

DUODECIM

A Structure for Point Scan Compression and Rendering

Jens Krüger, Jens Schneider, Rüdiger Westermann[†]

Computer Graphics and Visualization Group, Technical University Munich

Abstract

In this paper we present a compression scheme for large point scans including per-point normals. For the encoding of such scans we introduce a particular type of closest sphere packing grids, the hexagonal close packing (HCP). HCP grids provide a structure for an optimal packing of 3D space, and for a given sampling error they result in a minimal number of cells if geometry is sampled into these grids. To compress the data, we extract linear sequences (runs) of filled cells in HCP grids. The problem of determining optimal runs is turned into a graph theoretical one. Point positions and normals in these runs are incrementally encoded. At a grid spacing close to the point sampling distance, the compression scheme only requires slightly more than 3 bits per point position. Incrementally encoded per-point normals are quantized at high fidelity using only 5 bits per normal (see Figure 1). The compressed data stream can be decoded in the graphics processing unit (GPU). Decoded point positions are saved in graphics memory, and they are then used on the GPU again to render point primitives. In this way we render gigantic point scans from their compressed representation in local GPU memory at interactive frame rates (see Figure 2).

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: 3D Graphics and Realism

1. Introduction

Despite the advances in CPU and graphics hardware technology, for the largest available point scans point based rendering applications still cannot run at acceptable rates. As rendering capabilities continue to increase, so do the sizes of data being visualized as well as the resolutions of displays being used. Today, laser range scans comprised of almost a billion of vertices are available [Lev00, LPC*00], making even CPU processing difficult due to memory constraints. Figure 2 shows such gigantic scans, the largest of which consists of 250 millions of vertices and requires 6 GBytes to store positional and normal information. Because of the extraordinary richness of detail in these scans, the need for techniques able to reveal even the finest structures is becoming increasingly important. In addition to such scans, high resolution display systems of over 10 Mpixels are nowadays

available, letting the required bandwidth to transmit primitives to the GPU grow substantially. As these requirements will continuously increase in the future, there is a dire need for point rendering techniques that comprehensively address these issues.




Figure 1: Comparison of the original Atlas point scan [LPC*00] including normals (6 GB) to the point scan that was compressed by our method (231 MB). Note how the fine scale detail is preserved.

[†] jens.krueger@in.tum.de, jens.schneider@in.tum.de, westermann@in.tum.de




Figure 2: The three largest scans provided by the Michelangelo project are shown. All three scans, including per-point normals, have been compressed and now fit into 256 MB video memory on recent graphics cards. Images are generated by rendering the scans out of the compressed data stream on the GPU. Up to 50 million points per second can be decoded and rendered on consumer class graphics hardware.

In computer graphics, point based rendering has recently gained increasing popularity due to the simple and memory friendly nature of point rendering primitives. Such primitives do not require consistent topological information and they considerably reduce overdraw if high resolution models are rendered. A thorough discussion of these issues as well as a summary of recent point rendering techniques including various applications can be found in [GPA*, KB04].

First considered by Levoy and Whitted [LW85] and then revived by Grossmann and Dally [GD98], rendering systems based on point primitives have been proposed for both the hierarchical display of large point scan models [RL00] and high quality rendering of point sampled geometry [PZvBG00, ZPvBG01]. Due to the frequent use of such systems in practical applications, over the last few years there has been an ongoing improvement in this field with respect to rendering speed and quality, i.e., by exploiting graphics hardware and efficient GPU data structures [RPZ02, DVS03], by using high-quality point splats [BK03, ZRB*04], and through the use of point hierarchies [GM04] in combined with polygonal mesh representations [BK03, ZRB*04] to allow for efficient LOD rendering.

Besides rendering quality and speed, today's point rendering systems are facing the problem of continually increasing point sets. To keep up with this progress, several issues have to be considered: For large point scans the CPU might not be equipped with sufficient system memory. If the CPU works on a compressed data set, it might not be powerful enough to decode point positions and attributes at sufficient rates. If, on the other hand, a streaming representation is available that enables out-of-core rendering, disk access will most likely limit the overall performance. Even if the CPU could provide point rendering primitives at sufficient rate, the bandwidth of the communication channel connecting the CPU with the GPU might be too low as to allow for the transfer of point positions and attributes a fixed number of times per second. Finally, the GPU itself might not be able to render the points within the requested time interval.

