Volume Graphics (2007)
H. - C. Hege, R. Machiraju (Editors)

Applications of Optimal Sampling Lattices for Volume Acquisition via 3D Computed Tomography

Fang Xu Klaus Mueller

Center of Visual Computing, Computer Science, Stony Brook University

Abstract

There has been mounting evidence that the familiar Cartesian lattices, while convenient for signal processing and representation, are sub-optimal when it comes to signal fidelity. More suitable in this respect are optimal sampling lattices, such as the Hexagonal and Body Centered Cartesian (BCC) lattice, and recent work has employed these in the areas of volume rendering and image processing. In this paper we explore various applications of these lattices within the context of 3D Computed Tomographic Reconstruction, both in terms of the (2D) detector and the (3D) reconstructed object, and a theoretical analysis is provided. We combine this analysis with a practical application, that is, the use of these lattices within a real-time GPU-accelerated 3D reconstruction platform, in which performance is also of an immediate concern.

Categories and Subject Descriptors (according to ACM CSS): I.3.3 [Computer Graphics]: Display Algorithms.

1. Introduction

Regular lattices typically are Cartesian lattices, with grid samples distributed on a separable, orthogonal raster of most often equal grid spacing in all dimensions. This type of lattice is very convenient for representation, indexing, and interpolation, and it is also easy to conceptualize. It is mostly for these reasons that the regular Cartesian lattice has become the most dominantly used regular raster structure today. But recent years have brought an increased awareness with respect to the sub-optimality of this grid topology. Early efforts in this direction have mostly concentrated on efficiency, and the key observation motivating this is made in the frequency domain. When assuming a radially symmetric (spherical) frequency spectrum of the rasterized signal, then the optimal packing of the alias spectra is not a Cartesian lattice but a hexagonal one, since such a lattice packs the frequency spectrum spheres closest together (which directly follows from optimal sphere packing theory [CS99]). This, in turn, using the Fourier scaling theorem, stretches the samples in the co-domain, here the spatial domain, furthest apart, leading to the coarsest possible sampling pattern without risking (pre-) aliasing. This makes possible a reduction of grid points to 87% in 2D, 71% in 3D, and 50% in 4D, with a direct consequence being a reduction in storage by these amounts, which can affect cache behavior as well. But possibly more importantly, grid processing costs are reduced by these amounts as well, if these costs are strongly related to the number of grid points. This is the case whenever a point-based projection approach is used. We call such a scheme *voxel-driven*, in contrast to *pixel-driven* schemes in which rays originating from projection pixels traverse and interpolate the lattice.

Computed Tomography (CT) seeks to reconstruct an unknown object discretized into the reconstruction lattice from its X-ray projections. For this, CT reconstruction methods typically use a voxel-driven approach in which grid points interpolate the projection image and accumulate these contributions over all available projections. This procedure is called *back-projection*. Thus savings from optimizing the number of grid points immediately apply, and this was recognized for iterative CT reconstruction with the Simultaneous Algebraic Reconstruction Technique (SART) quite early by [ML95]. Similar observations have also been made later for the inverse case of CT, that is, volume projection and rendering [NM02][TMG01]. In fact, iterative CT consists of alternate projections and back-projections, where the estimates of the volume under reconstruction are iteratively projected onto the acquired images, the error is computed, and a normalized error is back-projected into

the lattice to bring it closer to the solution. In this paper, we restrict our discussion to analytical reconstruction methods, which are the most popular and consist only of one round of back-projection operations.

It turns out that cost savings are only one aspect of optimal lattices. They also provide a more uniform and isotropic sampling of the space [SBB03], under the condition of a tighter space sampling than that implied by the sphere-packing results. Our paper shows the important implications this has for CT, and in fact, these aspects come into play in CT on two occasions: (1) the initial data acquisition on the detector lattice, capturing the object-attenuated transmission X-ray radiation, and (2) the object formation on the reconstruction lattice accumulating the back-projected contributions. In the subsequent sections, we will give theoretical arguments as well as show practical examples for both.

But apart from reconstruction fidelity, computational performance is also an important issue. This is not only determined by the number of grid points, but also by the efficiency a back-projection can be produced using the lattice topology chosen. The GPU has become a major acceleration medium, at commodity pricing, and we have shown in other work [XM07] that GPU acceleration can enable a pipeline in which projection images are being back-projected at the same speed as they are acquired by a rotating flat-panel gantry. In that work it was also shown that hardware-accelerated bilinear interpolation is a major source of speedup.

