-1}, a_{2s-1,2t}, a_{2s,2t}, a_{2s,2t-1}) \quad (5.1)$$

for $1 \leq s, t \leq 2^{n-1}$. Compute the Fourier coefficients, $\widehat{v}_{s,t}(p)$, of a 4-point DFT on $v_{s,t}$:

$$\widehat{v}_{s,t}(p) = a_{2s-1,2t-1} + a_{2s-1,2t} W^p + a_{2s,2t} W^{2p} + a_{2s,2t-1} W^{3p} \quad (5.2)$$

where $p = 0, 1, 2, 3$. The three “highpass” DFT coefficients are inserted in $Q(f)$ in positions

$$\begin{aligned} b_{s,2^{n-1}+t} &= \widehat{v}_{s,t}(1) \\ b_{2^{n-1}+s,2^{n-1}+t} &= \widehat{v}_{s,t}(2) \\ b_{2^{n-1}+s,t} &= \widehat{v}_{s,t}(3). \end{aligned} \tag{5.3}$$

This fills out the three “outer” $2^{n-1} \times 2^{n-1}$ blocks U_1 , U_2 and U_3 of the quadtree matrix $(b_{p,q})$, and it remains to compute the upper lefthand $2^{n-1} \times 2^{n-1}$ block, U_0 . Proceed recursively by replacing the degree 2^n square matrix $(a_{p,q})$ in the preceding paragraph by the degree 2^{n-1} square matrix whose s, t entry is the 0^{th} DFT coefficient, $\widehat{v}_{s,t}(0)$, of the corresponding 2×2 submatrix block. This step then repeats for this smaller matrix, placing the resulting Fourier coefficients in the three “outer” $2^{n-2} \times 2^{n-2}$ blocks of the original upper lefthand U_0 .

This recursion continues, at the j^{th} stage filling the three outer $2^{n-j} \times 2^{n-j}$ blocks in a nested grid for $Q(f)$, but leaving empty the upper left $2^{n-j} \times 2^{n-j}$ block. At the last i.e., $(n-1)^{\text{st}}$ stage the 2×2 vacant block is filled with the four 4-point DFT coefficients computed from the 2×2 U_0 -matrix passed down from the preceding step. Note that the 1, 1 coefficient of $Q(f)$, $b_{1,1}$, represents the sum of all values of f . A normalized spectrum may also be determined by multiplying by a factor of $1/4$ at each stage of computation of the DFT. However, as we observe in [50], in order to retain the proper scaling of the multiresolution components to compute convolutions using $Q(f)$, there should *not* be a normalization. Note also that if all entries of f are real, then at each level the block matrices U_1 and U_3 —which are positioned symmetrically about the diagonal of $Q(f)$ —are complex conjugates of each other, and U_2 —which is along the main diagonal—is a real matrix.

The inverse transform may likewise be computed via a direct algorithm involving the same number of steps, based on the inverse 4-point DFT. Finally, note that different factorizations of a fixed N as $N = kn$ give different quadtree spectral decompositions of the same $2^N \times 2^N$ image. In the context of computing WPC convolution, we explore some of these different decompositions in [50].

5.2 Spectral invariance of $Z(1, 9)$ acting on 512×512 images

In this section we explore the effects of group operations on the vector spaces $L(k, n)$ and the resulting effect on a discrete signal f and on its spectrum. With the group-invariant property, we have group-invariance of the spectrum with respect to group operations. For a two-dimensional discrete image f we see that group operations can correspond to standard geometric operations on images. The WPC magnitude spectrum is shown to isolate lines and edges, and the phase is shown to provide additional information about image features. Comparisons are made between the wreath product transform (WPT) and both the one- and two-dimensional Haar transforms.

We first focus on specific examples—in particular, building on the 512×512 examples developed earlier—to explore WPC multiresolutions. In [50] we pursue these examples further to test the efficacy of WPC convolution and correlation for some classical problems in image processing.

With reference to Figure 7, we consider the WPC group $Z(1, 9)$ acting on the vector space $L(1, 9)$ of tree-structured 512×512 images. The WPC multiresolution of signals is visualized in this context as follows:

Let V_0, V_1, \dots, V_9 denote the group-invariant subspaces of $L(1, 9)$ associated to levels 0 through 9 of the tree respectively, as described in Theorem 4.4. Recall that V_j is the subspace of all images that are constant on each $2^{9-j} \times 2^{9-j}$ subimage in the collection of nested quadrants. Equivalently, V_j is the space of all functions on the leaves of the tree that take the same value on all leaves which lie below some common node at level j . In this level structure on $L(1, 9)$ the highest level space (in the tree-theoretic sense), V_9 , corresponds to the finest detail and the lowest level space V_0 to the coarsest.

The group-invariant subspace V_j has dimension 4^j since it may be identified as the vector space associated to the tree $Q(1, j)$, obtained from $Q(1, 9)$ via truncation at the j^{th} level; in this way elements of V_j may be associated to $2^j \times 2^j$ images. The projection of an image f onto the space V_j is then visualized as the $2^j \times 2^j$ image obtained from f by collapsing each block Q_i^j to a single pixel whose value is the sum of the values of f on all pixels within Q_i^j —this is the non-normalized Radon map from $L(1, 9)$ to V_j . Thus the map of f onto V_0 gives the sum of all values of f .

For $j \geq 1$, each V_j decomposes further as the direct sum of V_{j-1} and three irreducible, group-invariant subspaces of dimension 4^{j-1} , labeled $W_{j-1,1}$, $W_{j-1,2}$ and $W_{j-1,3}$. The quadtree spectral algorithm shows that the projection of an image f in V_j onto these irreducible subspaces is computed via 4-point DFTs. Note that V_0 is the one-dimensional space of constant functions at level zero, i.e., the set of images of totally constant color or intensity. Thus, for example, the subspaces $W_{1,1}$, $W_{1,2}$ and $W_{1,3}$ at level 2 are all of dimension 4, and lie within the 16-dimensional space V_2 :

$$V_2 = V_1 \oplus W_{1,1} \oplus W_{1,2} \oplus W_{1,3}. \quad (5.4)$$

Note that $V_1 = V_0 \oplus W_{0,1} \oplus \dots \oplus W_{0,3}$ is the 4-point DFT decomposition of a space of signals of length 4. For a real-valued signal f , its Fourier coefficients in subspaces $W_{i,1}$ and $W_{i,3}$ are complex conjugates while those in $W_{i,2}$ are real.

As an example of a multiresolution WPC spectrum, consider the 512×512 image depicted in Figure 11(a). For illustration purposes, and also for later use, we use an image containing vertical, diagonal, and circular edges at different intensities. Its (complex) quadtree spectrum is shown in Figures 11(b) and (c). The finest resolution spectrum occurring at level 9, defined by projections on subspaces $W_{8,1}$, $W_{8,2}$ and $W_{8,3}$, is shown in quadrants Q_1^1 , Q_2^1 , and Q_3^1 respectively of the figure. The image in quadrant Q_0^1 —which contains the lowpass coefficients or DC terms of all the 2×2 subblocks—is further resolved into detail spectra at the other 8 levels, culminating in the projection at level 0. As will be observed later, spectra in the three quadrants at various levels correspond to filtering the image with three complex directional filters.

5.2.1 The WPC group-invariance property

Invariance properties, particularly one like the invariance of the output to shifts in the input, are highly desirable for many important signal processing applications such as detection or estimation of signals with unknown arrival times. In orthogonal and biorthogonal discrete wavelet transforms, the time (space) varying operations preclude the time (space) invariance property. Unless the shifts are proportional to the subsampling rates, spectral coefficients may vary throughout the subspaces. The conventional way of addressing the problem has been to

remove the time (space) varying component after replacing the filters with equivalent filters. However, this results in redundancy and a corresponding loss of orthogonality. Other solutions to the translation-invariance problem have been suggested by many authors, a brief summary of which is presented by Del Marco [17]. Earlier work by Simoncelli, *et al.*, [61], that also entailed redundancy, suggested a weaker form of translation invariance in the sense that all information within a subband remained within that subband as the input is translated.

Invariance in the WPT exists, mainly due the nature of the underlying wreath product group. Invariance is in the sense of group operations, and a signal subject to a group operation τ results in a scaling of the spectrum at the point of application and a corresponding translation of a portion of the spectrum. Thus the spectrum scales and shifts as the image shifts. This general group-invariance property for the WPT includes some translation and rotation invariance, since some WPC group operations do correspond to translation and rotation operations. In addition, critical sampling and orthogonality are maintained in the WPT.