Up to now, only a few approaches have focussed on these issues explicitly. A popular technique is to quantize point positions with respect to a Cartesian grid hierarchy, either by absolute position encoding or relative to a parent node in this hierarchy [RL00, SK01, BWK02]. Although these approaches can significantly reduce the required number of bits to encode point positions, a similar compression ratio has not yet been shown for normals. Ochotta and Saupe [OS04] locally parameterized point sets as a height field and re-sampled the point sampled surface on a regular grid. This method achieves high compression ratio by using wavelet transforms to encode the resulting height fields, but it introduces additional smoothing artifacts and produces a non-uniform sampling of the surface. A progressive compression scheme for point sets including per-point attributes based on multiresolution predictive encoding was presented by Waschbüsch et al. [WWL*04]. This scheme yields effective compression rates, but it suffers from both the high complexity of the matching process to find similar points and the rather costly decoding process, which recursively traverses binary trees to calculate initial point positions.

In this paper, we present a novel point rendering system for gigantic point scans that comprehensively accounts for the aforementioned requirements. This system is based on a lossy compression scheme for point positions and normals. Compared to all previous compression schemes for point sets, point coordinates can be decoded on the GPU. Our scheme has the following particular properties:

- **Memory efficiency:** We present an effective compression scheme for large point scans based on an optimal sampling of these scans.
- **Decoding efficiency:** The compressed stream provides random access to encoded points and attributes, and it can be decoded using a few simple arithmetic and logical operations.
- **Bandwidth efficiency:** Due to its simplicity, decoding can be performed on the GPU. To render the point set only the compressed data stream has to be transmitted.

- **Rendering Efficiency:** On the GPU, decoded point positions and normals are used in turn to render the point scan, which results in a significant performance gain compared to previous approaches.

To enable these properties, we introduce closest sphere packing (CSP) grids as a new and effective spatial data structure for point clustering. From closest sphere packing theory [CSB87] we know, that an optimal sampling in the spatial domain corresponds to the tightest arrangement of spheres in frequency domain. This can be derived from the observation that the spectrum of the sampled signal contains the replicas of the primary spectrum, centered at the points of the dual (or reciprocal) of the sampling grid. Optimal sampling of the signal is achieved if there are no overlappings between the replicas.

The CSP grids we employ are composed of Trapezo Rhombic Dodecahedra (TRD) - the dual of the Johnson Solid 27 [Joh66]. A TRD is a space filling twelve (twelve=duodecim (latin)) sided polyhedron, which constitutes a base element for a closest sphere packing of 3D space. CSP grids consisting of TRDs in particular (HCP), have no second order neighbors, i.e. no cells share only an edge (see figure 3). If such a grid is sliced orthogonal to the y-axis, the resulting 2D grid is composed of regular hexagonal cells. In this 2D grid, all neighbors share an edge and the distance between adjacent cell centers is constant. These properties are illustrated in figure 7. Note that besides the CSP grid we use, there exists the Face Centered Cubic (FCC) grid composed of the Rhombic Dodecahedron [Mat04, NM02]. Although both grids are closed sphere packing grids, HCP grids have some beneficial properties compared to FCC grids. Most importantly, the dual grid of FCC, the Body Centered Cubic grid (BCC), which is used for sampling, is not composed of a single cell type. Therefore the volume of the different sampling cells is distributed non uniformly. This increases the maximum sampling error compared to HCP grids, although the same number of cells is required to closely pack 3D space. To our best knowledge, the HCP grid has never been employed in computer graphics applications so far.

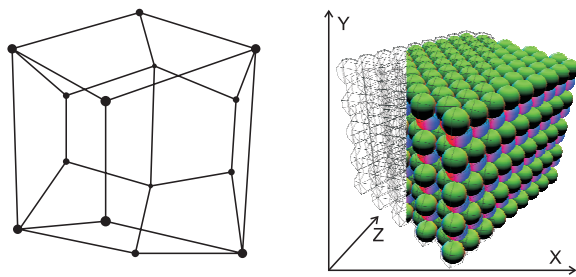


Figure 3: A trapezo rhombic dodecahedron and the corresponding closest sphere packing

We take advantage of HCP grids to establish an adjacency relation between points. Therefore, a binary spatial

data structure consisting of TRDs is generated, where a cell is full if it contains a point, and it is empty otherwise. The adjacency relation is exploited to incrementally encode runs of connected full cells in slices of the grid. This stage is very similar to the process described in [MH01] for the encoding of iso-surfaces in volumetric data sets. In contrast, however, we turn the problem of determining optimal runs into a graph theoretical one, which is posed as a search problem to find the minimum number of edges of specified length that cover an arbitrary graph.