Our paper is structured as follows. Section 2 presents relevant prior work. Then Section 3 gives a theoretical exposure on lattices and CT, in light of the arguments presented in this paper, while Section 4 discusses their practical implications on detector and reconstruction grid design. Finally, Section 5 gives experimental results, and Section 6 ends with conclusions.

2. Prior relevant work

As mentioned in the introduction, the optimal lattice in 2D is the hexagonal lattice (illustrated in Figure 1a) – it allows the sparsest sampling without a loss of frequency content. Note, however, that this assumes frequency spectra that fit into a radially symmetric hull (circle, sphere) with the radius given by the frequency bandwidth. This is not always the case, but as we shall see later, it is a given in CT.

The 3D equivalent in terms of space-optimality is the Body Centered Cartesian (BCC) lattice (illustrated in Figure 1b – in all of Figure 1, the grid distances have been expressed in terms of its frequency bandwidth-equivalent Cartesian (CC) lattice). The frequency transform of the hexagonal lattice is another hexagonal lattice. The frequency transform of the BCC lattice is the Face-Centered Cartesian (FCC) lattice, which has the tightest best sphere packing possible.

As mentioned above, optimal sampling lattices have been used for quite some time for discrete object representation, both for 3D [TMG01] and 4D volume datasets [NM02]. These works were mainly motivated by the reduction in the number of required lattice samples. Since the generation of projective, volume-rendered images requires the interpolation of lattice samples, now in this optimal lattice space, appropriate filters for this operation have been proposed, such as box-splines [EDM04], hex-splines [VNP*04], and a pre-filtering scheme that operates in conjunction with a Gaussian filter [Cse05]. Finally, the theoretical finding of space-optimality without quality loss was confirmed in a user study which compared images rendered from a Cartesian lattice and images rendered from the same object rendered from an optimal lattice [MSE*07].

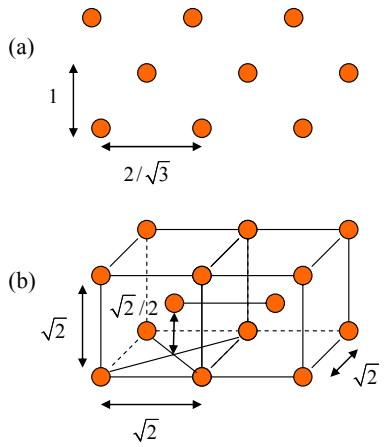

Figure 1: (a) hexagonal lattice, (b) BCC lattice

There were also researchers in the CT reconstruction arena who have used optimal lattices [ML95][MY96], and these works were all within iterative reconstruction frameworks. Just as in volume rendering, these efforts were mainly driven by the reduction of grid samples and the subsequent speedup in terms of reconstruction time. Using optimal lattices brought significant savings here since iterative algorithms typically project and back-project the evolving reconstruction many times, and thus a reduction in grid complexity can make a considerable difference in running time.

The work described in this paper looks at CT reconstruction and optimal lattices from a different perspective, that is, from the standpoint of lattice isotropy and uniformity. We will show that due to these intrinsic properties, optimal lattices can reconstruct and acquire fine detail significantly better than the standard Cartesian lattices, without a loss of performance. This is important, for example, when it comes to the detection of small lesions in CT. At the same time, optimal lattices can also provide a higher-fidelity sampling of the incoming X-ray signal on the detector plane, and this will explored in this paper as well.

3. Theory

In the following, we first give some insight with respect to the isotropy and uniformity of optimal lattices. Then we discuss the implications on CT.

3.1 Lattice uniformity

In a uniform lattice the maximum distance of an arbitrary sample in space to a lattice point is independent of direction. Thus, the Voronoi cell of such a uniform lattice must be a sphere. A collection of spheres, however, cannot be space-filling and thus optimal uniformity cannot be achieved in practice. We therefore seek a lattice with a Voronoi cell that is closest to a sphere. The sphere has the smallest surface area enclosing a given volume. Setting the volume of a sphere to unity, the surface area S_s is:

$$S_s = 4\pi(\frac{1}{\sqrt[3]{4/3\cdot\pi}})^2 = 4.83$$

The surface area S_{cc} of a unit cube is an obvious S_{cc}=6. The Voronoi cell of the BCC lattice is the truncated octrahedron, and the surface area S_{bcc} of its unit cell is:

