

A Granular Three Dimensional Multiresolution Transform

Alireza Entezari, Tai Meng, Steven Bergner and Torsten Möller[†]

School of Computing Science, Simon Fraser University, Vancouver, Canada

Abstract

We propose a three dimensional multi-resolution scheme to represent volumetric data in resolutions which are powers of two, resolving the rigidity of the commonly used separable Cartesian multi-resolution schemes in 3D that only allow for change of resolution by a power of eight. Through in-depth comparisons with the counterpart resampling solutions on the Cartesian lattice, we demonstrate the superiority of our subsampling scheme. We derive and document the Fourier domain analysis of this representation. Using such an analysis one can obtain ideal and discrete multidimensional filters for this multi-resolution scheme.

Categories and Subject Descriptors (according to ACM CCS): G.1.1 [Numerical Analysis]: Approximation
I.4.10 [Image Processing and Computer Vision]: Image Representation, Volumetric

1. Introduction

Multi-resolution schemes are important tools for dealing with large data. Different levels of detail of the data can be pre-computed and an appropriate level can be picked according to the available bandwidth of the display device or the transition channel (e.g. for online music, video, or graphics applications).

Ideally, one would like to find the appropriate level-of-detail, which preserves exactly as much data as can be handled by the underlying hardware or software constraints. In other words, we would like to have a continuous level-of-detail (LOD) slider, creating the LOD that is needed. However, it is only feasible to pre-compute finitely many LOD's. The choice of LOD's is often constrained by the underlying data structures and algorithms available to process these data structures.

One could easily use a continuous filter in order to reconstruct a continuous representation of the underlying data and to sample this continuous representation at any arbitrary resolution. Using this idea, any arbitrary granularity can be produced. However, this would require the use of a different filter mask for each sample point of the down-sampled grid. This is rather ineffective and computationally expensive. Hence, it is desirable to construct a multi-resolution pyramid using only discrete filters.

In image and volume processing it has been customary to work with image and volume pyramids that treat each dimension separately. In such a scenario the smallest possible granularity of sub-sampling would be to half the resolution of each dimension. This leads to the well-known quad-trees (in 2D) and octrees (in 3D). While we are only halving the resolution per coordinate axis during each iteration, we effectively reduce the overall data by a factor of four in 2D and eight in 3D.

While this is a convenient and widely used pipeline, the granularity may be too coarse and alternatives are of interest. Hence, quincunx filtering has been studied in the image processing community [Vai93,DM84] with the benefit of allowing a finer granularity for 2D image pyramids. Quincunx allows an overall data reduction by a factor of two in each iteration.

Van De Ville et al. [VDVBU05] have shown that a simple extension of the quincunx scheme to 3D is not possible and iterating through the commonly assumed Face Centered Cubic (FCC) subsampling cannot provide a suitable isotropic representation of the signal at various resolution levels. On the other hand, Linsen et al. [LPD*04] have studied a multi-resolution pipeline which has a change of resolution by a factor close to two in each step. In the first step of down-sampling, this method loses 3/8 of the data; in the second step it discards 3/5 of the data; and in the final step it loses 1/2 of the data to obtain a representation at 1/8 of the original data rate. Therefore, each resolution of the data is not

[†] {aentezar,tmeng,sbergner,torsten}@cs.sfu.ca

exactly halving the information. Another drawback is that this pipeline is made up of grids, for which no rendering algorithms with proper reconstruction filters exist. Moreover, it is not possible to analyze the representation of the data at various resolution levels in terms of the spectrum of the underlying signal. In other words, there are no signal processing tools available to analyze or predict the quality of the data representation and reconstruction on such grids.

In this paper we introduce a novel multi-resolution pipeline, which represents a volumetric dataset at one half of its original resolution on a Face Centered Cubic (FCC) lattice and at one quarter of its original resolution on a Body Centered cubic (BCC) lattice. The resolution of one eighths is again represented on a Cartesian lattice. Since proper interpolation and reconstruction schemes for these lattices exist we are able to render these 3D lattices properly and efficiently. Moreover, we can offer a Fourier analysis of this multi-resolution representation.