In addition to incremental encoding of point coordinates, the same compression scheme can be applied to per-point attributes like normals or colors. Because these attributes show only a slight variation along the selected point runs, high fidelity at low bit rate can be achieved by differential encoding. To support view frustum and back face culling, run-specific attributes like cone of normals and bounding boxes are computed.

The compressed point set can be decoded on the CPU, and point primitives can be sent to the GPU for rendering. Alternatively, the compressed stream can be decoded on the GPU. Therefore, runs of equal length are stored in 2D texture maps and can then be decoded incrementally. Decoded point positions are first saved in graphics memory, and they are then used on the GPU again to render point primitives. This is realized using recent functionality like vertex texture fetches in the Shader Model 3.0 [Mic04] or OpenGL `GL_EXT_framebuffer_objects` [ATI04]. For the rendering of large point sets that do not fit into graphics memory, the presented point rendering system can either avoid bus transfer at all, or it can reduce bandwidth requirements substantially if the meshes are so large that even in compressed format they do not fit in graphics memory.

2. Uniform Point Clustering

To exploit geometric coherence in unstructured point sets, an adjacency relation between points is established first. We employ a regular spatial data structure composed of TRDs, into which the original point set is sampled. Every TRD that contains at least one point is marked with 1, while any other cell is marked with 0.

Compared to other grids, the HCP grid leads to an optimal sampling density, and in particular compared to Cartesian grids it yields a significantly smaller sampling error if a region in 3D space is partitioned using the same number of cells. In Cartesian grids with grid spacing h the maximum distance – and thus the upper bound for the sampling error – between a cell center and any other point inside a cell is $\sqrt{3}/2h$. For a HCP grid, on the other hand, it is $\sqrt{3}/(2 \cdot \sqrt{2})h$.

In the sampling process, the cell spacing of the HCP grid has to be determined such that as many cells as possible contain only one original point. On the other hand, by using ever

smaller cells, connectivity between filled cells will be lost, letting the incremental encoding scheme become less efficient. Due to this reason, sampling is implemented as a two step procedure that takes in account both constraints.

2.1. Sampling

To sample a point set, we start with an initial resolution of the HCP grid. In the current implementation this resolution is equal to the sample spacing of the scanning device. If this spacing is not known, an arbitrary initial guess can be specified. During the sampling process, we count for every cell how many points are sampled into this cell. Let us call this value the *hit rate*. Now the initial resolution of the HCP grid is iteratively refined until the hit rate drops below a given threshold.

For the sake of simplicity we explain the algorithm for the 2D hexagonal slices, the extension to 3D HCP grids is canonical. A constant time sampling strategy can be derived from the following observation. Considering the odd and even rows separately generates two cartesian grids. The vertex in question is sampled into these two grids. This generates two cells that correspond to two hexagonal cell candidates. To choose from these cells the distance from the vertex to the two cell centers is computed and the cell with the smallest is the correct hexagon (see Figure 4). In 3D four cell candidates are generated and four distances are compared to find the smallest value. Hence, sampling into the HCP reduces to a fixed number of modulo operations and distance calculations.




Figure 4: This figure illustrates the sampling of points into hexagonal grids using two staggered Cartesian grids.

As the point sets and thus the grids we are concerned with are very large and can usually not be stored in main memory, the entire sampling process is performed out-of-core. Point subsets are sequentially sampled into the grid, and they are then sorted on disk with respect to increasing cell index along the x -, z -, and y -axis. The sorted list can then be traversed sequentially to determine duplicate samples in one cell, and to compute the average hit rate. At the end, the point

set is implicitly given by a set of TRDs – or more precisely by the coordinates of their centers – that contain at least one point of this set.