$$S_{bcc} = (6+12\sqrt{3})\cdot(\frac{1}{\sqrt[3]{8\sqrt{2}}})^2 = 5.31$$

Here the second term is the lattice parameter value setting for a unit cell. We see that the BCC lattice is about 10% worse than the sphere, but 12% better than the CC lattice. Finally, let us have a look at the Face-Centered Cartesian (FCC) lattice, which is the dual of the BCC lattice. Its Voronoi cell is the rhombic dodecahedron, and the surface area S_{fcc} of its unit cell is:

$$S_{fcc} = 8\sqrt{2}\cdot(\frac{1}{\sqrt[3]{(16/9)\sqrt{3}}})^2 = 5.34$$

Again, the second term is the lattice parameter value setting for a unit volume cell. Thus, the BCC lattice is slightly more isotropic than the FCC lattice, under this metric. We therefore choose the BCC lattice.

3.2 Computed tomography reconstruction

Let us consider the 2D case first, and assume we have collected a collection of M X-ray projections $p(k,\theta_m)$, each N samples wide, around the object, where k is the sample index on the projection acquired at angle θ_m. In order to reconstruct a function $f(x,y)$ from these projection, the following *filtered back-projection* formula is employed:

$$f(x,y) = \sum_{m=1}^{M}\sum_{k=-N/2}^{N/2} P(k,\theta_m)\cdot|k|\cdot e^{i2\pi k(x\cos(\theta_m)+y\sin(\theta_m))}$$

where $0\le\theta_m<\pi$. This equation indicates that all projections $p(k,\theta_m)$ are first Fourier transformed to yield $P(k,\theta_m)$. These are then ramp filtered and then inverse Fourier transformed. The last sum indicates the back-projection operation. Typically, this back-projection is accomplished by a mapping of the voxels to each filtered projection $p_f(k,\theta_m)$, where they interpolate the image using a bilinear filter.

Each projection essentially constitutes a radial slice of the object's Fourier spectrum (this is the *Fourier Slice Theorem*). As shown in Fig. 2, assuming an equi-angular projection acquisition, we obtain a uniform polar sampling (the outer areas are less tightly sampled as a linear function of frequency, and this is compensated for by the ramp filtering with factor $|k|$).

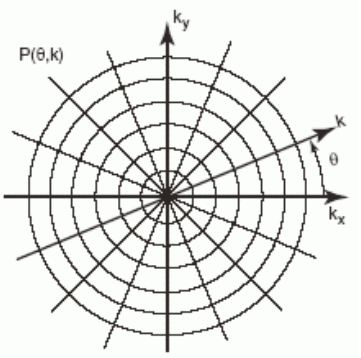

Figure 2: Fourier spectrum created by the set of acquired X-ray projections.

The polar grid has a number of important implications with regard to the sampling lattices used in reconstruction and acquisition. First, we see that the Fourier Slice Theorem results in the acquisition of a radially symmetric frequency spectrum. Therefore, an optimal sampling lattice, which assumes this type of spectrum (as noted before) is appropriate.

The next comment concerns the resolution itself. The resolution in the radial direction is given by the resolution of the projections, N. The maximum frequency k_{max} is then directly determined by N, that is, the spacing Δk in the frequency domain is:

$$\Delta k = \frac{k_{\max}}{N/2}$$

Thus, the there is no need to reconstruct onto a grid with more samples than N, since the maximal frequency and size are limited by this number. Also, we seek to acquire enough projections such that $\Delta k \ge \Delta\theta$, which yields the number of required projection $M=\pi N/2$.

What we have discussed just now was the 2D case, but we aim to reconstruct in 3D. This yields a sphere only in Radon space, which is the space of plane integrals [Gra91]. In the limit, for parallel-beam projection, the

frequency space is a stack of radially symmetric spectra, which only partially fulfills the optimal lattice assumption. We are using cone-beam projection, however (see Figure 3 and the descriptions in Section 4.3), which conforms somewhat better to the sphere.

4. Implementation

In this section we discuss the various CT reconstruction components affected by this work in closer detail, and describe our modifications.

