Eurographics Conference on Visualization (EuroVis) (2015)
E. Bertini, J. Kennedy and E. Puppo (Editors)

A Quality-Preserving Cartesian to Body-Centered Cubic Downsampling Transform

Usman R. Alim[1] and Thiago Valentin de Oliveira[2]

[1]University of Calgary, Calgary AB, Canada
[2]Universidade Federal do Rio de Janeiro, Rio de Janeiro RJ, Brazil

Abstract

The body-centered cubic lattice is the optimal sampling lattice in three dimensions. However, most volumetric datasets are acquired on the well-known Cartesian cubic lattice. In order to leverage the approximation capabilities of the body-centred cubic lattice, we propose a factor-of-four Cartesian to body-centered downsampling transform. We derive a Fourier domain post-aliasing error kernel and use it to optimize the cosine-weighted trilinear B-spline kernel. We demonstrate that our downsampling transform preserves fidelity when an oversampled function of interest is reconstructed with trilinear interpolation on the fine-scale Cartesian grid, and optimized cosine-weighted trilinear approximation on the coarse-scale body-centered cubic grid.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

The Cartesian cubic (CC) lattice $\mathbb{Z}^3$ is the de facto standard for the representation of volumetric data. It is conceptually simple, easy to implement, and is supported by modern graphics processing units (GPUs). Moreover, many classical univariate signal processing techniques can easily be extended to the CC lattice via a tensor product. It is well known that CC sampling is not optimal, and other lattices such as the body-centered cubic (BCC) lattice and the face-centered cubic (FCC) lattice are better, i.e. they can represent a signal with similar fidelity with fewer samples. Despite this, the CC lattice continues to enjoy tremendous support owing to the widespread availability of acquisition and processing tools. There has been a growing interest in studying non-Cartesian cubic lattices in the graphics and scientific visualization communities. Recent research advocates using non-Cartesian lattices instead of the ubiquitous CC lattice [ME10, Csé13, VCRG14]. In this work, we take a slightly different stand and promote the BCC lattice as a downsampled intermediate representation that achieves a similar quality as compared to a fine-scale CC lattice.

Most comparisons of the cubic lattices are based on a geometric argument that looks at the sphere packing efficiency of the dual lattice in the Fourier domain [TMG01, VCRG14]. Such a comparison is of theoretical interest but is not completely applicable to practical scenarios since it requires the use of infinitely supported multidimensional sinc kernels [YE12]. For practical scenarios, compact kernels are preferred since they result in efficient reconstruction schemes. Motivated by these goals, we propose a practical CC→BCC downsampling strategy that is firmly built upon approximation theoretic principles. The relevance of trilinear interpolation on the CC lattice can hardly be disputed. Therefore, our proposed scheme is designed to achieve similar quality on a coarse-scale BCC lattice as compared to trilinear interpolation on a fine-scale CC lattice. Among the plethora of kernels available for the BCC lattice, we choose the recently proposed cosine-weighted linear B-spline (CWLB) [Csé13] for its practical significance. In order to compare the approximation quality, we derive a Fourier domain post-aliasing error kernel and use it to show that, when the original CC data is sufficiently oversampled, using an *optimized* CWLB kernel for a downsampled representation on the BCC lattice leads to a remarkable reduction factor of 4 with very little loss of quality. We provide experimental results that validate our theoretical findings and provide additional pointers of practical relevance.

2. Related Work

We restrict attention to works that have focused on designing compact spline-like kernels for the 3D cubic lattices. The in-

DOI: 10.2312/eurovisshort.20151119

terested reader may also want to consult a multidimensional signal processing textbook such as [DM84].