The effect of group elements acting on a $2^9 \times 2^9$ image f and the resulting spectral invariance can be seen by evaluating the effect on the basis vectors of the generating elements described in (3.20). First note that if $\alpha_i^{(j)}$ represents an arbitrary member of this family acting on the i^{th} node at level j , then $\alpha_i^{(j)}$ affects only the irreducible subspaces at level j and higher (in the tree's indexing scheme). Thus, for example, with reference to Figure 7, consider the group element $\alpha_i^{(6)} \in Z(1,9)$ acting on the i^{th} node at the 6th level of the tree; then $\alpha_i^{(6)}$ affects only the subspace W_6 and all higher resolution subspaces W_7 , W_8 and W_9 , but no others.

We determine the effect of group operations as follows: As defined earlier we have $W = e^{2\pi j/4}$ with $j = \sqrt{-1}$ as the 4th root of unity. Starting with the highest resolution level 9, it is easily seen that $(\alpha_i^{(8)})^k$ scales the i^{th} basis vector in each of the subspaces $W_{8,1}$, $W_{8,2}$ and $W_{8,3}$ by $(W^1)^k$, $(W^2)^k$, $(W^3)^k$ respectively, $i = 0, 1, \dots, 4^8 - 1$, $k = 1, 2, 3$ leaving all other basis vectors unaffected. The corresponding i^{th} spectral coefficient at this level is scaled accordingly. At the next coarser resolution, level 8, $(\alpha_i^{(7)})^k$ scales the i^{th} basis vectors in each of the subspaces $W_{7,1}$, $W_{7,2}$, and $W_{7,3}$ by $(W^1)^k$, $(W^2)^k$, and $(W^3)^k$ respectively, while cyclically permuting only the i^{th} set of four basis vectors within each of the subspaces $W_{8,1}$, $W_{8,2}$ and $W_{8,3}$, $i = 0, 1, \dots, 4^7 - 1$ and $k = 1, 2, 3$. Corresponding spectral coefficients are scaled and translated accordingly. Similarly, at the next level $(\alpha_i^{(6)})^k$ scales the i^{th} basis vector in each of the subspaces $W_{6,1}$, $W_{6,2}$, $W_{6,3}$ by $(W^1)^k$, $(W^2)^k$, $(W^3)^k$ respectively, $i = 0, 1, \dots, 4^6 - 1$ and $k = 1, 2, 3$. It cyclically permutes the i^{th} set of four basis vectors within $W_{7,1}$, $W_{7,2}$, $W_{7,3}$ and the i^{th} set consisting of four subsets of four basis vectors each in $W_{8,1}$, $W_{8,2}$ and $W_{8,3}$. All other spaces are left unaffected.

Actions of other group elements on the subspaces follow in like fashion and may be determined accordingly. We have scaling at the level of application and in the corresponding subspace, and translation of coefficients within the subspaces at all other higher resolution subspaces. The sizes of the blocks affected, that is the translated blocks, increase linearly in powers of 4, the increase being proportional to the distance from the level of application. Hence, under the action of group elements, basis vectors within subspaces are scaled and translated but remain within the subspace, and spectral coefficients are adjusted accordingly. For the $Q(1,9)$ tree, the effect on the subspaces due to group operations at various levels is summarized in Table 1. We observe

here that scaling by powers of fourth roots of unity occurs only at the level of application while circular shifts occur at all higher levels.

As a simple illustration of the progressive effect of applying cyclic group elements on subspaces, consider the tree $Q(1, 4)$ and group elements

$$\alpha_0^{(0)}, \alpha_0^{(1)}, \alpha_0^{(2)}, \alpha_0^{(3)} \quad (5.5)$$

in $Z(1, 4)$ acting on the left-most subtrees at levels 0, 1, 2, 3 of the four-level tree, shown in Figure 10. We associate subspaces V_1, V_2, V_3, V_4 with levels 1, 2, 3, and 4 as before, where V_0 is the 1-dimensional subspace at level 0. The effects of the group elements are shown in Figures 12(a)–(d). In this figure, S_1, S_2, S_3 represent scaling by W, W^2, W^3 respectively, while the arrows indicate rotation of the sub-blocks. Starting with $\alpha_0^{(3)}$ at the finest level and continuing to $\alpha_0^{(0)}$ at the coarsest, we observe a progression from scaling to scaling and rotation within the subspaces, the size and number of the rotation blocks increasing the further the subspace resides from the point of application.

It should be noted that the non-expansive nature of the WPT implies a non-expansive progression of error due to local perturbations in the signal. Hence, edge effects and other localized features have very specific local influence. For the $Q(1, 9)$ tree, it is clear that changing one pixel value affects one spectral value in each of the three subspaces at all 9 levels and the DC term (at the 0th level), for a total of 28 pixel values. Furthermore, the locations of all affected pixels are known exactly: For example, referring to Figure 8 and writing all quadrant subscripts in base four, any changes in the original image in quadrant Q_{2333}^5 —which is the 16×16 block located in the lower lefthand corner of the quadrant $Q_{\frac{1}{2}}^1$ of the image—effect changes in the lower lefthand corners of each spectral component, i.e., the following blocks in the quadtree spectrum: for $i = 1, 2, 3$,

the 8×8 blocks Q_{i2333}^6 ,

the 4×4 blocks Q_{i233}^7 ,

the 2×2 blocks Q_{i23}^8 ,

the 1×1 blocks Q_i^9 , and the upper left corner Q_0^9 .

By comparison, in the two-dimensional discrete wavelet transform with a dyadic tree structure, a change in one pixel value of the image results in a change in the spectrum, of blocks of increasing size until a saturation is reached or until the block size is limited by downsampling. The size of the altered blocks depends on the filter size and the size of the perturbations in the image. For example, for all filters of size $2N$, and assuming that even samples are retained in the downsampling process, a one pixel change induces a change in the blocks of size $S_{i+1} \times S_{i+1}$. Letting $\lfloor \cdot \rfloor$ denote the greatest integer lower than its argument, we have

$$S_{i+1} = \lfloor \frac{S_i + 2N - 1}{2} \rfloor \quad i = 0, 1, 2, \dots, l \quad (5.6)$$

up to some level l , after which downsampling reduces the size of the affected blocks. Therefore, assuming $N = 8$, and a 9-level decomposition of images in $L(1, 9)$, altering one pixel value in a 512×512 image affects blocks of size 8, 11, 13, 14, 14, 8, 4, 2 and 1 in all three subspaces for a maximum number of 2515 pixel value changes. Note that retaining odd samples in downsampling rounds up the integers.

5.2.2 Effect on spectra of some geometric operations on $2^9 \times 2^9$ images

As an illustration of spectral invariance due to group operations that result in some standard geometric operations on subimages, Figure 13(a) shows a 512×512 gray scale image subject to two group operations: translation and rotation, as depicted in Figures 13(b) and (c) respectively. The spectra (amplitude component) of the image and transformed images are shown in Figures 13(d),(e) and (f). As indicated earlier, there is no need for pruning after each level of the decomposition. Hence we observe the non-expansive nature of the transform by noting that the image, which is confined to one-quarter of the area, results in a spectra also confined to one-quarter of the area in all the irreducible subspaces.

Referring to the corresponding tree $Q(1, 9)$ of Figure 7, the translated image results from the second-order cyclic group element $(\alpha_0^{(0)})^2$ applied at level 0. The spectrum of the translated image is then exactly the same as that of the original, except that at all subspaces other than V_0 and W_0 , it is translated (circularly shifted) within each subspace precisely as the image is translated (circularly shifted). At the lowest resolution level, the spectra in $W_{0,1}$, $W_{0,2}$, $W_{0,3}$ are scaled by $(W^1)^2$, $(W^2)^2$, $(W^2)^3$ while that at V_0 , being a multiple of the average is unchanged. Hence, as defined in Table 1, effects due to $(\alpha_0^{(0)})^2$ are a second-order scaling at the W_0 subspace level and second-order circular shifts in all other subspaces at levels above it.

Spectral changes are thus confined to precisely one-quarter of the spectra in each of the subspaces, with the remainder of the spectrum left unchanged. This may be verified again by recalling that in each of the subspaces, basis vectors are non-overlapping and orthogonal, with their translates covering the entire length of the signal. Therefore, transformations on specific portions of a signal affect only the corresponding portions of the spectra.