In Section 2 we review previous research on this topic and introduce our multi-resolution pipeline in Section 3. In Section 4 we are describing some details necessary for the implementation of our framework. Finally, in Section 5, we compare our novel pipeline to a pipeline based on comparable CC lattices only, which was computed using a continuous filter bank. We conclude our paper in Section 6 and point to some ideas on how to further improve our results.

2. Related Work

Since it is not always possible to interactively process data in the original sampling resolution, level-of-detail and multi-resolution techniques have been proposed by many researchers to balance between architectural constraints (e.g. performance and memory) and fidelity. Multi-resolution techniques have been employed in a variety of visual data representation approaches including geometric rendering [WG92, CDL*96, LWC*02] and volumetric data visualization [LHJ99], especially for out-of-core applications [SCM99].

Progressive Meshes [Hop96] probably come closest to a continuous LOD pipeline. This approach becomes practical through an efficient data structure, which minimizes the overhead. Unfortunately, nothing close to such elegance is known in image or volume processing. The de-facto standard is quadtree and octree structures. They are easy to implement and they do not alter the underlying lattice structure (this pipeline consists solely of CC lattices).

Octrees have been well studied [Sam90] and are widely used in scientific visualization and graphics [WG92]. They are the primary tool to deal with large and time-varying data [SCM99, YMC05].

In 2D image processing, the quincunx lattice provides more flexibility than a quadtree structure, since it preserves

a Cartesian lattice (through a rotation by 45°), while, the down-sampling rate is two rather than four as for quad-trees.

The FCC lattice is likened in structure to the quincunx lattice and provides a down-sampling by a factor of two in 3D. However, Van De Ville et al. show that no isotropic down-sampling pyramid, based on the dilation matrix induced by the FCC lattice, can be computed [VDVBU05]. Linsen et al. designed a wavelet multi-resolution pyramid which is based on an $\sqrt[3]{2}$ subdivision scheme [LPD*04]. Their multi-resolution pyramid goes from a CC to a modified FCC (with an additional sample in the middle of each cell) to a BCC lattice. Because of the existing filter banks associated with each level, it is potentially a very attractive pipeline. However, the modified FCC lattice of this pipeline is a new grid for which no proper interpolation filter is known. Further, no proper Fourier transform is known for this lattice and hence, it is not possible to analyze the spectrum of the data. Therefore, we cannot use this lattice in a volume visualization algorithm.

Recently we have devised interpolation and reconstruction kernels based on box splines for the BCC lattice and similar linear and cubic element kernels for the FCC lattice [QEE*05, EDM04]. We have demonstrated that efficient and fast linear reconstruction for both BCC and FCC lattices outperforms the trilinear interpolation for the Cartesian lattice. Also we have recently devised an extremely efficient evaluation method for the cubic box spline that will be published shortly.

The problem of optimal down-sampling relates to the classical cover problem. The cover problem is the problem of covering space with equal sized spheres while allowing the smallest amount of overlap [CS99]. When down-sampling, we are moving the copies of the spectra in the Fourier domain closer to each other. In order to minimize aliasing, we would like to minimize the amount of overlap of the aliased spectra. Hence, we would like to arrange the aliased spectra so that they incur the smallest amount of overlap. With the assumption of an isotropic signal spectrum, the Fourier domain interpretation of optimal down-sampling yields a tile of spheres where the overlap between the replicas of the spectrum is minimized. As discussed by Conway and Sloan [CS99], BCC attains the best (smallest) cover overlap. Second to that is the FCC pattern, and the Cartesian cubic ranks third in this measure. While these lattices give us insight over aliasing behavior in the frequency domain, their duals give the actual sampling pattern to use in the spatial domain. Hence, we first down-sample to FCC (which is dual to BCC), second we down-sample to BCC (which is dual to FCC) and finally we down-sample to Cartesian.