4.2 Data acquisition

Detectors are typically composed of an array of square pixels. There has been a recent trend to CMOS flat-panel detectors, and parallel to this, optical video cameras have also been introduced that adopted hexagonal lattices for more isotropic sampling. This has led us to explore these types of lattices also for X-ray detection. However, GPU-accelerated reconstruction (see Section 4.3) requires a standard Cartesian grid for the mapping of the voxels onto the projections. Else, the added overhead incurred by application of fragment shader-bound interpolation filters would cause a significant loss in performance, as this would occur at a complexity of $O(N^4)$, assuming we have $O(N)$ projections to reconstruct N^3 voxels. We therefore resample the acquired projections from the hexagonal grid onto a Cartesian grid before reconstruction begins. This enables fast hardware accelerated bilinear interpolation, and it also lowers the complexity of the hex-grid interpolation by an order of magnitude (now it is a pre-processing step). In fact, we couple this operation with a projection of the raw projection images into a standard axis-parallel configuration which allows for a faster voxel mapping in the back-projection stage [KKB06]. This yields projection images of twice the resolution which can be back-projected using nearest neighbor interpolation. However, since in the latest GPU hardware bilinear interpolation is only minutely less efficient as nearest-neighbor interpolation, we use the former.

4.3 Volume reconstruction

Our reconstruction framework uses the standard Filtered Backprojection cone-beam algorithm devised by Feldkamp, Davis, and Kress (we are only discussing the essential steps here, for detailed information please consult [FDK84]). The Feldkamp (FDK) algorithm consists of three stages: projection-space filtering, back-projection, and volume-space weighting. Figure 3 illustrates the cone-beam projection geometry and the notations used for the following discussion. A voxel to be reconstructed is denoted by a vector $\boldsymbol{r} = (x, y, z)$ in the (reconstruction) volume space defined by (X_v, Y_v, Z_v), with Y_v being the rotation axis. The elements on the projection (detector) plane oriented at angle ϕ are represented by $P_\phi(Y, Z)$, where Y and Z represent an element's spatial location in detector coordinates. The vector X_ϕ is orthogonal to the detector plane and connects the detector center with the source S. The two orthogonal vectors Y_ϕ and Z_ϕ complete the 3D coordinate system of the detector space, which is related to (X_v, Y_v, Z_v) by ways of a transformation matrix composed of gantry rotation and possible gantry warp. Finally, the distances from the source to the rotation center O and the detector center are defined as d and D, respectively.

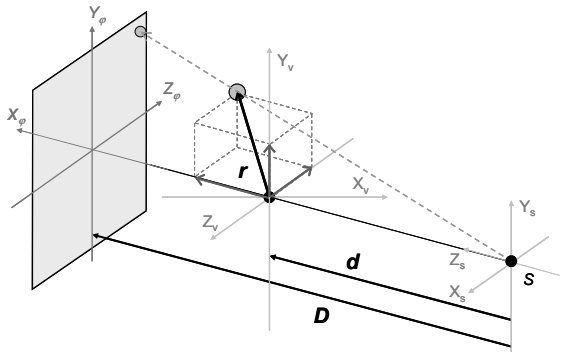

Figure 3: Cone-beam projection geometry

Let us assume we have M projections P_{ϕ_m} acquired over a range $0 \le \varphi_m < 2\pi$. The first step is the filtering of all projections in projection space:

$$\hat{P}_{\phi_m}(Y,Z) = \frac{D}{\sqrt{D^2+Y^2+Z^2}} P_{\phi_m}(Y,Z) ** g(Z)$$

where $g()$ is the spatial equivalent of the 1D frequency-space ramp filter mentioned in Section 3.2, and the first term represents an isotropic weighting about the detector origin.

The back-projection operation requires a mapping of each voxel onto the detector plane:

$$Y(\boldsymbol{r}) = \frac{\boldsymbol{r}\cdot y_\phi}{d+\boldsymbol{r}\cdot x_\phi} D, \quad Z(r) = \frac{\boldsymbol{r}\cdot z_\phi}{d+\boldsymbol{r}\cdot x_\phi} D$$

The back-projected value $\tilde{P}_{\phi_m}$ is then obtained by interpolating the samples on the filtered projection image $\hat{P}_{\phi_m}$. This value is then depth-weighed, and the contributions from all projections are accumulated into the reconstructed voxel:

$$f(\boldsymbol{r}) = \frac{1}{4\pi^2} \sum_{m=0}^{M} \frac{d^2}{(d+\boldsymbol{r}\cdot x_{\phi_m})^2} \tilde{P}_{\phi_m}(Y(\boldsymbol{r}), Z(\boldsymbol{r}))$$

The (depth-) weighting term is inversely proportional to the distance of the voxel's projection onto the source-detector axis and the source. These three equations then constitute a natural decomposition of the FDK algorithm,

forming a pipeline consisting of filtering, back-projection, and depth-weighting.