The 270° rotated image is due to the third-power rotation $(\alpha_0^{(1)})^3$ applied at the left-most node at level 1 followed successively by corresponding rotations of 270° at all nodes below it. In other words, this rotation is effected by applying to the image first $(\alpha_0^{(1)})^3$, then $(\alpha_0^{(2)})^3, \dots, (\alpha_3^{(2)})^3$, then $(\alpha_0^{(3)})^3, \dots, (\alpha_{15}^{(3)})^3$, and the like, until at the last stage the 2×2 blocks within the quadrant Q_0^1 are all rotated by $(\alpha_0^{(9)})^3, \dots, (\alpha_{4^9-1}^{(9)})^3$. Denote this group operation by τ . As per Table 1, τ results in both scaling and in cyclic shifts of the spectrum.

We now examine the specific effect of the various elements comprising the product τ . Since τ does not contain a group operation at level 0, there is no change to the spectra in subspaces $V_0, W_{0,1}, W_{0,2}, W_{0,3}$. To see the effect of $(\alpha_0^{(1)})^3$, recall that $\alpha_0^{(1)}$ generates scaling of the basis vectors in the W_1 subspaces and cyclic shifts of vectors in all the other higher resolution spaces. Specifically, this occurs as follows: In each of the 4-dimensional subspaces $W_{1,1}, W_{1,2}, W_{1,3}$, $\alpha_0^{(1)}$ affects only the first one-quarter of the support of each of the four basis vectors. However, since only the first basis vector in each of the subspaces is nonzero in that region, only that one is affected. Consequently only the first coefficient in the spectrum corresponding to the subspaces $W_{1,1}, W_{1,2}, W_{1,3}$ is scaled, with scaling by $(W^1)^3$, $(W^2)^3$, and $(W^3)^3$ respectively. In subspaces $W_{2,1}, W_{2,2}$, and $W_{2,3}$, each of dimension 16, only one-quarter, (or in this case the first four basis vectors in each of the subspaces) are affected by a cyclic rotation. Similarly, in each of the remaining subspaces $W_{3,1}, W_{3,2}, W_{3,3}, W_{4,1}, \dots, W_{8,3}$, one-quarter of the basis vectors are cyclically rotated in blocks of 4^k , $k = 1, 2, \dots, 6$. Corresponding coefficients in the spectrum in blocks of size $2^k \times 2^k$ are moved accordingly. Hence, the net effect of $(\alpha_0^{(1)})^3$ is that the

1-dimensional subspace W_0 is unaffected, subspace W_1 is scaled, while spectra in the remaining subspaces are rotated locally in blocks of $2^k \times 2^k$ for $k = 0, 1, 2, \dots, 6$.

The effect of group operations $(\alpha_0^{(2)})^3, (\alpha_1^{(2)})^3, \dots, (\alpha_{15}^{(2)})^3$ at level 2, is not unlike that of τ at level 1, except subspaces starting from the higher resolution spaces W_2 onwards are affected, while the lower resolution spaces remain unchanged. Accordingly, one-quarter of all vectors in the irreducible subspaces $W_{2,1}, W_{2,2},$ and $W_{2,3}$ are scaled by $(W^1)^3, (W^2)^3,$ and $(W^3)^3$ respectively while vectors in the higher resolution subspaces are rotated locally in blocks of size $2^k \times 2^k$, for $k = 0, 1, 2, \dots, 5$. The effect of other group operations $(\alpha_i^{(3)})^3, (\alpha_i^{(4)})^3, \dots, (\alpha_i^{(8)})^3$ may be interpreted in a similar fashion. Accordingly, the cumulative effect of the group operation τ on the spectrum is a local 270° rotation of one-quarter of the spectrum in all subspaces with dimension higher than four with scaling by $(W^i)^3$ in subspaces of dimension higher than one. Thus we conclude that τ effects cyclic shifts and scaling by $(W^1)^3, (W^2)^3, (W^3)^3,$ at all the eight higher resolution spaces W_2 through W_9 in each of the three irreducible subspaces at these levels.

Comparing the spectra of the original and rotated images, we observe that the spectrum in each of the subspaces (except subspaces V_0 and W_1) of the rotated image is a 90° counter-clockwise rotation of the corresponding spectrum of the original image, with all irreducible subspaces (except V_0) indexed by 1, 2, 3 scaled by W^1, W^2 and W^3 respectively.

5.3 Relationship to the Haar transforms

We have already noted that the one-dimensional Haar transform of a signal of length 2^N is the same as the WPT obtained from the binary tree with N levels i.e., where all the cyclic groups in the wreath product construction have order 2 (cf., Theorem 4.12). In this section we elucidate the relation of both the one- and two-dimensional Haar transforms to the WPT obtained from trees with higher order branching.

The WPT may be seen as a generalization of the one-dimensional Haar transform with entries containing roots of unity. As an illustration and by way of comparison with the $Z(1,2)$ group transform on signals of length 16 given explicitly in Section 4.3, consider the representation of a input signal \mathbf{f} of length 16 in terms of the Haar transform coefficients \mathbf{u} . Thus, $\mathbf{f} = \mathbf{H}\mathbf{u}$ where

$$\mathbf{H} = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & -1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & -1 & 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 1 & -1 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & -1 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 1 & -1 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & -1 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}$$

and as in Section 4.3 or display (4.21), the columns of the matrix \mathbf{H} are basis functions for the irreducible subspaces of the group acting on the vector space. We note the difference between the WPT and Haar transform bases: While the former involves scaled and shifted versions of a single prototype function (a “mother wavelet”), the WPT operates in a similar fashion on three prototype functions (whose group translates respectively span the

three irreducible subspaces at the finest detail level of the decomposition). We see that at the lowest resolution the Haar transform takes differences of sums of samples; as the resolution increases the differences become more local. The WPT, while generating a complex spectrum, also generates differences of sums which become more local as the resolution increases; but the differences are not for adjacent samples but next to adjacent samples.

A clearer and more interesting view emerges when the transforms are applied to two-dimensional signals. For a $2^N \times 2^N$ image the two-dimensional Haar transform operates on 2×2 subblocks, as does the WPT when applied to the $Q(1, N)$ quadtree scanned image. Accordingly, consider a prototype 2×2 subblock consisting of four elements x_0, x_1, x_2, x_3 scanned in clockwise order. Figure 14 portrays a 2×2 grid where the ellipses indicate summation and the arrows indicate differencing. The two-dimensional Haar transform operates on a function f defined on the 2×2 grid and generates sums and differences that may be considered as measuring averages as well as horizontal, vertical and diagonal edges:

$$\begin{aligned}
 w(0) &= f(x_0) + f(x_1) + f(x_2) + f(x_3) \\
 w(1) &= \{f(x_0) + f(x_1)\} - \{f(x_2) + f(x_3)\} \\
 w(2) &= \{f(x_0) + f(x_3)\} - \{f(x_1) + f(x_2)\} \\
 w(3) &= \{f(x_0) + f(x_2)\} - \{f(x_1) + f(x_3)\}.
 \end{aligned} \tag{5.7}$$

Likewise, the WPT generates its spectrum according to the following rules:

$$\begin{aligned}
 v(0) &= f(x_0) + f(x_1) + f(x_2) + f(x_3) \\
 v(1) &= \{f(x_0) - f(x_2)\} - W\{f(x_1) - f(x_3)\} \\
 v(2) &= \{f(x_0) + f(x_2)\} - \{f(x_1) + f(x_3)\} \\
 v(3) &= \{f(x_0) - f(x_2)\} + W\{f(x_1) - f(x_3)\}
 \end{aligned} \tag{5.8}$$

where $W = e^{2\pi j/4}$ and $j = \sqrt{-1}$. At this level of detail, the difference between the WPT and Haar transforms can be explained by the fact that these arise as Fourier transforms for the actions of two different groups on the space of signals of length 4. The WPT comes from the action of the cyclic group Z_4 acting by rotation (modulo 4) and the corresponding Fourier transform is simply the 4-point DFT. The Haar transform is the Fourier transform for the action of the non-cyclic (but still abelian) group $Z_2 \times Z_2$ acting by independent vertical and horizontal shifts (modulo 2) as described in Section 3.1, Example 2. This difference between the two transforms propagates to the coarser levels of resolution.