3. Subsampling Lattices

The goal of subsampling is to obtain a sub-lattice of the original lattice such that the sub-lattice has the appropri-

ate sampling density based on a desired subsampling ratio. Our construction of the multi-resolution transform is mainly motivated by suitable properties of the BCC and FCC lattices. Considering the original data given on the Cartesian lattice of \mathbb{Z}^3 , an FCC sub-lattice can be constructed, that has half its density. Moreover, a BCC sub-lattice of \mathbb{Z}^3 can be constructed whose density is a quarter of the original Cartesian lattice. The nesting structure of the FCC and BCC sub-lattices inside the original Cartesian lattice is depicted in Figure 1.




Figure 1: The Cartesian lattice and the structure of the sub-lattices. The FCC has half the sampling density and BCC has a quarter sampling density. The sampling volume for each sampling lattice is the inverse of the volume of the Voronoi cell of that sampling lattice.

3.1. Face Centered Cubic Lattice

The FCC lattice is commonly referred to as the checkerboard lattice in 3D. This property illustrates the structure of the FCC lattice as a sub-lattice of the \mathbb{Z}^3 with only half of the points. The lattice points belonging to the FCC lattice have the property that the sum of their coordinates is even.

In terms of signal processing theory of band-limited functions, the spectrum of a band-limited function represented on the FCC lattice is contained in the Voronoi cell of its dual lattice. The dual to the FCC lattice is a BCC lattice whose Voronoi cell is a truncated octahedron. Therefore, the spectrum of a properly band-limited function represented on the FCC lattice is contained in a truncated octahedron in the Fourier domain. This is illustrated in Figure 2b.

3.2. Body Centered Cubic Lattice

The BCC lattice points have the property that all three coordinates of each point have the same parity. In other words, a lattice point belongs to the BCC lattice if and only if all three of its coordinates are even or all three are odd.

As the FCC and BCC are dual lattices of each other, the spectrum of a properly band-limited function that is represented on the BCC lattice is contained in the Voronoi cell of the dual FCC lattice. The Voronoi region of an FCC lattice is a rhombic dodecahedron as illustrated in Figure 2a.




Figure 2: The Voronoi cells

3.3. Fourier Domain Representation

Assuming the original volumetric dataset is properly band-limited and sampled on a Cartesian lattice, and since the dual to the Cartesian is still a Cartesian lattice, the spectrum of the original signal is contained in a Voronoi cell of the Cartesian lattice, which is a cube.

In order to properly subsample this original function onto a sub-lattice, one needs to perform a filtering step to appropriately band-limit the original function for the sampling at the lower rate (i.e. the sub-lattice). The Nyquist region for the sub-lattice sampling is determined by the geometry of the sub-lattice. The Voronoi cell of the dual to the sub-lattice is precisely the Nyquist region for the sampling process. In other words, to perform the proper band-limiting process, one needs to filter out the parts of the original spectrum that fall outside the Voronoi region of the dual to the sub-lattice. In geometric terms, one needs to cut out and preserve the Voronoi cell of the dual to the sub-lattice from the original cubic spectrum.

The Nyquist region for subsampling the CC volume onto an FCC lattice is the Voronoi cell of the BCC lattice (a truncated octahedron). This polyhedron occupies exactly half the volume of the original cubic spectrum, as in Figure 3a. Similarly, when subsampling the Cartesian volume down to the BCC lattice, the Nyquist region is the Voronoi cell of an FCC lattice (a rhombic dodecahedron). This polyhedron occupies exactly a quarter volume of the original cube as in Figure 3b.




Figure 3: Nyquist regions for various subsampling steps. The gray cube indicates the support of the spectrum of the original Cartesian-sampled 3D function.

4. Implementation

The multi-resolution framework we describe here consists of two main parts. The *filtering process* can be regarded as a pre-processing step. It results in data representations at different scales (levels of detail). The key point of our method is that these representations are on different lattices and can be obtained from the original resolution Cartesian lattice by the use of discrete filters. The second subsection is concerned with *rendering* those lattices to the screen using ray-casting. Here, proper interpolation techniques are the key for a high-quality rendering process.