Our GPU-accelerated reconstruction framework (please refer to [XM07] for detailed description) uses a single PC equipped with a single high-end commodity graphics board (the NVIDIA 8800 GTX) to process clinically-sized projection data at speeds meeting and exceeding the typical flat-panel detector data production rates, enabling throughput rates of 40–50 floating-point precision projections/s for the reconstruction of 512^3 volumes. Due to this steady-flow pipeline we call this framework *streaming CT*. The back-projection operation on the GPU groups the voxels into volume slices and processes these together in one pass. It applies the perspective cone-beam transform to the slice (polygon) vertices, and a fragment program then uses the texture coordinates forwarded by the rasterizer to interpolate the projection images, using bilinear interpolation.

Essential for the work described here is the fact that the back-projection operations are independent of the underlying lattice. Lattice points (voxels) (i) are back-projected (mapped to a projection image), (ii) then interpolate their updates from this projection image, and (iii) finally receive their depth-weighting, all according to their individual coordinates. It is the topology of the lattice that decides if the signal so reconstructed will suffer from pre-aliasing. This is important for the reconstruction scheme described in this paper, which uses optimal lattices. As we have shown in Section 3.2, the main requirement of optimal lattices – a radially symmetric frequency spectrum – is (at least approximately) fulfilled for cone-beam CT. Thus, as long as we do not commit any aliasing in the mapping operations, the reconstructed signal will not be contaminated by pre-aliasing (disregarding aliasing stemming from the interpolation filters themselves).

The GPU-accelerated reconstruction algorithm itself only requires a slight change for handing the BCC lattice. Since for BCC each slice is a standard Cartesian lattice, we can employ our interpolation fragment program unchanged. The only item requiring change is the vertex program that needs to shift the slice polygon for each odd slice index.

We have chosen to keep the total number of voxels the same. This reduces the in-slice lattice spacing from the space-optimal $\sqrt{2}$ to $\sqrt[3]{2}$ = 1.26 (and an inter-slice distance of $\sqrt[3]{2}/2$), yielding a system where the reconstruction performance does not suffer at all. However, as we see next, this arrangement is able to resolve small feature better than the standard lattice with both in-slice and inter-slice spacing of 1.

5. Experimental Results

We have conducted experiments on both optimal detector lattices and reconstruction lattices. The results are reported next.

5.1 Detector lattice

We explored three configurations for the detector lattices, up-sampled these acquired images, and compared them with a directly acquired (via simulation) X-ray image at the same grid resolution. All simulation traced multiple (16) rays per pixel to simulate X-ray beams.

Since the detector elements in the optimal lattices are regular hexagons, let us first review some (regular) hexagon geometry metrics. The area of a hexagon $A_{hex} = 3\sqrt{3}/2 \cdot a^2$, where a is the side length of the hexagon, and the maximal diameter is $2a$ and the minimal diameter is $a\sqrt{3}$. The minimal diameter is the horizontal in-slice lattice spacing, and the vertical spacing is $\sqrt{3}/2$ of this. With this in mind, we explored the following configurations:

1. The standard Cartesian case, assuming a lattice spacing of 1, and a (square) detector element area of 1.
2. A standard Cartesian lattice with twice the resolution of #1, which gives a detector element area of 0.25.
3. The space-optimal (hexagonal) lattice version of #2, which has 86% ($\sqrt{3}/2$) of the elements of #2. Here, the detector element area is increased by 23% to 0.28.

All images were up-sampled to twice the resolution of the second configuration. We used a linear filter for all cases. For the Cartesian grids we used bilinear interpolation, while for the hexagonal grid we employed a filter based on barycentric coordinates. We performed our study using the Marschner-Lobb (ML) dataset [ML94] (see Figure 4), which has an almost uniform frequency content all the way up to the Nyquist limit. Measurements were taken for an angular range of 0-45° at 10° increments, and the projections were compared with the simulated projection obtained at the same angle in a numerical sense, via its RMS error. The results of this study are shown in Figure 5. There we observe that the space-optimal hexagonal detector lattice performs significantly better than the equivalent standard Cartesian lattice, at a greater detector element area. Thus, the hexagonal detector combines the better area efficiency, which has favorable implications with respect to the signal-to-noise ratio, with better sampling fidelity.