On the CC lattice, trivariate tensor-product extensions of the univariate B-splines [UAE93] are a popular choice. Non-separable spline kernels on the CC lattice have also been investigated [Chu88, dHR93, Wan01]. Most kernels on the BCC and FCC lattices are made up of box splines [dHR93]. Entezari *et al.* proposed second and fourth order box splines for function approximation on the BCC lattice [EDM04]. They later derived closed from polynomial representations [EVM08], proposed quasi-interpolants [EMK09] and implemented their efficient approximation algorithm on the GPU [FEVM10]. Kim *et al.* investigated box-spline generators on the FCC lattice [KEP08], and later generalized their construction scheme to the (non-Cartesian) *root* lattices in arbitrary dimension [KP10, KP11]. The idea of successively convolving the indicator function of the Voronoi cell has also been investigated. Mirzargar and Entezari proposed *Voronoi* splines [ME10], as well as their associated quasi-interpolants [ME11].

Multiresolution data representation on non-Cartesian cubic lattices is still a relatively unexplored area of research. Entezari *et al.* [EMBM06] proposed a subsampling strategy that works in the Fourier domain by eliminating the out-of-band portion of the spectrum. However, they did not account for the effect of the reconstruction kernel.

Owing to their non-cubic support, piecewise polynomial representation and GPU implementation of box-splines are tricky. Our work is inspired by another construction that aims to extend the trivariate tensor-product B-splines to the BCC lattice [Csé10, DC10, Csé13]. Of all the known compact admissible BCC kernels, these kernels provide the best quality and can easily be implemented on modern GPUs.

3. Background

3.1. Sampling and Reconstruction

A sampling lattice $\mathcal{L}$ is a set of sampling points generated by integer linear combinations of the columns of a generating matrix $\mathbf{L}$, i.e. $\mathcal{L} := \{\mathbf{Lk} : \mathbf{k} \in \mathbb{Z}^3\}$. For example, when $\mathbf{L} = \mathbf{I}$, the 3×3 identity matrix, we obtain the familiar CC lattice $\mathbb{Z}^3$. The BCC lattice $\mathcal{B}$ is obtained via the matrix

$$\mathbf{B} = \begin{bmatrix} 1 & -1 & -1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \end{bmatrix}. \tag{1}$$

It is also a sub-lattice of the CC lattice obtained by retaining those points whose coordinates are either all even or all odd. The even points form a Cartesian coset generated by $2\mathbf{I}$. The odd coset is given by $\{2\mathbf{k} + (1,1,1) : \mathbf{k} \in \mathbb{Z}^3\}$, i.e. it is obtained by shifting the even coset by $(1,1,1)$.

The density of a sampling lattice $\mathcal{L}$ is given by $|\det \mathbf{L}|$. Observe that, since $|\det \mathbf{B}| = 4$, the BCC lattice $\mathcal{B}$ is four times less dense as compared to the CC lattice $\mathbb{Z}^3$. The BCC lattice can therefore be normalized by multiplying $\mathbf{B}$ by $\sqrt[-3]{4}$. The normalized BCC lattice is denoted as $\mathcal{B}_0$. The dual of a lattice is denoted as $\mathcal{L}^\circ$; it is generated by the matrix $\mathbf{L}^{-T}$. The CC lattice is self-dual whereas the dual of a BCC lattice is an FCC lattice.

In order to approximate a function $f(\mathbf{x})$ via a lattice $\mathcal{L}$, one typically makes discrete measurements of f at the lattice sites $\mathbf{Lk}$ ($\mathbf{k} \in \mathbb{Z}^3$) to yield a coefficient sequence $c[\mathbf{k}]$. A linear combination of the lattice translates of a reconstruction kernel $\varphi(\mathbf{x})$ then yields the desired approximation, i.e.

$$\tilde{f}(\mathbf{x}) = \sum_{\mathbf{k}} c[\mathbf{k}]\varphi(\mathbf{x} - \mathbf{Lk}). \tag{2}$$

Even though the lattice dependence of $c[\cdot]$ is not explicitly indicated, it should be clear from the context established by the reconstruction equation (2).