We observe that at the first level of the Haar transform decomposition, (cf., Figure 14 and equation 5.7) the spectrum may be defined as consisting of “average” horizontal, vertical and diagonal edges where the average is over 2 pixels. Also noted is that the average $w(0)$ and diagonal edges represented by $w(3)$ are the same as $v(0)$ and $v(3)$ in the WPT. Furthermore $w(1)$ and $w(2)$ may be derived from $v(3)$ through the sum and difference of its real and imaginary parts. Accordingly, phase information present in the quadtree scanned WPT is implicitly available in the Haar transform.

Variables $w(1)$ and $w(3)$, (and similarly, $v(1)$ and its complex conjugate $v(3)$), are interpreted as approximations to various edge gradient functions. Consider the gradient of a continuous function $f(x, y)$ along polar

coordinates r in a direction θ . This is defined in [55] as

$$\frac{\partial f}{\partial r} = \frac{\partial f(x,y)}{\partial x} \cos(\theta) + \frac{\partial f(x,y)}{\partial y} \sin(\theta) = f_x \cos(\theta) + f_y \sin(\theta).$$

Here, the partial derivatives f_x and f_y are defined along the x - and y -axes, which in the discrete case are evaluated in terms of row edge gradients $g_r(j, k)$ and column edge gradients $g_c(j, k)$ where j, k denote row and column indices. Accordingly, the amplitude of the spatial gradient $g(j, k)$ is given by

$$|g(j, k)| = ([g_r(j, k)]^2 + [g_c(j, k)]^2)^{1/2}$$

while the orientation of the spatial gradient with respect to the row axis, which corresponds to the direction of the edge is

$$\theta(j, k) = \tan^{-1} \frac{g_c(j, k)}{g_r(j, k)}.$$

In the Haar transform, $w(1)$ and $w(2)$ may be considered approximations to the vertical and horizontal edge gradients. In contrast, variables $v(1)$ and $v(3)$ in (5.8) can be seen as

$$v(j, k) = g_1(j, k) \mp W g_2(j, k)$$

where $W = e^{2\pi j/4}$, $j = \sqrt{-1}$ and $g_1(j, k)$ and $g_2(j, k)$ are identified as measuring cross differences along orthogonal directions that are rotated $\pi/4$ radians from the x - and y -axes. Hence these are seen to be diagonal edge gradients with an amplitude

$$|v(j, k)| = ([g_1(j, k)]^2 + [g_2(j, k)]^2)^{1/2}$$

and an edge orientation

$$\theta(j, k) = \pi/4 + \tan^{-1} \frac{g_2(j, k)}{g_1(j, k)}$$

with respect to the x -axis.

The multiresolution spectra of the two transforms shown earlier may be compared. We note again that for the $Q(1, N)$ quadtree scan, one transform can be derived from the other and consequently they are inherently equivalent. The WPT however—based on a 4-point DFT—permits an explicit frequency interpretation and consequent analysis in terms of both amplitude and phase. In Figure 11, the WPT is shown in terms of its amplitude and phase. The gray scale intensity image for the phase shows the phase in the range $[-\pi, \pi]$ with $-\pi$ as black, π as white, and values between as intermediate shades of gray. The Haar transform is shown in absolute value. Since both transforms compute local spatial derivatives, they may be viewed in terms of their capabilities for line and edge detection, where a line is defined as one pixel thick and an edge is defined as a transition from a dark to a light region or conversely. The magnitudes of the derivatives reflect the line or edge intensities.

Alternatively, both transforms may each be viewed as sets of four two-dimensional filters consisting of a lowpass LP and three directional filters. For the Haar, we define HP_x , HP_y and HP_d as highpass filters in

the x , y and diagonal directions respectively, with the last of these generating directional filters sensitive to orientations at 45° and 135° in the image plane. The WPT directional filters consist of HP_d —same as the Haar—and two complex directional filters HP_{xy} and HP_{xy}^* which are highpass simultaneously in both the x and y directions. Figure 15 shows the frequency response of the two-dimensional filters.

The two-dimensional 9-level Haar transform is shown in Figure 11(d). The Haar transform detects local horizontal, vertical and diagonal lines and edges through a measurement of differences. Note that since the spectrum is displayed using absolute values, we observe in all quadrants of the spectrum edge occurrences that transition from both light to dark regions (positive slope) and conversely. In quadrants 2 and 4 of the spectrum we observe, respectively, the presence of horizontal and vertical edges in the image. Diagonal lines and edges are exhibited in quadrant 3. Here we observe the presence of both the light and dark diagonal lines in the image, and also the lower intensity diagonal edges of the ellipse. The WPT amplitude and phase are shown in Figures 11(b) and (c). We first note that since both the Haar and the WPT measure diagonal edges similarly, quadrant 3 of the WPT amplitude spectrum is identical to that of the Haar, identifying diagonal lines and edges, as described earlier. Quadrants 2 and 4 in the amplitude spectrum are identical since they measure complex conjugate quantities $v(1)$ and $v(3)$. The associated complex filter, HP_{xy} , responds to both horizontal and vertical edges (but not diagonal ones), and we therefore observe both such edges. In effect, this reflects the combined result of Haar filtering in both quadrants 2 and 4. Accordingly, the ellipse’s edges are seen almost in their entirety. Since this is an amplitude spectrum, there is no distinction at this point between edges that are transitions from one intensity region to another. However, gray scale levels indicate edge intensities.

In the phase spectrum, we note that quadrant 3 reflects the phase of $v(2)$, which is real; however, since negative numbers are displayed as π radians, we note a nonzero contribution in that quadrant which essentially reflects the difference in intensity, as measured by $v(1)$, of the two local diagonal edges. Recalling that the amplitude spectrum in quadrants 2 and 4 reflects the presence and intensity of horizontal and vertical edges, the phase spectrum in these two quadrants indicates the type of edge in terms of its transition from high (low) to low (high) intensity. Hence, referring again to the phase of $v(1)$, horizontal edges that transition from a low to a high intensity region are all portrayed equally with the same phase intensity (almost white), while those that transition the other way are all shown with a low intensity (almost black). Similarly, vertical edges from a low to high intensity appear almost white, while those occurring in the other direction are almost black. Quadrant 4 displays the complex conjugate phase of quadrant 2. We note further that the WPT complex spectrum amplitude $v(2)$ is the same as the approximation to the spatial gradient amplitude represented by the Roberts edge detector [54]. Other work with WPC groups and edge detection may be found in [14].

6 GENERALIZED MULTIREOLUTION ANALYSIS

As remarked earlier, in any situation in which invariant decompositions have some natural hierarchy, a multiresolution interpretation is possible, and hopefully, useful. The following provides three settings in which this approach seems to make sense.

6.1 Combinatorial Radon transforms and clumping

Let X be a set, and Y any set of subsets of X . Then the (*non-normalized*) *Radon transform* is a linear map \mathcal{R} from $L(X)$ to $L(Y)$ defined by

$$\mathcal{R}f(y) = \sum_{x \in y} f(x). \quad (6.1)$$

It is natural to think of \mathcal{R} as a “clumping” operator. Under the natural inner product on $L(X)$ and $L(Y)$, the adjoint of \mathcal{R} is the map \mathcal{S} defined by

$$\mathcal{S}h(x) = \sum_{y \ni x} h(y). \quad (6.2)$$

The image of \mathcal{S} is the set of functions constant on the subsets of Y (cf., [11] for a very nice introduction to this setting). If \mathcal{S} is injective, then we think of the image of \mathcal{S} as providing the “coarse scale” information for $L(X)$ and the orthogonal complement as providing the detail. Given a sequence consisting in X , subsets of X , subsets of subsets of X , and so on, then under analogously appropriate conditions a multiresolution of $L(X)$ is obtained. Notice that if, in addition, we have some sort of group action on X compatible with the indicated subsets (and subsets of subsets, etc.), then this multiresolution decomposition is G -invariant as well.

The sets $X_{\mathbf{m}}$ of course provide one such example, with the internal nodes indicating the sequence of (disjoint) sets of subsets. For a different example consider the following scenario. Let G act on X . Then we get an action of G on X_k , the set of all k -sets of X . Notice that here the subsets may have overlaps. This group action need not be transitive, even if the action on X is.