4.1. Filtering process

The main challenge in designing valid filters for the band-limiting process before subsampling, is to devise a band-pass filter whose support resembles the Voronoi cell of the dual lattice of the subsampling lattice. Therefore, we require filters whose support is contained within the polyhedron in Figure 3 and zero everywhere else. This would constitute the ideal subsampling filter for these sub-lattices. While these ideal filters are band-limited in the frequency domain, they have infinite support in the spatial domain and hence are impractical to use.

Our approach for designing compact support filters for the subsampling process was to take advantage of the linear and cubic reconstruction filters we recently introduced for the BCC and FCC lattices in [QEE*05, EDM04], respectively. These filters do have the proper support in the frequency domain.

For our experiments we used the linear element FCC filter that is a linear filter on the first neighbors cell of the FCC, where the first neighbors cell is a cub-octahedron (see Figure 4a). The cubic FCC filter is the convolution of the linear element FCC with itself [QEE*05]. Similarly, we used the linear box spline BCC filter whose support is a rhombic dodecahedron (see Figure 4b). The cubic box spline is the convolution of the linear box spline with itself [EDM04].

For evaluation purposes we also computed a comparable multi-resolution pyramid based solely on CC lattices. For this Cartesian pipeline, since there is no sub-lattice of a half

and a quarter resolution, we resort to resampling techniques. To obtain a properly band-limited function at half the resolution, we reconstructed the original (continuous) volume using a tri-cubic B-spline filter. The filter size was properly chosen, such that a proper band-limiting process occurred before the lower rate resampling. A similar computation was performed to obtain the CC lattice at one quarter of the resolution from the original lattice.




Figure 4: The first neighbor cells

4.2. Rendering Multi-resolution Lattices

In this section, we present our rendering framework for ray-casting volumes sampled on the Cartesian, BCC, and FCC lattices. First, we design a General Ray-caster that operates on continuous volumes. Then, we implement various reconstruction kernels for Cartesian, BCC, and FCC lattices to produce continuous volumes from datasets sampled on these lattices. This architecture enables the maximum reuse of ray-casting code common to all lattice types. Furthermore, minimal coding is needed to extend this framework to new lattice types, or new reconstruction kernels for existing lattice types.

Most parts of the standard volume rendering pipeline remain the same when rendering different lattices. The step that needs adjustment is the interpolation stage where a continuous volume is reconstructed from the given sampling lattice.

Datasets are sampled on a given lattice in world-space. We record the scaling and offset needed to arrive at the corresponding canonical-lattice-space. In particular, in the canonical Cartesian lattice, the distance between neighboring points along any (x , y , or z) axis is 1. In the BCC and FCC lattices, that distance is 2, as it can be seen in Figure 1. We now transform the canonical-lattice-space into a compacted-Cartesian-space, in preparation of data storage. Finally, coordinates in the compacted-Cartesian-space are transformed to 1-dimensional indices in memory-space, and samples are stored on file accordingly.

All the afore-mentioned transformations are well understood, except the one going from canonical-lattice-space to compacted-Cartesian-space. For the Cartesian lattice this




Figure 5: *BCC indexing scheme. The sequence of images shows a simple $2 \times 2 \times 4$ BCC dataset being mapped from canonical-lattice-space to compacted-Cartesian-space. Green (dotted) lines illustrate even z -slices of the BCC dataset. Red (solid) lines illustrate odd z -slices. Left column: canonical-lattice-space for BCC. Right column: compacted-Cartesian-space for CC; note that the size of each compacted Cartesian slice has been decreased by a scaling factor of two (and thus the concept of 'compaction'). Top row: front view. Bottom row: side view.*

step is just the identity transformation. For the BCC lattice, as shown in Fig. 5, we shift down all odd-numbered z -slices (colored red, connected by solid lines) by an offset of 1 along the x and y axes, so that they align with the even slices along the z axis. Then, we compact the x and y axes by a factor of 2, so that the resulting compacted-Cartesian-lattice becomes a canonical-Cartesian-lattice. For the FCC lattice, illustrated




Figure 6: *FCC index compaction into Cartesian scheme (given in xyz -coordinate tuples) used for memory storage.*

in Fig. 6, we decompose each z -slice into two sub-slices, push the second sub-slice to be halfway between the current z -slice and the z -slice directly behind it, and then bring all sub-slices into alignment along the z axis. Finally, we perform the appropriate scaling to transform these z -aligned sub-slices into a canonical-Cartesian-lattice.