Trilinear Interpolation on CC Trilinear interpolation on the CC lattice is well-known. A 3D tensor product of the univariate linear B-spline $\beta^1(x) := \max(0, 1 - |x|)$ yields the trivariate reconstruction kernel

$$L(\mathbf{x}) := \beta^1(x_1)\beta^1(x_2)\beta^1(x_3), \tag{3}$$

where $\mathbf{x} = (x_1, x_2, x_3)$. Trilinear interpolation on the CC lattice is therefore achieved by

$$f(\mathbf{x}) \approx \tilde{f}_1(\mathbf{x}) := \sum_{\mathbf{k} \in \mathbb{Z}^3} f[\mathbf{k}]L(\mathbf{x} - \mathbf{k}), \tag{4}$$

where $f[\mathbf{k}] := f(\mathbf{k})$ are the Cartesian samples of f. The kernel is completely supported within a cube of volume 8 which implies that upto 8 sample values contribute to the reconstruction at a general position.

CWLB Approximation on BCC Cosine-Weighted triLinear B-spline (CWLB) approximation [Csé13] on the BCC lattice makes use of the kernel

$$C_\lambda(\mathbf{x}) := L(\frac{\mathbf{x}}{2}) \underbrace{\left(\tfrac{1}{2} + \tfrac{\lambda}{6}(\cos \pi x_1 + \cos \pi x_2 + \cos \pi x_3)\right)}_{W_\lambda(\mathbf{x})}. \tag{5}$$

The corresponding approximation is therefore given by

$$\tilde{f}_2(\mathbf{x}) := \sum_{\mathbf{k} \in \mathbb{Z}^3} c[\mathbf{k}]C_\lambda(\mathbf{x} - \mathbf{Bk}), \tag{6}$$

where $c[\mathbf{k}]$ is a coefficient sequence associated with the BCC lattice, and may or may not be the same as the samples of the function f. Observe that the BCC approximation (6) is at a coarser scale as compared to the CC approximation (4). The kernel C_λ is compactly supported with a cube of volume 64. Since $|\det \mathbf{B}| = 4$, this means that upto 16 coefficients contribute to the reconstruction at a general position. As shown by Csébfalvi [Csé13], the reconstruction is easily implemented by first trilinearly interpolating on the even and odd cosets separately, and then combining the results using the weighting function W_λ. Thus, this reconstruction scheme is more expensive but is easily implemented on GPUs by taking advantage of GPU texture lookup capabilities. When $\lambda = 1$, the kernel C_λ is interpolating. In this case, we can use the lattice samples for interpolation, i.e. $c[\mathbf{k}] = f(\mathbf{Gk})$. On the other hand, when $\lambda \neq 0$, a prefilter needs to be applied to the samples of f to obtain the necessary coefficients $c[\cdot]$.

3.2. Post-aliasing Error

In order to characterize the aliasing error for an approximation scheme, we assume that the Fourier transform of f is isotropically bandlimited, i.e. $\hat{f}(\mathbf{u}) = 0$ whenever $\|\mathbf{u}\| > \frac{1}{2}$. Recall that $\hat{f}(\mathbf{u}) := \int_{\mathbb{R}^3} f(\mathbf{x}) \exp(-2\pi\imath\mathbf{u}\cdot\mathbf{x})$. The overall reconstruction error can be quantified in the Fourier domain by measuring the deviation of the reconstruction kernel $\hat{\varphi}$ from the ideal kernel. Recall that the spectrum of the ideal kernel is unity within the Voronoi cell of the dual $\mathcal{L}^\circ$ centered at the origin, and zero elsewhere. Assuming that the lattice under consideration is properly normalized ($|\det \mathbf{L}| = 1$), the overall reconstruction error $\|\hat{f} - \hat{\tilde{f}}\|_{L_2}^2$ is given by