3. Run Generation

The HCP grid provides a structure to generate contiguous runs of filled cells. This step is the transition from pure clustering to coherence based compression. The goal is to determine runs that are as long as possible, and to incrementally encode the cells in these runs. Starting with the grid index of the first cell in such a run, following cells can be encoded by storing the face they share with the predecessor. In this work, we restrict ourselves to the generation of runs within slices orthogonal to the grid y -axis. Because in such a slice every element only has 6 faces, adjacency information can be encoded in 2.25 bits.

Run generation proceeds layer by layer, reading all cells in the current layer from disk. By connecting these cells, an undirected graph is constructed. Run generation now operates on this graph, which makes the algorithm independent of the underlying grid structure. Ideally, the algorithm finds the minimum number of edge runs of given length SL that cover the entire graph and do not contain any edge twice. Since the solution to this problem is NP-hard, we present a linear-time 2-approximation, i.e., one that contains every cell at most twice. Note that this 2-approximation property is a very conservative upper bound. In all meshes we processed, the runs never contained more than 5 percent points more than once.




Figure 5: The Round-Trip Generation

At first, for every connected subgraph a single run is generated, so-called long-runs, which are then cut into pieces of length SL . Starting with an arbitrary node in the graph, as long as this node has exactly one neighbor that has not yet been visited, the current node is appended to the long-run and the neighbor becomes the current node. If there are no such neighbors, the run is terminated. Otherwise, for every neighbor that is encountered a new long-run including the current node as start element is generated. For every new run but the longest one we now produce a round-trip, i.e. a run that returns to its starting point. If all new runs, with the longest of these runs considered last, are appended to the current run, a contiguous path is generated that traverses every branch but the longest forth and back. What remains to

be done is to eliminate redundant pieces, i.e. round-trips that begin a run and parts that appear in another short-run. Figure 5 illustrates this algorithm. For the algorithm we never need to consider more than three neighbors. Figure 6 shows that if more than three neighbors are filled, these additional neighbors are traversed by the child calls already and do not need to be traversed by the parent anymore.




Figure 6: The images show a cell with four neighbors. Note that independent of the choice of the first child to be traversed, all neighbors are handled before the recursion returns to the parent node. For more than four neighbors this procedure is alike.

Due to the construction of runs, in every long-run an index can appear at most three times. However, for an index that appears more than twice there always exist at least one index that appears only once. This property can be proven by the means of complete induction over all T junctions. If a child run at a T junction was chosen to be serialized, it must have been the shorter run of the cell. Therefore the cells of the longer child are not duplicated. This always duplicates less than half of the cells while only the T junction itself could be tripled. Consequently, the algorithm computes a 2-approximation of the optimal solution.

This procedure generates for every connected subgraph of one slice a single run. These long-runs are cut into pieces of given length SL ; so-called short-runs. While this cutting takes place several optimizations are performed to further reduce the number of redundant cells. One of these optimization is to skip the way back of a previously serialized child run.

By restricting the maximum run length, the number of generated runs is increased at the same time decreasing the variation of their lengths. This kind of construction accommodates perfectly to GPU rendering in that SIMD streaming computations on such architectures can be exploited. If multiple processing units, i.e. fragment units, decode a number of runs in parallel, it is desirable for every run to contain exactly the same number of encoded points.

To generate a LOD hierarchy of the point set, sampling and run generation is repeated with decreasing resolution of the HCP grid. Starting with the optimal resolution, at every hierarchy level the resolution is reduced by a factor of two. Because the resolution at every coarser level is now fixed, grid size optimization of runs does not have to be carried out.

4. Optimization

To determine the optimal grid resolution for sampling and run generation, we consider the average length of short-runs in addition to the average hit rate. The first value measures how many points are lost due to the sampling process. The second value is a measure of the compression efficiency. A perfect sampling would result in an average hit rate of 1 and an average length of short-runs equal to SL . Making the grid cells smaller results in a lower hit rate but reduces the average run length. To find the optimal cell size we first start with an initial size, which is 1.5 times the average distance of points in the point set. If the average hit rate is above a limit, usually 1.6, the cell size is reduced according to ratio between the maximum hit rate and the current hit rate.