For any $0 \leq j < k \leq n/2$ we have a Radon transform multiresolution analysis via the Radon transform maps

$$R_k^j : L(X_j) \longrightarrow L(X_k)$$

defined by

$$(R_k^j f)(y) = \sum_{x \subset y} f(x)$$

and the adjoint of R_k^j is S_k given by

$$(S_k^j F)(x) = \sum_{y \supset x} F(y)$$

These maps are G -equivariant and represent a natural refinement of $L(X)$. Of particular interest is the action of the symmetric group S_n on 1-sets, 2-sets, ..., k -sets of $\{1, \dots, n\}$ ($k \leq n/2$). In this case there is a nesting of the permutation representations,

$$L(X) = L(X_1) \hookrightarrow L(X_2) \hookrightarrow \dots \hookrightarrow L(X_k)$$

such that

$$L(X_j) = L(X_{j-1}) \oplus V_j$$

and V_j is irreducible. Lottery data can be (and has been) studied from this point of view (cf., [19]). In this situation, $f(x)$ will be the number of contestants choosing a particular k -set of numbers for their “pick.” The multiresolution described above is the natural consideration of the pure effect of various individual numbers, pairs of numbers, etc. on the overall patterns of choice. This sort of analysis can also be applied in the setting of experimental designs. Here a standard analysis of variance approach will consider the projection of data onto a “blocks-ignoring-treatments” space. The decomposition described above is one way of looking for a finer stratification of this contribution to the overall sum-of-squares [3].

6.2 Induced representations

The wreath product filtration of the spatial domain $L(X)$ in Theorem 4.4 is a special case of decompositions that result from chains of subgroups of a group. We describe the general situation, noting that it, in turn, is a special case of multiresolutions obtained from clumping (defined above).

Section 3.1 shows that each transitive action of a finite group G on a set X is equivalent to the action of G by left translation on the set G/H of left cosets of H in G , where H is the stabilizer of a point in X (and X may then be identified with G/H). Conversely, one may begin with any subgroup H of G and obtain a transitive permutation representation of G on the set G/H in which H is the stabilizer of the coset $1H$. The linear representation of G on $L(X) = L(G/H)$ obtained from this permutation action is the trivial (identity representation of degree 1) representation of H induced to G , namely $Ind_H^G(I_H)$ (cf., Section 3.3 in [59]). Decomposing $L(X)$ into G -invariant components to obtain a multiresolution is thus equivalent to decomposing this induced representation.

When H is contained in a larger subgroup K there is a natural clumping of the coset space $X = G/H$ by the set Y of left cosets of K : Here two cosets $\alpha_1 H$ and $\alpha_2 H$ belong to the same element of Y precisely when they are contained in the same coset αK of K . Equivalently, the space $L(G/K)$ of functions constant on the left cosets of K is a G -invariant subspace of $L(G/H)$. The Radon map from $L(G/H)$ to $L(G/K)$ sends any function constant on the cosets of H to the function whose value on a coset αK is the sum of all its values on the left cosets of H contained in αK . (In representation-theoretic terms, the Frobenius Reciprocity Theorem guarantees that $Ind_K^G(I_K)$ is a constituent of $Ind_H^G(I_H)$.) Thus there is a G -invariant decomposition:

$$L(X) = L(G/H) = L(G/K) \oplus W,$$

where the space W need not be irreducible (as was the case in Theorem 4.8).

More generally, any chain of subgroups $H = K_m \leq K_{m-1} \leq \cdots \leq K_0 = G$ gives a G -invariant filtration of $L(G/H)$: Let $V_i = L(G/K_i)$ for $0 \leq i \leq m$. Then

$$0 \subseteq V_0 \subseteq V_1 \subseteq \cdots \subseteq V_{m-1} \subseteq V_m = L(G/H)$$

where V_0 is the one-dimensional space of constant functions in $L(G/H)$. Also, $V_i = V_{i+1} \oplus W_i$ for some G -invariant subspace W_i . If each K_i is a proper subgroup of K_{i-1} then the corresponding subspace containments are also proper (by counting dimensions). Thus any subgroup chain containing H gives a multiresolution of $L(G/H)$.

For arbitrary finite groups this rationale reduces the problem of finding “group invariant multiresolutions” to treating cases where the stabilizer subgroups are *maximal* in their overgroups (in which case the permutation action is called *primitive*), and then decomposing the complementary spaces W into irreducible constituents. For wreath products of cyclic groups this accomplished in Theorem 4.8.

6.3 Finite approximants to profinite groups

Lastly, we note that by taking inverse limits of finite groups it is possible to extend some of our analysis to the setting of profinite groups acting on profinite sets (see e.g. [6, 27]) or on infinite dimensional topological vector spaces (see e.g. [53]). This approach may lead to appropriate “continuous” signal spaces admitting group actions and multiresolutions.

A sequence of finite groups acting on rooted trees effectively gives the action of an automorphism group of an infinite SHRT on the infinite paths, or *ends* of the tree. In this setting, there is a natural notion of “bandlimited function,” defined as any function on the ends whose value at a given end only depends on a finite subtree which the end extends. Our analysis can be viewed as deriving a bandlimited approximation to a function on the ends of an infinite SHRT.

In general, an inverse limit of finite groups may be given the structure of a (profinite) totally disconnected, compact topological group; and admissible representations of such a group form an important class of representations that generalizes the finite dimensional representation theory. Moreover, within this family of representations there are analogues of the group algebra, induced representations, Radon maps (distributions), and the like (see [60]) that lay a foundation on which a useful signal processing theory may be developed.

6.4 Wavelet packets and the best basis algorithm

The construction of Haar wavelets can be written as the successive averaging (lowpassing) and differencing (highpassing) of neighboring data points in an input sequence of size 2^n . Each time we separate out the highpass part, and then make a new (half-sized) data sequence of the lowpassed part. This procedure is then repeated on the lowpass.

Recall (cf., Theorem 4.12) that the above procedure is precisely the decomposition of functions defined on the rooted binary tree of depth n , decomposed into its projections onto irreducible subspaces,

$$L^2(X(n)) = L^2(X_{2,2,2,\dots,2}) = V_0 \oplus V_1 \oplus \dots \oplus V_n.$$

The subspace V_j is of dimension 2^{j-1} (with V_0 of dimension 1) and is the highpass information returned at step $n - j$ in the averaging and differencing.

In this scheme, *wavelet packets* are constructed by additionally, operating on the highpass output by repeating the application of the highpass and lowpass operation to this sequence as well (and recursing). This yields projections onto an overcomplete set of functions, as at each stage there is a preferred choice of basis functions (since one keeps splitting each of the V_j , and computes projections on all the subspaces).

The *Best Basis Algorithm* [13] takes “the standard” bases for all of the subspaces and computes the projections onto each of the basis vectors. From this overcomplete set of data a representation of the original function is then chosen to minimize some cost function on the projections. For example, it may be useful to “split” the projections into some of the V_j .

The connection with our point of view is as follows: Let G_k be the symmetry group for a binary tree of height k , so G_n is the original group. Then we have a tower of groups

$$G_n > G_{n-1} > \cdots > G_1 > G_0 = \{1\}. \quad (6.3)$$

Let G act on X , and hence on $L(X)$, with irreducible decomposition

$$L(X) = V_0 \oplus \cdots \oplus V_k.$$

We may now consider this as a representation space for G_{n-1} as well. Each V_j is G_{n-1} -invariant and so we further decompose these under G_{n-1} and obtain

$$L(X) = V_{0,0} \oplus \cdots \oplus V_{0,n_0} \oplus \cdots \oplus V_{k,0} \oplus \cdots \oplus V_{k,n_k}.$$

Of course we can iterate this through the chain of subgroups (6.3). If each decomposition comes with a preferred basis, then we can perform an analogous best basis algorithm on this group-theoretically motivated decomposition. Notice that in fact, this approach makes sense for any sort of representation, permutation or not. It is conceivable to attempt to do this over a class of subgroup chains.

7 CONCLUSION

In this paper we have shown how group theory, especially the use of permutation actions and their associated representations, gives rise to some of the tools and techniques of modern digital signal processing. Classical DFT and FFT based approaches come from the use of commutative groups. By broadening our horizons to noncommutative groups we expand upon the original work initiated by Holmes, Karpovsky and Trachtenberg, who first brought these ideas to the subject of signal processing. In particular, we have paid special attention to groups obtained as wreath products of cyclic groups and in so doing we are able to rederive and give a group-theoretic context to Haar wavelets, as well as develop a more general class of multiresolution expansions for discrete signals. The multiresolution is compatible with the group action and in this way generalizes the classical spectral analysis approach for analyzing a discrete signal. With specific application to 2-dimensional images in mind, we describe a quadtree scanning scheme for converting a discrete $2^N \times 2^N$ signal into a row vector in such a way that the wreath product group acts on a nested grid of subimages, and we give a quadtree scheme for displaying the group-based multiresolution spectrum of such images.