$$\int_{\mathbb{V}_{\mathcal{L}^\circ}} |\hat{f}(\mathbf{u})|^2 \Big((1-\hat{\varphi}(\mathbf{u}))^2 + \sum_{\mathbf{r}\in\mathcal{L}^\circ\setminus\{\mathbf{0}\}} |\hat{\varphi}(\mathbf{u}-\mathbf{r})|^2 \Big) d\mathbf{u} \quad (7)$$

where $\mathbb{V}_{\mathcal{L}^\circ}$ denotes the Voronoi cell of the dual lattice centered at the origin. The first term quantifies smoothing whereas the second term is a measure of post-aliasing. When the data is oversampled, smoothing error is negligible and post-aliasing error dominates. Using some simple algebraic manipulations, we can express the post-aliasing term as follows.

$$\begin{aligned}
&\int_{\mathbb{V}_{\mathcal{L}^\circ}} |\hat{f}(\mathbf{u})|^2 \sum_{\mathbf{r}\in\mathcal{L}^\circ\setminus\{\mathbf{0}\}} |\hat{\varphi}(\mathbf{u}-\mathbf{r})|^2 d\mathbf{u} \\
&= \int_{\mathbb{V}_{\mathcal{L}^\circ}} |\hat{f}(\mathbf{u})|^2 \Big(\sum_{\mathbf{r}\in\mathcal{L}^\circ} |\hat{\varphi}(\mathbf{u}-\mathbf{r})|^2 - |\hat{\varphi}(\mathbf{u})|^2 \Big) d\mathbf{u} \\
&= \int_{\mathbb{V}_{\mathcal{L}^\circ}} |\hat{f}(\mathbf{u})|^2 \underbrace{\Big(\hat{A}_\varphi(\mathbf{u}) - |\hat{\varphi}(\mathbf{u})|^2 \Big)}_{E(\mathbf{u})} d\mathbf{u}
\end{aligned}$$

Here $\hat{A}_\varphi(\mathbf{u}) := \sum_{\mathbf{r}\in\mathcal{L}^\circ} |\hat{\varphi}(\mathbf{u}-\mathbf{r})|^2$ is the discrete-time Fourier transform of $a_\varphi[\mathbf{k}]$: the autocorrelation of φ with respect to the primal lattice $\mathcal{L}$. Recall that, for a symmetric kernel φ, $a_\varphi[\mathbf{k}] = \int_{\mathbb{R}^3} \varphi(\mathbf{x})\varphi(\mathbf{x}-\mathbf{L}\mathbf{k})d\mathbf{x}$.

Similar error kernels have appeared in the signal processing literature before [BU99, BTU01, CVDVB05], but their use in the graphics/visualization literature is rather sparse [AMC10]. For the sake of completeness, we have provided an intuitive alternate derivation. Note that our post-aliasing error kernel $E(\mathbf{u})$ is asymptotically equivalent to the minimum-error kernel of Blu and Unser [BU99]. It is a powerful tool as it can be evaluated in closed form provided that the Fourier transform $\hat{\varphi}$ and the autocorrelation spectrum $\hat{A}_\varphi$ are known. We can use it to compare the quality of different lattice-generator combinations and to optimize parametric reconstruction kernels such as C_λ.

4. CC to BCC Downsampling

In order to derive a high quality CC → BCC downsampling filter, we first optimize the kernel C_λ and then orthogonally project the fine scale trilinear reconstruction to the coarse scale CWLB reconstruction space.

Optimized CWLB Approximation on BCC Since the CWLB kernel C_λ is parametric, we seek a value of λ that minimizes the integral $\int_{\mathbb{V}_{\mathcal{B}_0^\circ}} E_{\mathsf{bcc}}(\mathbf{u})d\mathbf{u}$. Note that the integral is over the Voronoi cell of the normalized dual lattice $\mathcal{B}_0^\circ$ (a rhombic dodecahedron), and $E_{\mathsf{bcc}}(\mathbf{u})$ denotes the normalized post-aliasing error kernel. In order to simplify the optimization problem, we take advantage of the fact that C_λ is a second-order kernel. In other words $E_{\mathsf{bcc}}(\mathbf{u}) = O(\|\mathbf{u}\|^4)$ as $\mathbf{u} \to 0$ [BU99]. Consequently, the dominant contribution is due to the fourth-order Taylor series coefficients of E_{bcc} and we may as well replace E_{bcc} by its fourth-order coefficients. Using a computer algebra system (see supplementary material), we evaluated E_{bcc} is closed-form, determined its fourth-order Taylor developments in terms of λ and evaluated the integral analytically over the rhombic dodecahedron. This yields a quadratic polynomial in λ that can be easily minimized. The optimal value of λ thus obtained is