5. Results

In order to examine the multi-resolution scheme discussed in this paper, we have implemented two pipelines - one pipeline based on CC/FCC/BCC down-sampling and one solely based on Cartesian lattices. Both pipelines have a sample reduction of one-half at each step. Normal estimation, needed for shading, was based on central differencing

of the reconstructed continuous function both in the Cartesian and FCC/BCC case. Central Differencing is easy to implement and there is no reason to believe that it performs any better or worse than taking the analytical derivative of the reconstruction kernel [MMMY97].

We have chosen the synthetic dataset first proposed in [ML94] as a benchmark for our comparisons. The function was sampled at the critical resolution of 40^3 on the Cartesian lattice. We then obtained the corresponding down-sampled FCC and BCC data volumes using the filtering process described in Section 4.1. The FCC volume has half the samples of the original volume and the BCC volume has a quarter of the original samples. Therefore, we obtained a pair of linearly down-sampled and cubically down-sampled FCC and BCC volumes. Similarly, we constructed CC lattices at one half and one quarter of the original resolution via a filtering step discussed in Section 4.1. Again we produced a pair of linearly and cubically re-sampled Cartesian datasets at each resolution.

The rendering pipeline renders images of the FCC volumes using the linear and cubic element filters introduced in [QEE*05] and the BCC renderer makes use of the linear and cubic box splines as in [EDM04]. The Cartesian images were rendered using the popular tri-linear and tri-cubic B-spline interpolations. For each sub-sampled dataset rendered, we produced an error image to visualize the errors incurred in this lower resolution representation of the original volume. Since the gradient error is perceptually significant, we mapped the angular error of the computed gradient (via central differencing) to gray-scale. We mapped the maximum angular error to white and zero error to black. In our experiments we used an angular error of 2 radians as our maximum error.

The left image of Figure 7 depicts the original ML dataset at 40^3 resolution (using tri-linear interpolation), and on the right is the corresponding error image. The error image is black when there is no error in estimating the gradient and the white and gray areas convey the angular error of the gradient estimation.




Figure 7: *ML Data at a CC resolution of 40^3 (left) and its error image (right). The max angular error depicted is 2 radians.*

Figure 8 displays the half resolution volumes. The top row

is the image of the FCC volume (left) and its error (right). The bottom row is the corresponding resolution on the Cartesian pipeline. These volumes are down-sampled using the corresponding linear filters and rendered using the linear reconstruction filters. The appearance of brighter tints of gray in the error image of the Cartesian representation indicates the appearance of more errors in the Cartesian pipeline. Figure 9 displays the half resolution volumes like Figure 8




Figure 8: Top row: FCC at half resolution (left) and its error image (right); Bottom row: CC at half resolution (left) and its error image (right). The max angular error depicted is 2 radians. We used linear down-sampling and linear reconstruction filters.

using a linear down-sampling filter. These volumes were rendered using cubic reconstruction filters. We can still observe the dominance of dark and black regions in the FCC error image that indicates fewer errors in the FCC representation.

Figure 10 displays the quarter resolution volumes. The top row depicts the image of the BCC volume (left) as well as the gradient error of the BCC volume (right). The bottom row is the corresponding resolution on the Cartesian pipeline. These volumes are down-sampled using the corresponding linear filters and rendered using the linear filters for reconstruction. Figure 11 displays the quarter resolution volumes using linear down-sampling filters. These volumes were rendered using the cubic reconstruction filters. The appearance of brighter tints of gray in the error image of the Cartesian indicates that more errors exist in the Cartesian pipeline. We can still observe the dominance of dark regions in the BCC error image that indicate existence of fewer errors in the BCC representation.