This paper initiates an analysis which is concluded in Part II of this sequence of papers, [50]. There we show that the additional algebraic structure imposed by our group-theoretic approach provides a natural group-based convolution, and hence a natural collection of (noncommutative) convolution filters. Explicit examples demonstrating convolution on standard test images using the quadtree scanning method developed in this paper are given, and applications to image recognition problems are explored.

There are many possible avenues for future work: The wreath product groups are but one class of non-commutative groups. Other noncommutative groups might yield a similarly rich set of tools for digital signal processing tasks. A “continuous theory” for finite and infinite groups may exist, in a manner analogous to Shannon’s sampling theorem, that would formalize a new theoretical bases for further application. Part II explores some applications of the existing theory of WPC-based signal processing. As discussed in the conclusion there, application of the WPC group-invariant noncommutative filters to other signal processing tasks such as matched filtering, pattern recognition and noise removal are but some of the many standard applications awaiting investigations.

References

- [1] G. Apple and P. Wintz, *Calculation of Fourier transforms on finite Abelian groups*, *IEEE Trans. Inf. Theory*, *IT-16*, March 1970, 233–234.
- [2] L. Auslander and R. Tolimieri, *Is computing with the finite Fourier transform pure or applied mathematics?* *Bull. Amer. Math. Soc. (N.S.)* **1**(6) (1979), 847–897.
- [3] R. Bailey, P. Diaconis, D. Rockmore and C. Rowley, *A spectral analysis approach to data on designs*, in preparation.
- [4] R. Bailey, C. E. Praeger, C. A. Rowley, and T. Speed, *Generalized wreath products of permutation groups*, *Proc. London Math. Soc.* **47**(3) (1983), 69–82.
- [5] K. Balasubramanian, *Graph theoretical characterization of NMR group nuclear spin species and the construction of symmetry adapted NMR spin functions*, *J. Chem. Phys.* **73**(7) (1980), 3321–3337.
- [6] H. Bass, M. V. Otero-Espinar, D. N. Rockmore, and C. Tresser, *Cyclic renormalization and automorphism groups of rooted trees*, *Lecture Notes in Mathematics*, Vol. 1621, Springer-Verlag, Berlin, (1996).
- [7] L. Beckett and P. Diaconis, *Spectral analysis for discrete longitudinal data*, *Adv. Math.* **103** (1994), 107–128.
- [8] R. Bernadini and J. Kovacevic, *Designing local orthogonal bases on groups I: Abelian case*, Preprint (1997).
- [9] R. Bernadini and J. Kovacevic, *Designing local orthogonal bases on groups II: Nonabelian case*, Preprint (1997).
- [10] R.E. Blahut, *Algebraic methods for signal processing and communications coding*, Springer Verlag, NY (1991).
- [11] E. Bolker, *The finite Radon transform*, *Contemporary Math.* **63** (1987), 27–50.
- [12] T. W. Cairns, *On the fast Fourier transform on finite abelian groups* *IEEE Trans. Comp.*, May 1971, 569–571.
- [13] R. R. Coifman and M. V. Wickerhauser, *Wavelets and adapted waveform analysis*, in *Wavelets: Mathematics and Applications*, J. Benedetto and M. Frazier, eds., CRC Press, Boca Raton, FL 1993, pp. 399–423.

- [14] V. Chickanosky and G. Mirchandani, *Wreath products for edge detection*, Proceedings of the 1998 ICASSP, Vol. 5 (1998), 2953–2956.
- [15] G. S. Chirikjian and I. Ebert-Uphoff, *Numerical convolution on the Euclidean group with applications to workspace generation*, Tech Rep. RMS 9-95-10. Dept. of Mechanical Eng., Johns Hopkins University, Sept., (1995).
- [16] J. W. Cooley and J. W. Tukey, *An algorithm for machine calculation of complex Fourier series*, Math. Comp., **19** (1965), 297–301.
- [17] S. Del Marco and J. Weiss, *Improved transient signal detection using a wavepacket-based detector with an extended translation-invariant wavelet transform*, IEEE Trans. on Signal Processing, 45(4), 841–850, (1997).
- [18] P. Diaconis, *A generalization of spectral analysis with applications to ranked data*, Ann. Stat. **17** (1989), 949–979.
- [19] P. Diaconis, *Group representations in probability and statistics*, IMS, Hayward, CA, (1988).
- [20] J. R. Driscoll and D. Healy, *Computing Fourier transforms and convolutions on the 2-sphere*, (extended abstract) Proc. 34th IEEE FOCS, (1989) 344–349; Adv. in Appl. Math., **15** (1994), 202–250.
- [21] J. R. Driscoll, D. Healy and D. Rockmore, *Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs*, SIAM J. Comput. **26**(4) (1997) 1066–1099.
- [22] D. Dummit and R. Foote, *Abstract Algebra*, Second Edition. John Wiley & Sons, New Jersey (1999).
- [23] A. Dutt, M. Gu, and V. Rokhlin, *Fast algorithms for polynomial interpolation, integration, and differentiation*, SIAM J. Numer. Anal. **33**(5) (1996), 1689–1711.
- [24] D. Eberly and D. Wenzel, *Adaptation of group algebras to signal and image processing*, CVGIP: Graphical Models and Image Processing **53**(4) (1991), 340–348.
- [25] I. Ebert-Uphoff and G. S. Chirikjian, *Inverse kinematics of discretely actuated hyper-redundant manipulators using workspace density*, Proceedings 1996 IEEE International Conference on Robotics and Automation, Minneapolis, MN April 1996.
- [26] A. Fässler and E. Stiefel, *Group theoretical methods and their applications*, Birkhäuser, Boston, MA, (1992).
- [27] A. Figà-Talamanca and C. Nebbia, *Harmonic Analysis and Representation Theory for Groups Acting on Homogeneous Trees*, LMS Lecture Note Series 162, Cambridge U. Press (1991).
- [28] Y. Fisher, E.W.Jacobs and R.D. Boss, *Fractal Image Compression using iterated transforms* in Image and Text Compression, ed. James A. Storer, Kluwer Academic Publishers, 35–61.
- [29] K. Flornes, A. Grossman, M. Holschneider, and B. Torrèsani, *Wavelets on discrete fields*, Appl. and Comp. Harmonic Analysis, **1**(2) (1994), 137–146.

- [30] A. Grossman, J. Morlet, and T. Paul, *Transforms associated to square-integrable group representations, I, General results*, J. Math. Phys., **26** (1985), 2473–2479.
- [31] E. J. Hannan, *Group representations and applied probability*, J. Appl. Prob., **2** (1965), 1–68.
- [32] D. Healy, G. Mirchandani, T. Olson and D. Rockmore, *Wreath products for image processing*, Proceedings of the 1996 ICASSP, Vol. 6, 3582–3586.
- [33] D. Healy, S. Moore, and D. Rockmore, *An FFT for the 2-sphere and applications*, Proceedings of the 1996 ICASSP, Vol. 3, 1323–1326.
- [34] D. Healy, P. Kostelec, S. Moore and D. Rockmore, *FFTs for the 2-sphere — improvements and variations*, Advances in Applied Mathematics, to appear.
- [35] R. Holmes, *Signal processing on finite groups*, Technical Report 873, MIT, Lincoln Laboratory, (1990).
- [36] R. Holmes, *Mathematical foundations of signal processing, II*, Technical Report 781, MIT, Lincoln Laboratory, (1987).
- [37] C. Johnston, *On the pseudodilation representations of Flornes, Grossmann, Holschneider, and Torresani*, J. Fourier Anal. Appl. **3**(4) (1997), 377–385.
- [38] T.A.C.M. Kalker and I. A. Shah, *A group theoretic approach to multidimensional filter banks: theory and applications*, IEEE Trans. Signal Processing, Vol. 44, No. 6, June 1996, 1392–1405.
- [39] M. Karpovsky and E. Trachtenberg, *Filtering in a communication channel by Fourier transforms over finite groups*, in Spectral Techniques and Fault Detection, M. Karpovsky (ed.) Academic Press, NY (1985), 179–212.
- [40] A. Kerber. *Representations of Permutations Groups I, II*, Lecture Notes in Mathematics, Vols. 240 and 495, Springer-Verlag, Berlin (1971) and (1975).
- [41] A. B. Kyatkin and G. Chirikjian, *Template matching as a correlation on the discrete motion group*, Department of Mechanical Engineering, Johns Hopkins University, Technical Report RMS-97-3, (1997).
- [42] J. Lafferty and D. Rockmore, *Spectral techniques for expander codes*, Proc. of Symp. on Theory of Computing, (1997), 160–167.
- [43] R. Lenz, *Using representations of the dihedral groups in the design of early vision filters*, Proc. ICASSP '93, Vol. 5, 165–168.
- [44] R. Lenz, *Group theoretical methods in image processing*, Springer-Verlag Lecture Notes in Computer Science, Vol. 413, New York, (1987).
- [45] F.J. MacWilliams, *Codes and ideals in group algebras*, in Combinatorial Mathematics and its Applications, R.C. Bose and T.A. Dowling, eds., University of North Carolina Press, Chapel Hill, NC (1969), 317–328.