$$\lambda_o = -\frac{864(\pi^2-9)}{27\pi^2(\pi^2+2)-416} \approx -0.273512. \quad (8)$$

Note that the optimal value is negative; this case was overlooked by Csébfalvi [Csé13] even though they acknowledged that any non-zero value of λ leads to a valid kernel.

Comparison of Trilinear and CWLB Kernels The question of how trilinear interpolation on CC compares with *optimized* CWLB approximation on BCC can also be answered by investigating the normalized post-aliasing error kernels E_{cc} and E_{bcc}. We fix the sampling rate on the CC lattice at unity and look for a sampling rate h that yields the same error, i.e. $\int_{[-1/2,1/2]^3} E_{\mathsf{cc}}(\mathbf{u})d\mathbf{u} = \int_{\mathbb{V}_{\mathcal{B}_0^\circ}} E_{\mathsf{bcc}}(h\mathbf{u})d\mathbf{u}$. Equivalently, at what sampling rate h does the optimized CWLB approximation on the BCC lattice — generated by the matrix $\sqrt[3]{4}h\mathbf{B}$ — yield a similar quality as compared to trilinear interpolation on a unit density CC lattice? The answer is obtained via numerical integration (see supplementary material) and found to be $h = 1.61908$. Note that this is remarkably close to $\sqrt[3]{4} \approx 1.5874$, i.e. a BCC lattice which is 4 times less dense should provide a similar quality.

Downsampling Given sample values on a CC lattice, a straightforward CC→BCC downsampling strategy is to simply retain the BCC points, and use prefilterd CWLB interpolation [Csé13] (corresponding to C_{λ_o}) on the resulting BCC sample values. However, since our error bound is somewhat loose, we compensate for it by seeking an error-minimizing downsampling strategy that is inspired by the work of Hossain *et al.* [HAM11]. Given a fine-scale CC reconstruction $f_1(\mathbf{x})$ (see (4)), a minimum-error approximation on the coarse-scale BCC lattice $\mathcal{B}$ is obtained by orthogonally projecting f_1 to the target reconstruction space. This is tantamount to computing the coarse-scale BCC coefficients in (6) according to [Uns00]

$$c[\mathbf{k}] = \int_{\mathbb{R}^3} f_1(\mathbf{x}) C_\lambda^\circ(\mathbf{x}-\mathbf{B}\mathbf{k})d\mathbf{x}, \quad (9)$$

where C_λ° denotes the biorthogonal dual of C_λ. Observe that f_1 lies in the trilinear reconstruction space, and $C_\lambda^\circ = \sum_{\mathbf{k}} a_{C_\lambda}^{-1}[\mathbf{k}] C_\lambda(\mathbf{x}-\mathbf{B}\mathbf{k})$, where $a_{C_\lambda}^{-1}[\cdot]$ is the inverse autocorrelation sequence. Therefore, the coarse-scale coeffi-