Figure 12 depicts the data at a resolution of one-eighth of the original resolution.

We also compared the two multi-resolution pipelines on an experimental data set. Although we can not quantify the




Figure 9: Top row: FCC at half resolution (left) and its error image (right); Bottom row: CC at half resolution (left) and its error image (right). The max angular error depicted is 2 radians. We used linear down-sampling and cubic reconstruction filters.




Figure 10: Top row: BCC at one quarter resolution (left) and its error image (right); Bottom row: CC at one quarter resolution (left) and its error image (right). The max angular error depicted is 2 radians. We used linear down-sampling and linear reconstruction filters.

accuracy of the representations, we included the actual rendering of the dataset at each resolution for comparison purposes. The results can be seen in Figure 13. In the middle rows, the left images belong to FCC and BCC representation and the right image to the corresponding Cartesian multi-resolution pipeline. The original skull dataset was sampled at the resolution of 128^3 . It becomes clear from the images that our new multi-resolution pipeline is preferable since it




Figure 11: Top row: BCC at one quarter resolution (left) and its error image (right); Bottom row: CC at one quarter resolution (left) and its error image (right). The max angular error depicted is 2 radians. We used linear down-sampling and cubic reconstruction filters.




Figure 12: CC at one eighth of the original resolution (left) and its error image (right). The max angular error depicted is 2 radians. We used linear down-sampling and linear reconstruction filters.

preserves more detail. However, some rendering artifacts are visible for the FCC image, which are due to a linear FCC filter that exhibits girdering artifacts (see [CMS01]). In our experimental setting we have experienced with filters that would get rid of these artifacts. These filters will be published shortly.

6. Conclusions and Future Work

We introduced a multi-resolution algorithm that allows the change of resolution in 3D with powers of two. Moreover, we have demonstrated that the representation of the data on these lattices is a more accurate representation than re-sampling to a lower resolution Cartesian lattice. This was illustrated through the error images in the results section. Analyzing the images in the results section and other possible combinations of down-sampling and reconstruction filters, we noted that linear type down-sampling along with cu-



Figure 13: The skull data set at the original resolution of 128^3 at the top. First row, at half the resolution, left image on FCC, right on re-sampled Cartesian lattice. Second row at quarter resolution, left image on BCC and right image on re-sampled Cartesian. The bottom row, the dataset at one eighth resolution. These volumes are down-sampled using linear type filter and reconstructed using cubic type reconstructions.

bic type reconstruction produces the most suitable results in terms of the smoothness of reconstruction, while this configuration keeps the angular gradient estimation errors low.

Since the FCC lattice constitutes the data representation at the half resolution, it is important to design accurate filters with desirable smoothness properties. The linear and cubic element filters on the FCC lattice do not have any proven smoothness properties. We are currently investigating filters for the FCC step with variable guaranteed smoothness. We have some preliminary results on a B-spline family of filters for the FCC lattice. Using these filters, we could eliminate

the artifacts introduced on the FCC resolution image of Figure 13.

Currently, we are able to quantitatively visualize the errors in the various multi-resolution steps, through comparison with the analytical ML function. For other real life datasets, we also would like to quantitatively visualize the errors of various resolution datasets with respect to their original high resolution volume. This error analysis will enable us to examine our multi-resolution scheme against the Cartesian pipeline on a number of other real life datasets.

Possible applications for our fine-grained multi-resolution pyramid include LOD or progressive rendering, 3D multi-scale feature extraction, or multi-scale spatio-temporal feature extraction in video data, to just name a few directions. Other future development could go towards a generalized lattice rendering beyond 3D. Key issues here are appropriate linear or higher order filtering kernels, and a unified indexing scheme. Apart from that GPU accelerated lattice rendering of the 3D lattices (BCC, FCC) using higher-order kernels is currently being investigated.