- [46] S. Mallat, *Multiresolution approximations and wavelet orthonormal bases of $L^2(\mathbf{R})$* , Trans. Amer. Math. Soc. **315**(1) (1989), 69–87.
- [47] D. Maslen, *Fast transforms and sampling for compact groups*, J. Fourier Analysis and its Applications, (to appear).
- [48] D. Maslen and D. Rockmore, *Generalized FFTs*, DIMACS Series in Disc. Math. and Theor. Comp. Sci., Vol. 28, L. Finkelstein and W. Kantor (eds.), (1997), 183–237 .
- [49] D. Maslen and D. Rockmore, *Separation of variables and the efficient computation of Fourier transforms on finite groups, I*, J. of A. M. S, **10**(1) (1997), 169–214.
- [50] G. Mirchandani, R. Foote, D. Rockmore, D. Healy and T. Olson, *A wreath product group approach to signal and image processing: Part II—Convolution, Correlation, and Applications*, submitted to the IEEE Trans. on Signal Processing, May 1999.
- [51] M. Özaydin and T. Przebinda, *Platonic, orthonormal wavelets*, ACHA **4** (1997), 351–365.
- [52] S.-M. Phoong and P. P. Vaidyanathan, *Paraunitary filter banks over finite fields*, IEEE Trans. Sig. Proc. **45**(6) (1997), 1443–1457.
- [53] L. S. Pontrjagin, *Topological Groups*, Gordon and Breach, NY, (1966).
- [54] W. K. Pratt, *Digital Image Processing*, Second Edition, John Wiley & Sons, Inc., New York, (1991), 497–503.
- [55] J. M. S. Prewitt, “Object Enhancement and Extraction,” in *Picture Processing and Psychopictorics*, B. S. Lipkin and A. Rosenfeld (eds.), Academic Press, New York, (1970).
- [56] D. Rockmore, *Applications of generalized FFTs* DIMACS Series in Disc. Math. and Theor. Comp. Sci., Vol. 28, L. Finkelstein and W. Kantor (eds.), (1997), 329–369 .
- [57] H. Samet, *Region representation: quadtrees from boundary codes*, Comm. ACM **23** March 1980, 163–170.
- [58] W. R. Scott, *Group Theory*, Prentice-Hall, New Jersey (1964).
- [59] J. P. Serre, *Linear Representations of Finite Groups*, Springer-Verlag, New York (1977).
- [60] A. J. Silberger, *Introduction to Harmonic Analysis on Reductive P-Adic Groups*, Princeton University Press, NJ (1979).
- [61] E. P. Simoncelli, W. T. Freeman, E. H. Adelson and D. J. Heeger, *Shiftable multiscale transforms*, IEEE Trans. on Information Theory, **38**(2), 587–607, March 1992
- [62] M. Sipser and D. Spielman, *Expander codes*, IEEE Trans. on Inf. Theory, **42**(6) November 1996, 1710–1722.
- [63] A. Srivastava, M. I. Miller, and U. Grenander, *Lie group parameterization for dynamics based prior in ATR*, Sixth Digital Signal Processing Workshop, November 1995

- [64] A. Srivastava, M. I. Miller and Ulf Grenander, *Ergodic algorithms on special Euclidean groups for ATR*, Proceedings MTNS 96, Birkhauser.
- [65] M. Swanson and A. Tewfik, *A binary wavelet decomposition of binary images*, IEEE Trans. Image Proc., **5**(12) (1996), 1637–1650.
- [66] M. Unser, *On the approximation of the discrete Karhunen-Loève transform for stationary processes*, Signal Processing, **7** (1984), no. 3, 231–249.
- [67] R. J. Valenza, *A representation-theoretic approach to the DFT with noncommutative generalizations*, IEEE Trans. Sig. Proc., **40**(4) (1992), 814–822.
- [68] N. Vilenkin, *Special functions and the theory of group representations*, Translations of Mathematical Monographs, **22**, A.M.S., Providence RI, (1968).
- [69] J.B. Weaver, D.M. Healy, Jr., Sumit Chawla, and D.W. Warner, *Alternative norms for image compression and imaging with reduced encodes*, (Abstract) Proceedings of the Society of Magnetic Resonance, Vancouver, Canada, August, 1997, p. 1999.
- [70] L. G. Weiss, *Wavelets and wideband correlation processing*, Signal Processing Magazine, **11**(1) (1992), 13–32.
- [71] A. S. Willsky, *On the algebraic structure of certain partially observable finite-state Markov processes*, Inform. Contr. **38**, (1978), 179–212.
- [72] C. M. Woodman, *The symmetry groups of non-rigid molecules as semi-direct products*, Mol. Phys. (6) **19** (1970), 753–780.




Figure 1: The SHRTs $\mathcal{T}_{(3,4)}$ and $\mathcal{T}_{(3,2,3)}$




Figure 2: The SHRT $\mathcal{T}_{(3,2,3)}$ and illustrations of some SHRT terminology




Figure 3: The first truncation of $\mathcal{T}_{(3,2,3)}$




Figure 4: Two symmetries of $\mathcal{T}_{(3,4)}$




Figure 5: A schematic of the nested decomposition

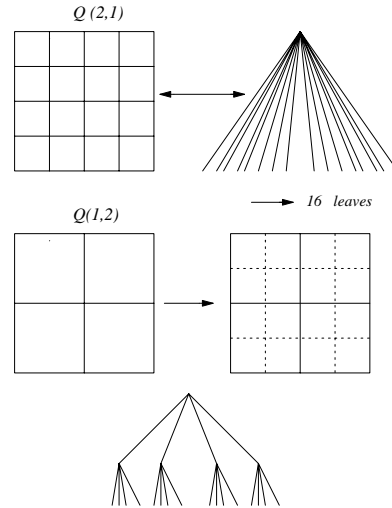


Figure 6: Two quadtree indexings




Figure 7: Quadtree structure of tree $Q(1,9)$ and some preferred group elements of $Z(1,9)$





Figure 8: Image quadrants



(a) Tree $Q(1,2)$

(b) Decomposition subspaces

Figure 9: Tree $Q(1,2)$ and decomposition associated subspaces


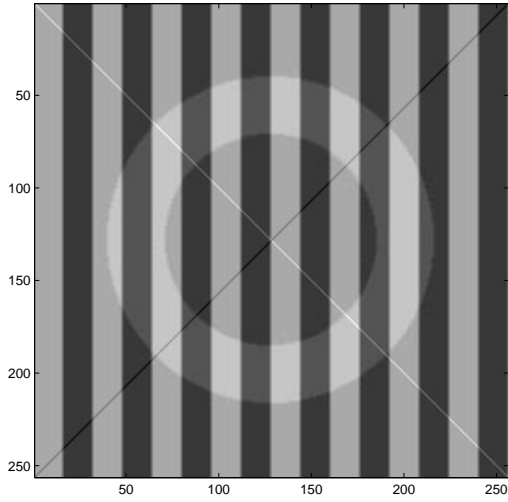



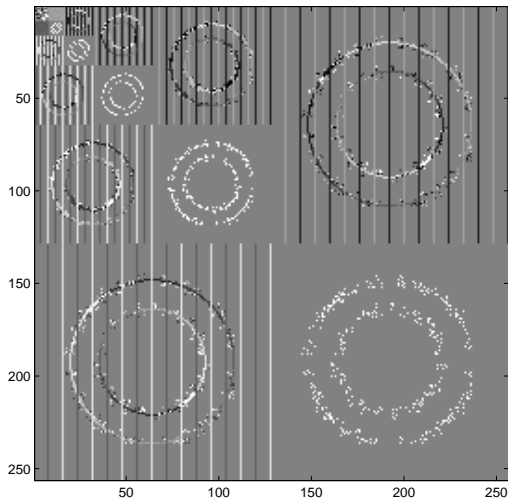
Figure 10: Tree $Q(1,4)$ with subspace decomposition



