A *Fast* Fourier Transform with Rectangular Output on the BCC and FCC Lattices

Usman R. Alim [(1)] and Torsten Möller [(1)]

(1) School of Computing Science, Simon Fraser University, Burnaby BC V5A 1S6, Canada.
ualim@cs.sfu.ca, torsten@cs.sfu.ca

Abstract:

This paper discusses the efficient, non-redundant evaluation of a Discrete Fourier Transform on the three dimensional Body-Centered and Face-Centered Cubic lattices. The key idea is to use an axis aligned window to truncate and periodize the sampled function which leads to separable transforms. We exploit the geometry of these lattices and show that by choosing a suitable non-redundant rectangular region in the frequency domain, the transforms can be efficiently evaluated using the Fast Fourier Transform.

1. Introduction

The Discrete Fourier Transform (DFT) is an important tool used to analyze and process data in an arbitrary number of dimensions. Most applications of the DFT in higher dimensions, however, rely on a tensor product extension of a one-dimensional DFT, with the assumption that the underlying data is sampled on a Cartesian lattice. This extension has the advantage that it allows for a straightforward application of the Fast Fourier Transform (FFT).

The Cartesian lattice is known to be sub-optimal when it comes to sampling a band-limited function in two or higher dimensions [6]. In 3D, for instance, the Body-Centered Cubic (BCC) lattice is the optimal sampling lattice and yields a 30% savings in samples as compared to the Cartesian lattice [8]. The Face-Centered Cubic (FCC) lattice, although not optimal, is still better than the Cartesian lattice and is also the lattice that yields the minimum amount of Fourier-domain aliasing when sampling a general trivariate function [2].

From the perspective of continuous signal reconstruction, both the BCC and FCC lattices have received considerable attention because of their many applications in Visualization and Computer Graphics. Entezari et al. have devised a set of Box-Splines that can be used for signal approximation on the BCC [3] as well as the FCC [5] lattices. However, very little effort has gone into the development of discrete processing tools that are suitable for these non-Cartesian lattices.

The idea of a multidimensional DFT (MDFT) on non-Cartesian lattices is not new. Mersereau provided a derivation of a DFT for a hexagonally periodic sequence and designed other digital filters suitable for a 2D hexagonal lattice [6]. Later, the idea was extended to higher dimensions and a MDFT for arbitrary sampling lattices was proposed [7]. Guessoum et al. proposed an algorithm for evaluating the MDFT that has the same computational complexity as the Cartesian DFT [4].

Recently, Csébfalvi et al. [1] applied the MDFT to the BCC and FCC lattices by choosing a Cartesian periodicity in the spatial domain which leads to a Cartesian sampling of the Fourier transform. This allows the MDFT to be written in a separable form that can be evaluated via the FFT. However, their representation is redundant and leads to inefficient transforms. The aim of this paper is to revisit these transforms and show that they can be computed much more efficiently by exploiting the geometric properties of the BCC and FCC lattices to eliminate the redundancy.

The paper is organized as follows. We provide a basic review of multidimensional sampling in Section 2. which is later used in the derivation of a fast DFT for BCC and FCC lattices in Section 3. Some properties of these transforms are discussed in Section 4. and a summary is presented in Section 5.

2. Optimal Trivariate Sampling

Let $f_c(\boldsymbol{x})$ be a continuous trivariate function and $F_c(\boldsymbol{\xi})$ be its Fourier transform defined as

$$F_c(\boldsymbol{\xi}) = \int_{\mathbb{R}^3} f_c(\boldsymbol{x}) \exp[-2\pi j \boldsymbol{\xi}^T \boldsymbol{x}] d\boldsymbol{x} \qquad (1)$$

where T denotes the transpose operation. Let $f(\boldsymbol{n})$ be the sampled sequence obtained by sampling the function through

$$f(\boldsymbol{n}) = f_c(\boldsymbol{L}\boldsymbol{n}) \qquad (2)$$

where $\boldsymbol{L}$ is a 3×3 sampling matrix and $\boldsymbol{n}$ is an integer vector. Sampling on the lattice defined by the matrix $\boldsymbol{L}$ amounts to a periodization of the Fourier spectrum on a reciprocal lattice generated by the matrix $\boldsymbol{L}^{-T}$. In particular, the spectrum of the sampled sequence is given by [9]

$$\hat{F}(\boldsymbol{\xi}) = \frac{1}{|\det \boldsymbol{L}|} \sum_{\boldsymbol{r}} F_c(\boldsymbol{\xi} - \boldsymbol{L}^{-T} \boldsymbol{r}) \qquad (3)$$

where $\boldsymbol{r}$ is any integer vector.