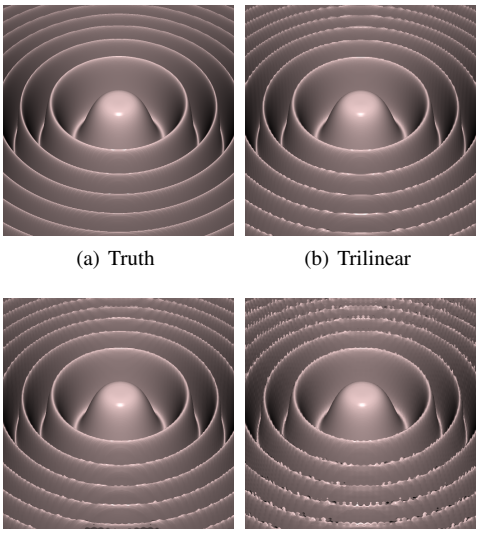

(a) Truth (b) Trilinear

(c) CWLB ($\lambda = \lambda_o$) (d) CWLB ($\lambda = 1$)

Figure 1: *Isosurface renderings of the ML test function: (a) ground truth, (b) sampled on a* 128^3 *CC grid, and (c–d) CC representation downsampled on* $64^3 \times 2$ *BCC grids.*

cients can also be obtained via the following two-step filtering operation:

1. Convolve the CC samples $f[\cdot]$ with the CC filter $\sigma[\mathbf{k}] := \int_{\mathbb{R}^3} L(\mathbf{x}) C_\lambda(\mathbf{x}-\mathbf{k}) d\mathbf{x}$. Subsample the resulting sequence to BCC to yield the intermediate BCC sequence $r[\mathbf{k}] := (f *_{\mathsf{cc}} \sigma)[\mathbf{Bk}]$.
2. Convolve the intermediate BCC sequence with the inverse auto-correlation sequence to yield the result, i.e. $c[\mathbf{k}] = (r *_{\mathsf{bcc}} a_{C_\lambda^{-1}})[\mathbf{k}]$.

It should be emphasized that the convolution in the first step is on the CC lattice, whereas that in the second step is on the BCC lattice. In practice, these convolutions can be performed in the Fourier domain by using the multidimensional fast Fourier transform (FFT) on CC and the BCC-FFT [AM09] on the BCC lattice.

5. Results and Discussion

In order to validate our downsampling scheme, we point sampled the popular test function of Marschner and Lobb (ML) [ML94] on a 128^3 CC grid. The calibration procedure proposed by Vad *et al.* [VCG12] is not applied since the ML function is well-above the Nyquist rate at this resolution. Using our downsampling scheme, we obtained $64^3 \times 2$ BCC representations that are suitable for both interpolating ($\lambda = 1$) and optimized ($\lambda = \lambda_o$) CWLB reconstruction. For all cases, we rendered the isosurface corresponding to an isovalue of 0.5. In order to approximate the gradient for shading, we used the revitalization procedure of Alim *et. al* [AMC10] to estimate the gradient using derivative filters. On the BCC lattice, this procedure was applied to each CC coset separately, and the results were combined using the weighting function W_λ.

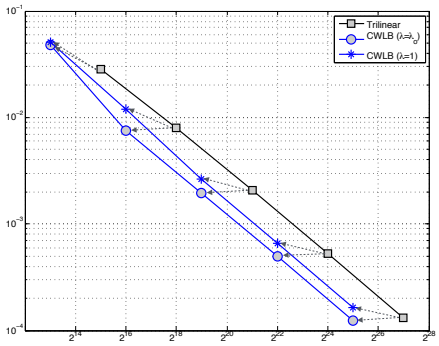

Figure 2: *RMS error vs. the total number of grid points.*

Fig. 1 shows our results as compared to a ground truth rendition obtained using the synthetic function itself. Even though the data is oversampled, trilinear interpolation on CC shows some post-aliasing. As expected, the optimized downsampled BCC rendition ($\lambda = \lambda_o$) is very close to the fine-scale CC rendition and remarkably different from the unoptimized downsampled rendition ($\lambda = 1$). Fig. 2 compares the root mean square (RMS) trilinear reconstruction error (measured using 10^6 points randomly distributed within the cube $[-0.5, 0.5]^3$) for the ML function as the CC grid resolution is increased from 32^3 to 512^3. The corresponding downsampled CWLB reconstruction errors (for the cases $\lambda = \lambda_o$ and $\lambda = 1$) are also shown. Observe that, when the data is sampled below the Nyquist limit (32^3), downsampled CWLB reconstruction is worse as compared to trilinear interpolation. On the other hand, when the data is oversampled, downsampled CWLB interpolation ($\lambda = 1$) is only slightly worse than trilinear interpolation but, as predicted, downsampled CWLB approximation ($\lambda = \lambda_o$) is actually slightly better even though the data is four times coarser.