Further, perfect reconstruction filter banks could be designed for these non-separable lattices.

7. Acknowledgement

The research in this paper has been supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada. as well as the Canada Foundation of Innovation (CFI).

References

- [CDL*96] CHAMBERLAIN B., DEROSE T., LISCHINSKI D., SALESIN D., SNYDER J.: Fast rendering of complex environments using a spatial hierarchy. In *GI '96: Proceedings of the Conference on Graphics Interface '96* (1996), pp. 132–141.
- [CMS01] CARR H., MÖLLER T., SNOEYINK J.: Simplified subdivisions and sampling artifacts. In *VIS '01: Proceedings of the conference on Visualization '01* (Washington, DC, USA, 2001), IEEE Computer Society, pp. 99–106.
- [CS99] CONWAY J., SLOANE N.: *Sphere Packings, Lattices and Groups*, 3rd ed. Springer, 1999.
- [DM84] DUDGEON D. E., MERSEREAU R. M.: *Multidimensional Digital Signal Processing*, 1st ed. Prentice-Hall, Inc., Englewood-Cliffs, NJ, 1984.
- [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.
- [Hop96] HOPPE H.: Progressive meshes. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1996), ACM Press, pp. 99–108.
- [LHJ99] LAMAR E., HAMANN B., JOY K. I.: Multiresolution techniques for interactive texture-based volume visualization. In *Proceedings of the Conference on Visualization '99* (1999), pp. 355–361.
- [LPD*04] LINSEN L., PASCUCCI V., DUCHAINEAU M. A., HAMANN B., JOY K. I.: Wavelet-based multiresolution with nth-root-of-2 subdivision. *Journal on Computing special edition* (2004).
- [LWC*02] LUEBKE D., WATSON B., COHEN J. D., REDDY M., VARSHNEY A.: *Level of Detail for 3D Graphics*. Elsevier Science Inc., New York, NY, USA, 2002.
- [ML94] MARSCHNER S. R., LOBB R. J.: An evaluation of reconstruction filters for volume rendering. In *Proceedings of the IEEE Conference on Visualization 1994* (Los Alamitos, CA, USA, Oct. 1994), Bergeron R. D., Kaufman A. E., (Eds.), IEEE Computer Society Press, pp. 100–107.
- [MMMY97] MÖLLER T., MACHIRAJU R., MUELLER K., YAGEL R.: A comparison of normal estimation schemes. In *Proceedings of the IEEE Conference on Visualization 1997* (Oct. 1997), pp. 19–26.
- [QEE*05] QIAO W., EBERT D., ENTEZARI A., KORRUSINSKI M., KLIMECK G.: Volq: Direct volume rendering of multi-million atom quantum dot simulations. In *Proceedings of the IEEE Conference on Visualization* (Oct. 2005), pp. 319–326.
- [Sam90] SAMET H.: *The design and analysis of spatial data structures*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990.
- [SCM99] SHEN H.-W., CHIANG L.-J., MA K.-L.: A fast volume rendering algorithm for time-varying fields using a time-space partitioning (TSP) tree. In *Proceedings of IEEE Visualization '99* (1999), pp. 371–377.
- [Vai93] VAIDYANATHAN P.: *Multirate Systems and Filter Banks*. Signal Processing. Prentice Hall, Englewood Cliffs, New Jersey 07632, 1993.
- [VDVBU05] VAN DE VILLE D., BLU T., UNSER M.: On the multidimensional extension of the quincunx subsampling matrix. *IEEE Signal Processing Letters* 12, 2 (February 2005), 112–115.
- [WG92] WILHELMS J., GELDER A. V.: Octrees for faster isosurface generation. *ACM Transactions on Graphics* 11, 3 (July 1992), 201–227.
- [YMC05] YOUNESY H., MÖLLER T., CARR H.: Visualization of time-varying volumetric data using differential time-histogram table. In *Workshop on Volume Graphics 2005 (VG05)* (Jun. 2005), pp. 21–29.