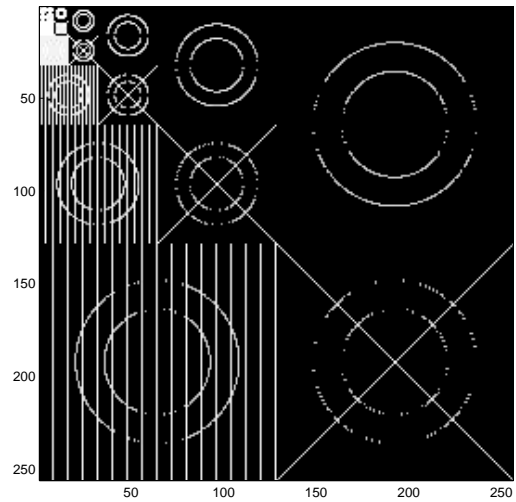
(a) Image



(b) WPT spectrum - amplitude

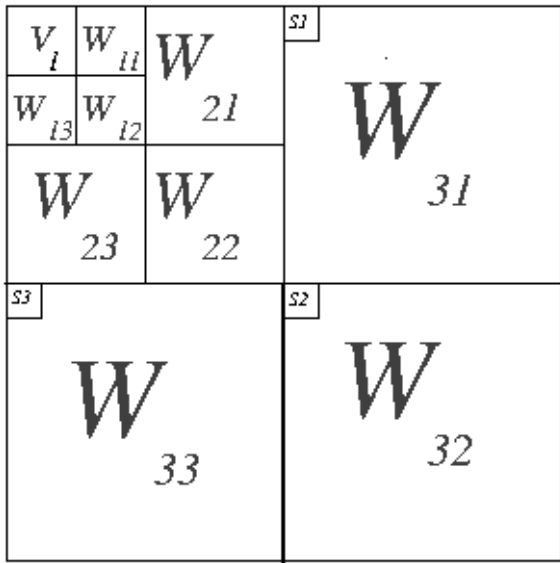


(c) WPT spectrum - phase

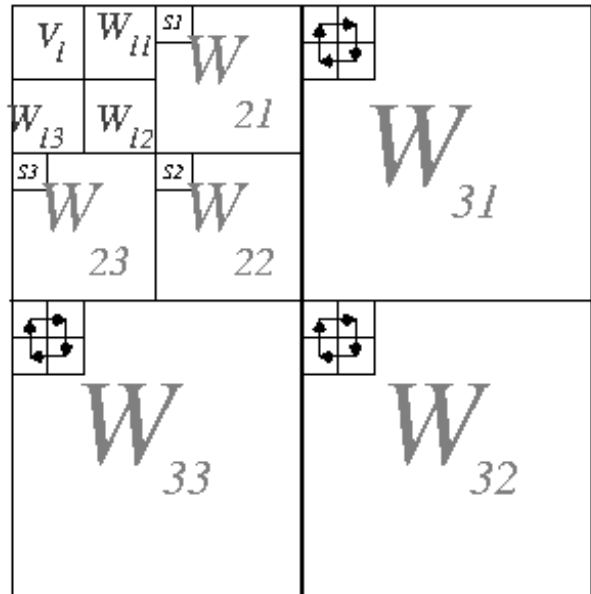


(d) Haar spectrum - amplitude

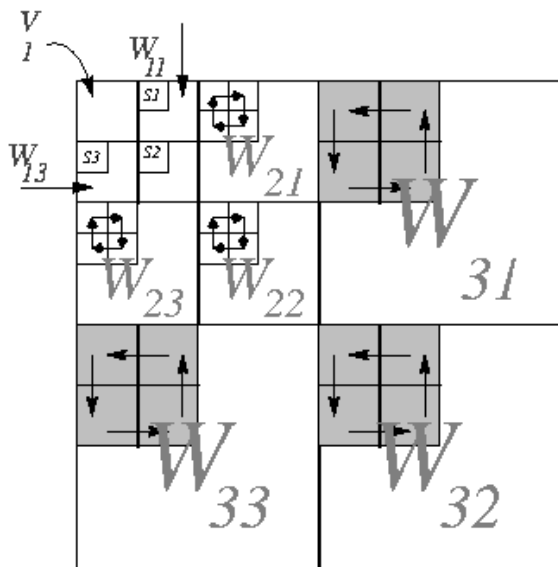
Figure 11: Multiresolution spectra for the Wreath Product and Haar transforms



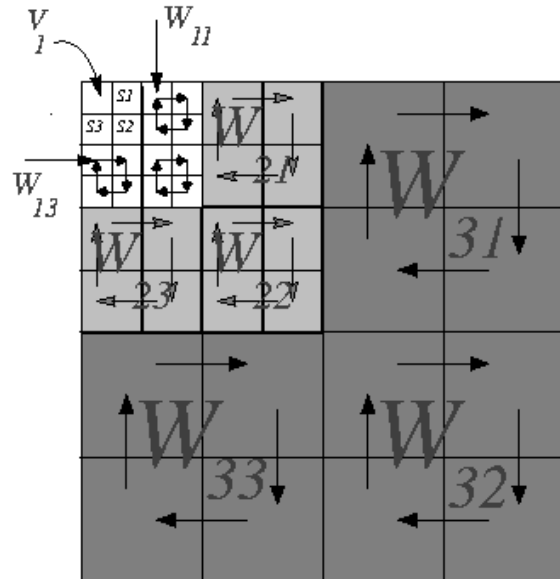
(a) Effect of $\alpha_0^{(3)}$



(b) Effect of $\alpha_0^{(2)}$



(c) Effect of $\alpha_0^{(1)}$




(d) Effect of $\alpha_0^{(0)}$

Figure 12: Spectral invariance with tree $Q(1,4)$

Subspaces	Group Elements								
	$a_i^{(0)}$	$a_i^{(1)}$	$a_i^{(2)}$	$a_i^{(3)}$	$a_i^{(4)}$	$a_i^{(5)}$	$a_i^{(6)}$	$a_i^{(7)}$	$a_i^{(8)}$
w_0	S	-	-	-	-	-	-	-	-
w_1	C	S	-	-	-	-	-	-	-
w_2	C	C	S	-	-	-	-	-	-
w_3	C	C	C	S	-	-	-	-	-
w_4	C	C	C	C	S	-	-	-	-
w_5	C	C	C	C	C	S	-	-	-
w_6	C	C	C	C	C	C	S	-	-
w_7	C	C	C	C	C	C	C	S	-
w_8	C	C	C	C	C	C	C	C	S


S - scaling
 C - cyclic shift

Table 1: Effect of group operations on subspaces




(a) Image

(b) Translated Image



(c) Rotated Image


(d) WPT Spectrum - Image




(e) WPT Spectrum - Translation

(f) WPT Spectrum - Rotation

Figure 13: Illustration of WPT spectral invariance

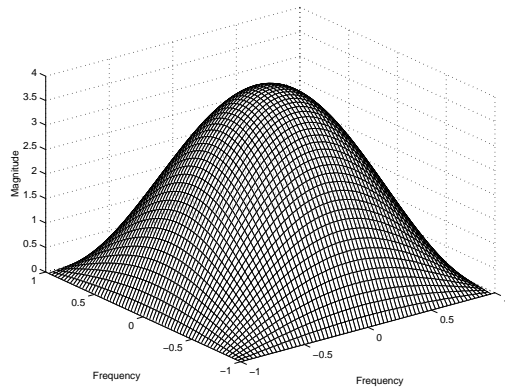


(a) Haar




(b) WPT


Figure 14: Effect of the Haar and WPT on a 2×2 grid




(a) Haar- Lowpass Filter




(b) Haar- Highpass Filter (Horizontal Edges)




(c) Haar- Highpass Filter (Vertical Edges)



(d) Haar- Highpass Filter (Diagonal Edges)



(e) WP- Highpass Filter (Complex)



(f) WP- Highpass Filter (Complex Conjugate)

Figure 15: Frequency response of the 2-dimensional filters