6. Conclusion

We optimized the CWLB kernel on the BCC lattice using a post-aliasing error kernel and derived a CC→BBC downsampling scheme that leads to a 75% data reduction when the original CC data is sufficiently oversampled. Our downsampling filter is infinitely supported and needs to access the entire volume. However, it is straightforward to derive compactly supported approximations that would be more suitable for on-the-fly downsampling for streaming data. Furthermore, since CWLB reconstruction makes use of trilinear interpolation on the constituent CC cosets, it is also possible to take advantage of existing CC bricking or compression schemes to handle large volumetric datasets. Thus, our proposed downsampling scheme can potentially lead to a significant data reduction with minimal loss of quality.

Acknowledgements

Support for this work was provided by the Faculty of Science, University of Calgary, and the Science without Borders (SwB) program.

References

[AM09] ALIM U. R., MÖLLER T.: A fast Fourier transform with rectangular output on the BCC and FCC lattices. In *Proceedings of the Eighth International Conference on Sampling Theory and Applications (SampTA'09)* (Marseille, France, May 2009). 4

[AMC10] ALIM U. R., MÖLLER T., CONDAT L.: Gradient estimation revitalized. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization 2010) 16*, 6 (Nov. 2010), 1494–1503. 3, 4

[BTU01] BLU T., THÉVENAZ P., UNSER M.: MOMS: Maximal-order interpolation of minimal support. *IEEE Transactions on Image Processing 10*, 7 (2001), 1069–1080. 3

[BU99] BLU T., UNSER M.: Quantitative Fourier analysis of approximation techniques: Part I—Interpolators and projectors. *IEEE Transactions on Signal Processing 47*, 10 (Oct. 1999), 2783–2795. 3

[Chu88] CHUI C.: *Multivariate Splines*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 1988. 2

[Csé10] CSÉBFALVI B.: An evaluation of prefiltered B-spline reconstruction for quasi-interpolation on the body-centered cubic lattice. *IEEE Transactions on Visualization and Computer Graphics 16*, 3 (May 2010), 499–512. 2

[Csé13] CSÉBFALVI B.: Cosine-weighted b-spline interpolation: A fast and high-quality reconstruction scheme for the body-centered cubic lattice. *IEEE Transactions on Visualization and Computer Graphics 19*, 9 (Sept 2013), 1455–1466. 1, 2, 3

[CVDVB05] CONDAT L., VAN DE VILLE D., BLU T.: Hexagonal versus orthogonal lattices: A new comparison using approximation theory. In *Image Processing, 2005. ICIP 2005. IEEE International Conference on* (2005), vol. 3, IEEE, pp. III–1116. 3

[DC10] DOMONKOS B., CSÉBFALVI B.: DC-splines: Revisiting the trilinear interpolation on the body-centered cubic lattice. In *Proceedings of Vision, Modeling, and Visualization 2010* (Nov. 2010), pp. 275–282. 2

[dHR93] DE BOOR C., HÖLLIG K., RIEMENSCHNEIDER S.: *Box Splines*. Springer Verlag, 1993. 2

[DM84] DUDGEON D. E., MERSEREAU R. M.: *Multidimensional Digital Signal Processing*, 1st ed. Prentice-Hall, Inc., Englewood-Cliffs, NJ, 1984. 2