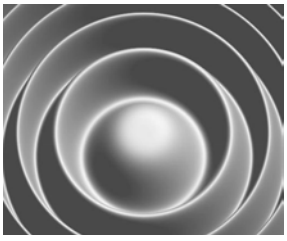

Figure 4: A volume rendering of the ML dataset

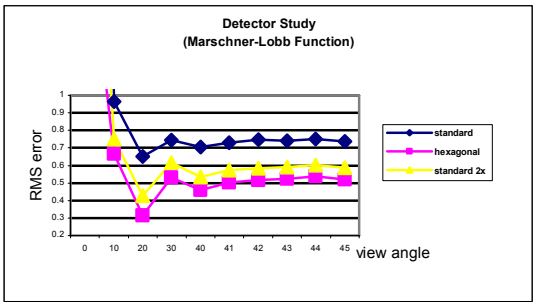

Figure 5: Projection results for the ML function, for various detector lattice configurations.

5.2 Reconstruction Lattice

We then used the up-sampled detector images for 3D reconstruction. We first reconstructed the Toes dataset, both onto a standard Cartesian lattice and then on the BCC lattice with the same number of elements, using our streaming CT reconstruction application. Fig 6 shows visualizations from the same viewpoint and at the same iso-surface setting. We used direct volume rendering (DVR) with a few compositing steps for iso-surface visualization (we did not extract a polygonal mesh like [CTM03]). We notice a slightly better recovery of the features with DVR and an overall smoother surface quality. But the differences are not overly dramatic.

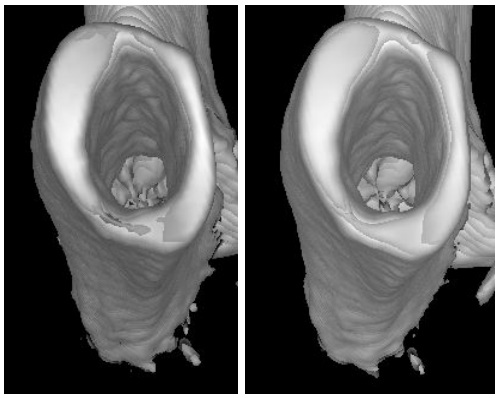

Figure 6: Iso-surface rendering of the Toes dataset: (left) CC-lattice reconstruction, (right) BCC-lattice reconstruction.

The true advantages of optimal lattices lie in their ability to recover small features better, and with less sensitivity to orientation (due to its more isotropic sampling). A common task in medical imaging is the detection of small lesions and tumors. To explore the performance of the two candidate lattices in this context, we generated a 3D phantom dataset consisting of blobs of various sizes and generated (standard Cartesian) projections of these. The projections had a resolution such that the smallest blobs would still be detectable to at least 1-2 pixels. We then reconstructed these blobs on both a CC and the equivalent BCC lattice at a resolution matching that of the projections. At this lattice resolution the size of the smallst blobs amounted to about 1.5 voxels. The reconstruction results are shown in Fig. 7. We see that only the BCC lattice is able to recover the smaller blobs, and overall the blobs appear better refined.

Next, we generated a tumor brain phantom by embedding a selection of randomly distributed small tumors into an existing brain volume dataset. We generated projections and reconstructed them as well, using the same protocol than for the blob phantom. Fig. 8 shows a slice of this reconstruction. Again, we observe that the BCC lattice is able to recover almost all of the tumors, while the CC lattice fails in many cases.

6. Conclusions and Future Work

We have demonstrated that BCC lattices compare favorably with CC lattices for CT data acquisition and reconstruction. For reconstruction, we have focused on accuracy only, that is, we have used the same number of lattice points, but distributed them into the more isotropic configuration of the BCC lattice. We found that this results in somewhat minor overall quality improvements, at least for the dataset we have tested, But, more importantly, we have also found that the BCC lattice increases the recovery and detectability of small features, for example, small tumors in the brain. This is an important aspect in CT practice.

We have also investigated the use of optimal lattices on X-ray detectors, and our findings were that the better isotropy of hexagonal lattices enables the detector elements to be made larger, but yet obtain more accurate projections.