If we assume that $f_c(\boldsymbol{x})$ is isotropically band-limited (i.e $F_c(\boldsymbol{\xi}) = 0$ for $\|\boldsymbol{\xi}\| > \xi_0$ for some band-limit ξ_0), then one of the lattices that achieves the tightest possible packing of the spectrum replicas (spheres) in the Fourier domain is

the FCC lattice. Thus, in order to sample a trivariate band-limited function optimally, the function should be sampled on the reciprocal of the FCC lattice, i.e. the BCC lattice.
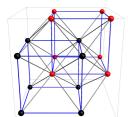
3. Discrete Fourier Transform

If the sequence $f(\boldsymbol{n})$ is non-zero within a finite region, it can be periodically extended spatially and represented as a Fourier series which is a sampled version of the transform (3) [7]. The pattern with which the continuous transform (3) is sampled in the Fourier domain depends on the periodicity pattern in the spatial domain. Merserau et al. [7] used a periodicity matrix to define the periodic extension of the finite sequence. Here, we use a somewhat different approach by splitting the sampled sequence into constituent Cartesian sequences [1].

The BCC and FCC lattices $\mathcal{L}_B$ and $\mathcal{L}_F$ are generated by the integer sampling matrices

$$\boldsymbol{L}_B = \begin{bmatrix} 1 & -1 & 1 \\ -1 & 1 & 1 \\ 1 & 1 & -1 \end{bmatrix} \quad \text{and} \quad \boldsymbol{L}_F = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

respectively. Both these lattices are based on a cubic sampling pattern whereby, in addition to samples at the eight corners of a cube, $\mathcal{L}_B$ has an additional sample in the center of the cube and $\mathcal{L}_F$ has six additional samples on the faces. Both these lattices can also be built from shifts of a Cartesian sublattice as shown in Fig. 1. In particular, samples that lie on the corners of cubes form the sublattice $2\mathbb{Z}^3$. The quotient group $\mathcal{L}_B/2\mathbb{Z}^3$ is isomorphic to $\mathbb{Z}_2$ and the quotient group $\mathcal{L}_F/2\mathbb{Z}^3$ is isomorphic to $\mathbb{Z}_4$. Therefore, $\mathcal{L}_B$ can be partitioned into two Cartesian cosets while $\mathcal{L}_F$ has four Cartesian cosets (Fig. 1).
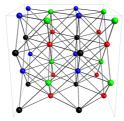

Figure 1: Left, the BCC lattice, a 16 point view. Right, the FCC lattice, a 32 point view. Lattice sites that are Voronoi neighbors are linked to each other. Cosets are indicated by different colors.

3.1 BCC DFT

The BCC lattice with arbitrary scaling is obtained via the sampling matrix $h\boldsymbol{L}_B$ where h is a positive scaling parameter. The Voronoi cell is a truncated octahedron having a volume of $|\det h\boldsymbol{L}_B| = 4h^3$. The Voronoi cell of the reciprocal FCC lattice is a rhombic dodecahedron having a volume of $\frac{1}{4h^3}$. Since $\mathcal{L}_B$ has two Cartesian cosets, a sampled sequence can be split up into two subsequences given by

$$f_0(\boldsymbol{n}) = f_c(2h\boldsymbol{I}\boldsymbol{n}) \quad \text{and} \quad f_1(\boldsymbol{n}) = f_c(2h\boldsymbol{I}\boldsymbol{n} + h\boldsymbol{t}),$$