[EDM04] ENTEZARI A., DYER R., MÖLLER T.: Linear and cubic box splines for the body centered cubic lattice. In *Proceedings of the IEEE Conference on Visualization* (Oct. 2004), pp. 11–18. 2

[EMBM06] ENTEZARI A., MENG T., BERGNER S., MÖLLER T.: A granular three dimensional multiresolution transform. In *Proceedings of the Eurographics/IEEE-VGTC Symposium on Visualization* (May 2006), pp. 267–274. 2

[EMK09] ENTEZARI A., MIRZARGAR M., KALANTARI L.: Quasi-interpolation on the body centered cubic lattice. *Computer Graphics Forum 28*, 3 (2009), 1015 – 1022. 2

[EVM08] ENTEZARI A., VAN DE VILLE D., MÖLLER T.: Practical box splines for reconstruction on the body centered cubic lattice. *IEEE Transactions on Visualization and Computer Graphics 14*, 2 (2008), 313 – 328. 2

[FEVM10] FINKBEINER B., ENTEZARI A., VAN DE VILLE D., MÖLLER T.: Efficient volume rendering on the body centered cubic lattice using box splines. *Computers & Graphics 34*, 4 (2010), 409–423. 2

[HAM11] HOSSAIN Z., ALIM U. R., MÖLLER T.: Toward high quality gradient estimation on regular lattices. *IEEE Transactions on Visualization and Computer Graphics 17*, 4 (Apr. 2011), 426–439. 3

[KEP08] KIM M., ENTEZARI A., PETERS J.: Box spline reconstruction on the face-centered cubic lattice. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (2008), 1523–1530. 2

[KP10] KIM M., PETERS J.: Symmetric box-splines on the $\mathcal{A}_n^*$ lattice. *Journal of Approximation Theory 162*, 9 (2010), 1607–1630. 2

[KP11] KIM M., PETERS J.: Symmetric box-splines on root lattices. *Journal of Computational and Applied Mathematics 235*, 14 (2011), 3972–3989. 2

[ME10] MIRZARGAR M., ENTEZARI A.: Voronoi splines. *IEEE Transactions on Signal Processing 58*, 9 (2010), 4572–4582. 1, 2

[ME11] MIRZARGAR M., ENTEZARI A.: Quasi interpolation with Voronoi splines. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 1832–1841. 2

[ML94] MARSCHNER S. R., LOBB R. J.: An evaluation of reconstruction filters for volume rendering. In *Proceedings of the IEEE Conference on Visualization* (Oct. 1994), pp. 100–107. 4

[TMG01] THEUSSL T., MÖLLER T., GRÖLLER E.: Optimal regular volume sampling. In *Proceedings of the IEEE Conference on Visualization 2001* (Oct. 2001), pp. 91–98. 1

[UAE93] UNSER M., ALDROUBI A., EDEN M.: B-Spline signal processing: Part I—Theory. *IEEE Transactions on Signal Processing 41*, 2 (Feb. 1993), 821–833. 2

[Uns00] UNSER M.: Sampling-50 years after Shannon. *Proceedings of the IEEE 88*, 4 (2000), 569–587. 3

[VCG12] VAD V., CSÉBFALVI B., GABBOUJ M.: Calibration of the marschner-lobb signal on cc, bcc, and fcc lattices. In *Proceedings of EuroVis* (2012), pp. 19–23. 4

[VCRG14] VAD V., CSÉBFALVI B., RAUTEK P., GRÖLLER E.: Towards an unbiased comparison of cc, bcc, and fcc lattices in terms of prealiasing. *Computer Graphics Forum 33*, 3 (2014), 81–90. 1

[Wan01] WANG R.: *Multivariate Spline Functions and their Applications*. Mathematics and its Applications. Kluwer Academic Publishers, 2001. 2

[YE12] YE W., ENTEZARI A.: A geometric construction of multivariate sinc functions. *IEEE Transactions on Image Processing 21*, 6 (2012), 2969–2979. 1