Future work on reconstruction lattices will focus on "smart" multi-resolution grids, which adapt their resolution to the projection images provided. This, in some ways, is a continuation of earlier work on D^2VR [RCG*06][XM06], where volume rendered images are generated directly from the projections and no volume lattice is ever generated.

Future work on hexagonal detector lattices should focus on further explorations of the tradeoff in lattice resolution and detector element size. More experiments will help to define the optimal configurations in that respect, given the task of the application. This goes in hand with more research on better interpolation filters for the subsequent hexagonal-Cartesian regridding.

Acknowledgments

This work was funded by NIH grant R21 EB004099-01.

References

[CTM03] H. Carr, T. Theußl, and T. Möller, "Isosurfaces on optimal regular samples," *Eurographics/TCVG Symposium on Data Visualization*, pp. 39-48, 2003.

[CS99] J. Conway and N. Sloane. *Sphere Packings, Lattices and Groups*. Springer, 3rd edition, 1999.

[Cse05] B. Csebfalvi, "Prefiltered Gaussian reconstruction for high-quality rendering of volumetric data sampled on a body-centered cubic grid," *IEEE Visualization*, pp. 40-48, 2005.

[EDM04] A. Entezari, R. Dyer, and T. Möller, "Linear and Cubic box splines for the body centered cubic lattice," *Proc. IEEE Visualization*, pp. 11–18, 2004.

[FDK84] L. A. Feldkamp, L. Davis and J. W. Kress, "Practical cone beam algorithm," *J. Optical Society of America,* pp. 612–619, 1984.

[Gra91] P. Grangeat, "Mathematical framework of cone-beam 3D reconstruction via the first derivative of the Radon transform," in *Mathematical Methods in Tomography*, Springer-Verlag, 1991, vol. 1497, Lecture Notes in Mathematics, pp. 66–97.

[KKB06] M. Kachelrieß, M. Knaup, and O. Bockenbach, "Hyperfast perspective cone–beam backprojection," *IEEE Medical Imaging Conference 2006*.

[ML94] S. Marschner, and R. Lobb, "An evaluation of reconstruction filters for volume rendering," *Proc. of IEEE Visualization*, pp. 100-107, 1994.

[ML95] S. Matej, R.M. Lewitt, Efficient 3D grids for image-reconstruction using spherically-symmetrical volume elements, *IEEE Trans. Nuclear Science*. 42:1361-1370, 1995.

[MSE*07] T. Meng, B. Smith, A. Entezari, A. E. Kirkpatrick, D.Weiskopf, L. Kalantari, and T. Möller, "On visual quality of optimal 3D sampling and reconstruction," *Graphics Interface,* pp. 265-272, 2007.

[MY96] K. Mueller and R. Yagel, "The use of dodecahedral grids to improve the efficiency of the Algebraic Reconstruction Technique (ART)," *Annals Biomedical Engineering, Special issue, 1996 Annual Conference of the Biomedical Engineering Society*, p. S-66, 1996.

[NM02] N. Neophytou and K. Mueller, "Space-time points: 4D Splatting on effcient grids," *ACM/IEEE Symposium on Volume Visualization and Graphics*, pp. 97-106, 2002.

[RCG*06] P. Rautek, B. Csébfalvi, S. Grimm, S. Bruckner, E. Gröller, "D^2VR: High-quality volume rendering of projection-based volumetric data," *Eurographics/TCVG Symposium on Visualization (EuroVis)*, pp. 211-219, 2006.

[SBB03] M. Stijnman, R. Bisseling, and G. Barkema, "Partitioning 3D space for parallel many-particle simulations," *Computer Physics Communications*, 149(3):121-134, 2003.

[TMG01] T. Theußl, T. Möller, and E. Gröller, "Optimal regular volume sampling," *IEEE Visualization*, pp. 91–98. 2001.

[VNP*04] D. Van De Ville, T. Blu,M. Unser,W. Philips, I. Lemahieu, and R. Van de Walle. "Hex-Splines: A novel spline family for hexagonal lattices*," IEEE Trans. Image Processing*, vol. 13, no. 6, pp. 758–772, 2004.

[XM06] F. Xu and K. Mueller, "GPU-Accelerated D2VR," *Volume Graphics Workshop*, pp. 23-30, 2006.

[XM07] F. Xu and K. Mueller, "Real-time 3D computed tomographic reconstruction using commodity graphics hardware," *Physics in Medicine & Biology*, vol. 52, pp. 3405-3419, July 2007.