where $\boldsymbol{I}$ is the 3×3 identity matrix, $\boldsymbol{t}$ is the translation vector $(1,1,1)^T$ and $\boldsymbol{n} = (n_1, n_2, n_3)^T$ is any integer vector. $f_0(\boldsymbol{n})$ is the sequence associated with the first coset while $f_1(\boldsymbol{n})$ is associated with the second. Since these sequences are sampled on a Cartesian pattern, a straightforward truncation of the original sequence is to

choose a cuboid shaped fundamental region generated by limiting $\boldsymbol{n}$ to the set $\mathcal{N} := \{\boldsymbol{n} \in \mathbb{Z}^3 : 0 \leq n_1 < N_1, 0 \leq n_2 < N_2, 0 \leq n_3 < N_3\}$ for some positive integers N_1, N_2 and N_3. This region consists of $2N_1N_2N_3$ data points (i.e. Voronoi cells) and has a total volume of $8N_1N_2N_3h^3$. If we define $\boldsymbol{N}$ to be the diagonal matrix $\text{diag}(N_1, N_2, N_3)$, then the two subsequences $f_0(\boldsymbol{n})$ and $f_1(\boldsymbol{n})$ contained within the fundamental region can be periodically extended on a Cartesian pattern such that they satisfy

$$f_0(\boldsymbol{n} + \boldsymbol{N}\boldsymbol{r}) = f_0(\boldsymbol{n}) \quad \text{and} \quad f_1(\boldsymbol{n} + \boldsymbol{N}\boldsymbol{r}) = f_1(\boldsymbol{n}),$$

for all $\boldsymbol{n}$ and $\boldsymbol{r}$ in $\mathbb{Z}^3$.

This Cartesian periodic extension in the spatial domain amounts to a Cartesian sampling in the Fourier domain. In particular, the continuous transform (3) is sampled at the frequencies $\boldsymbol{\xi} = \frac{1}{2h}\boldsymbol{N}^{-1}\boldsymbol{k}$ yielding the sequence

$$
\begin{aligned}
F(\boldsymbol{k}) =& \hat{F}(\boldsymbol{\xi})\big|_{\boldsymbol{\xi}=\frac{1}{2h}\boldsymbol{N}^{-1}\boldsymbol{k}} \\
=& \sum_{\boldsymbol{n}\in\mathcal{N}} f_0(\boldsymbol{n}) \exp\left[\frac{-2\pi j}{2h}\boldsymbol{k}^T\boldsymbol{N}^{-1}2h\boldsymbol{I}\boldsymbol{n}\right] + \\
& f_1(\boldsymbol{n}) \exp\left[\frac{-2\pi j}{2h}\boldsymbol{k}^T\boldsymbol{N}^{-1}(2h\boldsymbol{I}\boldsymbol{n}+h\boldsymbol{t})\right] \\
=& \sum_{\boldsymbol{n}\in\mathcal{N}} \left(f_0(\boldsymbol{n})+f_1(\boldsymbol{n})\exp\left[-\pi j\boldsymbol{k}^T\boldsymbol{N}^{-1}\boldsymbol{t}\right]\right)\cdot \\
& \exp\left[-2\pi j\boldsymbol{k}^T\boldsymbol{N}^{-1}\boldsymbol{n}\right],
\end{aligned}
\tag{4}
$$

where $\boldsymbol{k} = (k_1, k_2, k_3)^T \in \mathbb{Z}^3$ is the frequency index vector. The above equation defines a forward BCC DFT. Since it is a sampled version of a continuous transform that is periodic on an FCC lattice, it should be invariant under translations that lie on the reciprocal lattice generated by the matrix $(h\boldsymbol{L}_B)^{-T} = \frac{1}{2h}\boldsymbol{L}_F$. This property is easily demonstrated as follows. If $\boldsymbol{r} \in \mathbb{Z}^3$, then after substituting $\boldsymbol{\xi} = \frac{1}{2h}(\boldsymbol{N}^{-1}\boldsymbol{k} + \boldsymbol{L}_F\boldsymbol{r})$ in (4) and simplifying, we get

$$
\begin{aligned}
& \hat{F}\left(\frac{1}{2h}(\boldsymbol{N}^{-1}\boldsymbol{k} + \boldsymbol{L}_F\boldsymbol{r})\right) \\
=& \sum_{\boldsymbol{n}\in\mathcal{N}} \left(f_0(\boldsymbol{n})+f_1(\boldsymbol{n})\exp\left[-\pi j(\boldsymbol{k}^T\boldsymbol{N}^{-1}+\boldsymbol{r}^T\boldsymbol{L}_F)\boldsymbol{t}\right]\right)\cdot \\
& \exp\left[-2\pi j(\boldsymbol{k}^T\boldsymbol{N}^{-1}+\boldsymbol{r}^T\boldsymbol{L}_F)\boldsymbol{n}\right] \\
=& \hat{F}\left(\frac{1}{2h}\boldsymbol{N}^{-1}\boldsymbol{k}\right),
\end{aligned}
$$

since $\boldsymbol{r}^T\boldsymbol{L}_F\boldsymbol{n}$ is always an integer and $\boldsymbol{r}^T\boldsymbol{L}_F\boldsymbol{t}$ is always even.

One fundamental period of the BCC DFT is contained within a rhombic dodecahedron of volume $\frac{1}{4h^3}$. The sampling density in the frequency domain is given by $|\det\frac{1}{2h}\boldsymbol{N}^{-1}| = (8N_1N_2N_3h^3)^{-1}$. Thus, the fundamental period consists of a total of $2N_1N_2N_3$ distinct frequency samples which is the same as the number of distinct spatial samples.

The inverse BCC DFT is obtained by summing over all the distinct sinusoids and evaluating them at the spatial sample locations. This gives

$$f_0(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k}\in\mathcal{K}} F(\boldsymbol{k}) \exp\left[2\pi j\boldsymbol{k}^T\boldsymbol{N}^{-1}\boldsymbol{n}\right] \tag{5a}$$

$$f_1(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k}\in\mathcal{K}} F(\boldsymbol{k}) \exp\left[2\pi j\boldsymbol{k}^T\boldsymbol{N}^{-1}(\boldsymbol{n}+\frac{1}{2}\boldsymbol{t})\right] \tag{5b}$$

where $N = 2N_1 N_2 N_3$ is the number of samples and $\mathcal{K} \subset \mathbb{Z}^3$ is any set that indexes all the distinct frequency samples. It is easily verified that both the sequences (5a) and (5b) are periodic with periodicity matrix $\boldsymbol{N}$.

3.1.1 Efficient Evaluation

Since $\boldsymbol{N}$ is diagonal, the kernel in both equations (4) and (5) is separable. This suggests that the transform can be efficiently computed via the rectangular multidimensional FFT, provided that a suitable rectangular index set $\mathcal{K}$ can be found. Observe that the Cartesian sequence $F(\boldsymbol{k})$ is periodic with periodicity matrix $2\boldsymbol{N}$, i.e $F(\boldsymbol{k} + 2\boldsymbol{N}\boldsymbol{r}) = F(\boldsymbol{k})$ for all $\boldsymbol{r} \in \mathbb{Z}^3$. Therefore, one way to obtain a rectangular index set is to choose $\mathcal{K}$ such that it contains all the frequency indices within one period generated by the matrix $2\boldsymbol{N}$. This consists of a total of $|\det 2\boldsymbol{N}| = 4N$ indices and hence contains four replicas of the fundamental rhombic dodecahedron.