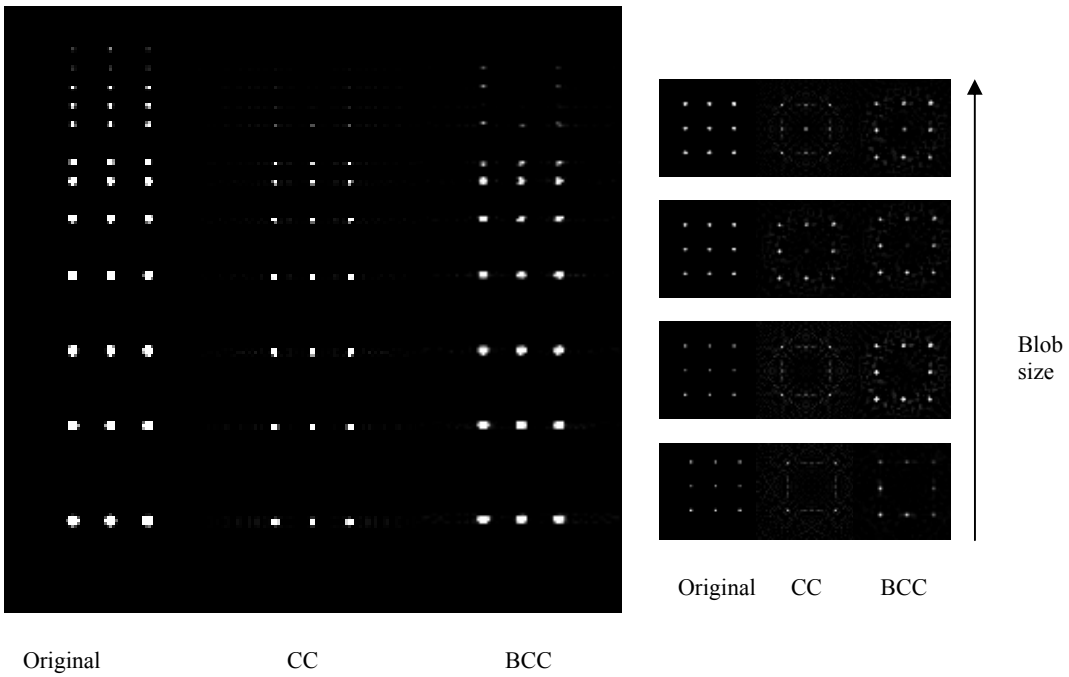

Figure 7: Blob phantom study. (Left): a row of blobs of decreasing size extends diagonally across space. The projections had a resolution such that the smallest blobs would still be detectable to at least 1-2 pixels. We then reconstructed these blobs on both a CC and the equivalent BCC lattice at a resolution matching that of the projections. At this lattice resolution the size of the smallest blobs amounted to about 1.5 voxels. (Right): The CC and BCC lattice reconstructions of these are shown as cross-sections: first the blob phantom, and next to it its reconstructions. We observe that the BCC lattice recovers the small blobs significantly better, and in some cases is the only lattice to recover them.

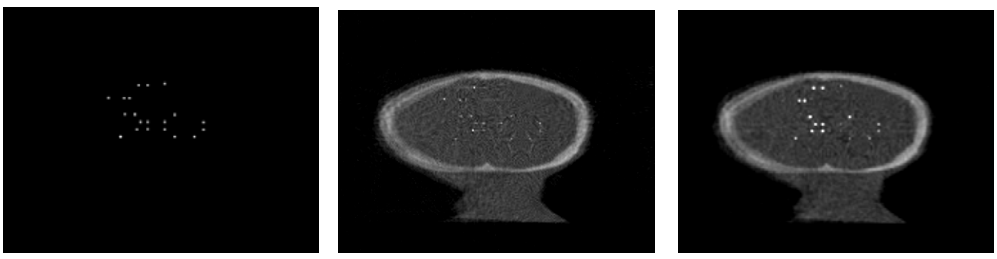

Figure 8: Tumor phantom study, which is a more realistic experiment than the blob study. A set of small tumors was distributed in 3D (a slice of these is shown on the left) and was embedded into a brain volume dataset. A corresponding slice of the reconstruction results using the CC and the BCC lattices are shown to the right, respectively. We see that the BCC lattice cannot recover all of them either (due to their small size), but it recovers significantly more of them and at higher intensity.