A non-redundant rectangular index set can be found by exploiting the geometric properties of the FCC lattice. If we consider the first octant only, $4N$ samples are contained within a cube formed by the FCC lattice sites that have even parity. This cube also contains six face-centered sites. By joining any two axially opposite face-centered sites, we can split the cube into four rectangular regions such that each region consists of non-redundant samples only. Six rhombic dodecahedra contribute to such a region as illustrated in Fig. 2. The non-redundant region shown in Fig. 2b is obtained by limiting $\boldsymbol{k}$ to the index set given by $\mathcal{K} = \{\boldsymbol{k} \in \mathbb{Z}^3 : 0 \leq k_1 < N_1, 0 \leq k_2 < N_2, 0 \leq k_3 < 2N_3\}$.
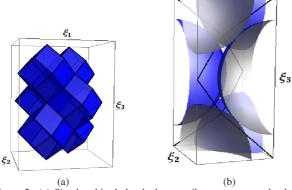

(a)   (b)

Figure 2: (a) Six rhombic dodecahedra contribute to a non-redundant rectangular region. (b) Zoomed in view of the non-redundant rectangular region that contains the full spectrum split into six pieces. ξ_1, ξ_2 and ξ_3 indicate the principal directions in the frequency domain.

This region can further be subdivided into two cubes stacked on top of each other, each containing $N_1 \times N_2 \times N_3$ samples. The forward transform (4) can then be evaluated in the two cubes separately by appropriately applying the Cartesian FFT to the two sequences $f_0(\boldsymbol{n})$ and $f_1(\boldsymbol{n})$ and combining the results together. After rearranging terms in (4), the forward transform in the bottom cube becomes

$$F_0(\boldsymbol{k}) = F(\boldsymbol{k}) = \sum_{\boldsymbol{n}\in\mathcal{N}} f_0(\boldsymbol{n}) \exp\left[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right] + \exp\left[-\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{t}\right] \sum_{\boldsymbol{n}\in\mathcal{N}} f_1(\boldsymbol{n}) \exp\left[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right], \quad (6)$$

where $\boldsymbol{k}$ is now restricted to the set $\mathcal{N}$. Since this equation is valid for all $\boldsymbol{k} \in \mathbb{Z}^3$, the forward transform in

the top cube can be computed from (6) by $F_1(\boldsymbol{k}) = F_0(\boldsymbol{k} + (0, 0, N_3)^T)$ which simplifies to

$$F_1(\boldsymbol{k}) = \sum_{\boldsymbol{n}\in\mathcal{N}} f_0(\boldsymbol{n}) \exp\left[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right] - \exp[-\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{t}] \sum_{\boldsymbol{n}\in\mathcal{N}} f_1(\boldsymbol{n}) \exp\left[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right], \quad (7)$$

for $\boldsymbol{k} \in \mathcal{N}$. Equations (6) and (7) are now in a form that permits a straightforward application of the Cartesian FFT. Since the two equations are structurally similar, only two $N_1 \times N_2 \times N_3$ FFT computations are needed, one for the sequence $f_1(\boldsymbol{n})$ and one for $f_2(\boldsymbol{n})$.

In a similar fashion, the inverse transform (5) can be computed using two inverse FFT computations. Splitting the summations in (5) into the two constituent cubes gives

$$f_0(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k}\in\mathcal{N}} (F_0(\boldsymbol{k}) + F_1(\boldsymbol{k})) \exp\left[2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right],$$

$$f_1(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k}\in\mathcal{N}} \left((F_0(\boldsymbol{k}) - F_1(\boldsymbol{k})) \exp[\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{t}]\right) \cdot \exp\left[2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}\boldsymbol{n}\right]. \quad (8)$$

3.2 FCC DFT

The FCC lattice with arbitrary scaling is generated by the sampling matrix $h\boldsymbol{L}_{\mathsf{F}}$. The rhombic dodecahedral Voronoi cell has a volume of $|\det h\boldsymbol{L}_{\mathsf{F}}| = 2h^3$. The frequency spectrum is replicated according to (3) on a reciprocal BCC lattice that has a truncated octahedral Voronoi cell having a volume of $\frac{1}{2h^3}$.

A sequence sampled on the FCC lattice can be split up into four Cartesian subsequences corresponding to the four Cartesian cosets. Each subsequence is given by

$$f_i(\boldsymbol{n}) = f_c(2h\boldsymbol{I}\boldsymbol{n} + h\boldsymbol{t}_i),$$

where $i \in \{0, 1, 2, 3\}$ and $\boldsymbol{t}_i$ are the integer shift vectors $(0, 0, 0)^T$, $(1, 0, 1)^T$, $(0, 1, 1)^T$ and $(1, 1, 0)^T$ respectively. Analogous to the BCC case, let us choose a rectangular truncation of the original sequence by limiting $\boldsymbol{n}$ to the set $\mathcal{N}$ and extend the sequences periodically so that they satisfy $f_i(\boldsymbol{n} + \boldsymbol{N}\boldsymbol{r}) = f_i(\boldsymbol{n})$. This truncation yields a rectangular fundamental region in the spatial domain consisting of a total of $N = 4N_1 N_2 N_3$ distinct samples. Therefore, each truncated octahedron in the frequency domain tessellation will consist of N distinct points that are sampled in a Cartesian fashion at the frequencies $\boldsymbol{\xi} = \frac{1}{2h}\boldsymbol{N}^{-1}\boldsymbol{k}$ where $\boldsymbol{k} \in \mathbb{Z}^3$. The sampled sequence in the frequency domain is thus given by

$$F(\boldsymbol{k}) = \hat{F}(\frac{1}{2h}\boldsymbol{N}^{-1}\boldsymbol{k})$$

$$= \sum_{\boldsymbol{n}\in\mathcal{N}} \sum_{i=0}^{3} f_i(\boldsymbol{n}) \exp\left[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}(\boldsymbol{n} + \frac{1}{2}\boldsymbol{t}_i)\right]. \quad (9)$$

This defines a forward FCC DFT. Like the BCC case, it is invariant under shifts of the type $\boldsymbol{\xi} = \frac{1}{2h}(\boldsymbol{N}^{-1}\boldsymbol{k} + \boldsymbol{L}_{\mathsf{B}}\boldsymbol{r})$ making it periodic on a BCC lattice with one fundamental period contained in a truncated octahedron.

The inverse FCC DFT is obtained by summing over all the distinct sinusoids evaluated at the spatial sample locations

$$f_i(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k}\in\mathcal{K}} F(\boldsymbol{k}) \exp\left[2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1}(\boldsymbol{n} + \frac{1}{2}\boldsymbol{t}_i)\right], \quad (10)$$

where $\mathcal{K} \subset \mathbb{Z}^3$ is any set that indexes all the N distinct sinusoids.

3.2.1 Efficient Evaluation

Since N is diagonal, the key to efficiently evaluating the FCC DFT pair (9) and (10) is to choose a suitable rectangular region in the frequency domain that contains N distinct samples. Similar to the BCC DFT, the sequence (9) is $2N$ periodic with one complete rectangular period containing $|\det 2N| = 2N$ samples and hence two complete spectrum replicas. These $2N$ samples are contained within a cube, the corners of which lie at the even parity points of the BCC lattice. This cubic region can be split into two by halving along any of the three principal directions yielding a rectangular region that contains only non-redundant samples as illustrated in Fig. 3. The index set that spans the region depicted in Fig. 3b is given by $\mathcal{K} = \{\boldsymbol{k} \in \mathbb{Z}^3 : 0 \le k_1 < 2N_1, 0 \le k_2 < 2N_2, 0 \le k_3 < N_3\}$.
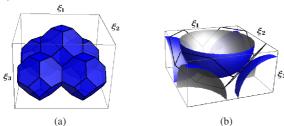

(a)   (b)

Figure 3: (a) Five truncated octahedra contribute to a non-redundant rectangular region. (b) Zoomed in view of the rectangular region that contains the full spectrum.

The non-redundant region can be split into four $N_1 \times N_2 \times N_3$ cubic subregions and the forward transform (9) can be evaluated in each of the subregions separately by appropriately applying the FFT to each of the subsequences $f_i(\boldsymbol{n})$ and combining the output. The derivation is very similar to the BCC case and we leave the details to the reader. The forward transform in each subregion can be written as

$$F_m(\boldsymbol{k}) = \sum_{i=0}^{3} H_{im} \exp[-\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1} \boldsymbol{t}_i] \cdot \left(\sum_{\boldsymbol{n} \in \mathcal{N}} f_i(\boldsymbol{n}) \exp[-2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1} \boldsymbol{n}] \right), \quad (11)$$

where $m \in \{0, 1, 2, 3\}$, $\boldsymbol{k} \in \mathcal{N}$ and H_{im} is an element of the 4×4 Hadamard matrix $\boldsymbol{H} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}$. The four subregions $F_m(\boldsymbol{k})$ have their bottom left corners at the frequency index vectors $(0, 0, 0)^T$, $(N_1, 0, 0)^T$, $(0, N_2, 0)^T$ and $(N_1, N_2, 0)$ respectively.

Likewise, the inverse transform (10) can be evaluated using four inverse FFT evaluations, one for each of the subsequences. This yields

$$f_i(\boldsymbol{n}) = \frac{1}{N} \sum_{\boldsymbol{k} \in \mathcal{N}} \left(\exp[\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1} \boldsymbol{t}_i] \sum_{m=0}^{3} H_{im} F_m(\boldsymbol{k}) \right) \cdot \exp[2\pi j \boldsymbol{k}^T \boldsymbol{N}^{-1} \boldsymbol{n}]. \quad (12)$$

4. Discussion

The decomposition of the non-redundant region in the frequency domain into cubes leads to transforms that are much more efficient. Both the BCC and FCC DFTs proposed by Csébfalvi et al. [1] are redundant and require the FFT of a $2N_1 \times 2N_2 \times 2N_3$ sequence. In contrast, our proposed evaluation strategy eliminates the redundancy and computes only two $N_1 \times N_2 \times N_3$ FFTs for the BCC case and four $N_1 \times N_2 \times N_3$ FFTs for the FCC case.

Any operation in the frequency domain must respect the arrangement of the different portions of the spectrum. The BCC DFT splits the spectrum into six parts as illustrated by the six pieces (two lunes and four spherical triangles) of the sphere in Fig. 2b. The FCC transform splits the frequency spectrum into five parts as indicated by the hemisphere and the four spherical triangles in Fig. 3b.

5. Summary

In this paper, we have shown that a MDFT of a Cartesian periodic sequence sampled on the BCC or FCC lattices can be efficiently evaluated using the FFT. The BCC lattice can be represented as two shifted Cartesian lattices. This representation leads to a separable transform that is efficiently computed via two non-redundant FFT evaluations of the Cartesian subsequences. Similarly, the FCC lattice consists of four shifted Cartesian lattices and the MDFT requires four non-redundant FFT evaluations.

References:

[1] B. Csébfalvi and B. Domonkos. Pass-Band Optimal Reconstruction on the Body-Centered Cubic Lattice. In *Vision, Modeling, and Visualization 2008: Proceedings, October 8-10, 2008, Konstanz, Germany*, page 71. IOS Press, 2008.

[2] A. Entezari. *Optimal Sampling Lattices and Trivariate Box Splines*. PhD thesis, Simon Fraser University, Vancouver, Canada, July 2007.

[3] A. Entezari, D. Van De Ville, and T. Möller. Practical box splines for volume rendering on the body centered cubic lattice. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):313 – 328, 2008.

[4] A. Guessoum and R. Mersereau. Fast algorithms for the multidimensional discrete Fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-34(4):937–943, 1986.

[5] M. Kim, A. Entezari, and J. Peters. Box Spline Reconstruction on the Face Centered Cubic Lattice. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization/Information Visualization 2008)*, 14(6):1523–1530, 2008.

[6] R. Mersereau. The Processing of Hexagonally Sampled Two-dimensional Signals. *Proceedings of the IEEE*, 67(6):930–949, June 1979.

[7] R. Mersereau and T. Speake. The processing of periodically sampled multidimensional signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, (1):188–194, 1983.

[8] T. Theußl, T. Möller, and M. Gröller. Optimal regular volume sampling. In *Proceedings of the conference on Visualization'01*, pages 91–98. IEEE Computer Society Washington, DC, USA, 2001.

[9] P. Vaidyanathan. Fundamentals of multidimensional multirate digital signal processing. *Sadhana*, 15(3):157–176, 1990.