The sampling process is repeated with the new grid resolution, until the hit rate is below the maximum hit rate. Then, the run generation process is started. If the average length of short-runs is below a given threshold, usually 70% of the maximum length, we first try to close disconnected runs by inserting new cells. If this does not bring the average run length above 70%, the grid cell size is enlarged, sampling is repeated, and new runs are generated. The process terminates and outputs all current runs if the average run length is reached. Note that the 70% value is an arbitrary starting value for the algorithm based on our experience with different scans, even with other starting values the algorithm will still find an optimal value but it will possibly take longer to converge.




Figure 7: Geometric properties of a hexagonal 2D slice of the HCP with the one and two-ring neighborhood

To close disconnected runs, we search for filled cells in the 2-ring neighborhood (see figure 7) of those cells that are contained in runs shorter than 50% of the maximum length.

We do not try to connect longer runs because this could result in runs that are so long that they are cut into short runs later. If such a cell is found and both cells belong to different runs, the cell in between is set to connect the two runs. As the examples below demonstrate, this process adds far less than 10% of the initial cells on all models we tested.

5. Encoding

The proposed scheme generates a large set of runs less or equal to a given length. To every run, three 16 Bit values are associated: the first two values are used to encode the start cell of each run within the current slice. Positions in slices of a resolution of up to $2^{16} \times 2^{16}$ can thus be encoded. The third 16 Bit value is used as index into a codebook that contains quantized normals. The computation of this index as well as the codebook is described below.

5.1. Point Encoding

In a run, every point but the first one is encoded relative to its predecessor using 2.25 bits to encode a step to one of the 6 adjacent neighbors. In our current implementation however, we use 3 bits per vertex to keep the decoding process as simple as possible, and thus to enable the GPU to efficiently decode point runs. If decoding is to be performed on the CPU, memory requirements can be reduced about 25% by using all bits in the data stream.

5.2. Normal Encoding

Normals are either given for the original point set, or they are computed prior to the compression stage, e.g., by computing normals on a given triangulation or by moving least squares [ABCO*01].

As adjacent normals in a run only show a slight variation, they can be encoded incrementally. Every normal but the first one is expressed in spherical coordinates relative to its predecessor. Let θ and ϕ be the azimuth and the longitude coordinates, respectively, of the current normal. To avoid suboptimal compression at the poles we compute both the negative and the positive angles and use the one that leads to a smaller delta. If the difference to the following normal in spherical coordinates is $\Delta\theta, \Delta\phi$, then the new normal in Euclidean space coordinates is given by:

$$\begin{aligned} x &= \cos(\theta + \Delta\theta) \cdot \sin(\phi + \Delta\phi) \\ y &= \sin(\theta + \Delta\theta) \cdot \sin(\phi + \Delta\phi) \\ z &= \cos(\phi + \Delta\phi) \end{aligned}$$

As this computation requires trigonometric functions to be evaluated, we employ trigonometric relations

$$\begin{aligned} \sin(\alpha + \beta) &= \sin(\alpha)\cos(\beta) + \cos(\alpha)\sin(\beta) \\ \cos(\alpha + \beta) &= \cos(\alpha)\cos(\beta) - \sin(\alpha)\sin(\beta) \end{aligned}$$

to express Euclidean space coordinates in terms of pre-computed sine and cosine values. More precisely, given $\sin(\Delta\theta)$, $\cos(\Delta\theta)$, $\sin(\Delta\phi)$ and $\cos(\Delta\phi)$, as well as the respective values for the previous normal, Euclidean coordinates for the current normal can be decoded using simple arithmetic.

During run generation, we collect all normal increments in spherical coordinates that occur in the entire data set. These increments are then clustered using vector quantization [AG92]. The two angular increments in the codebook are stored as four sine and cosine values for each entry. In the current work, 16 bits are used to quantize the start normal of every run, and 5 bits are used to quantize normal increments.

6. Rendering

To render the compressed point set, the encoded data is traversed slice by slice. For each run, the start position is decoded from the associated grid coordinate, and the start normal is fetched from the quantization codebook. All other point coordinates can then be decoded incrementally from the relative offsets that are stored with respect to the underlying grid structure. Only for the incremental decoding of normals an additional lookup into the delta normal codebook is required. This process consecutively generates pairs of coordinates and normals, which are written to a vertex and a normal array. Via graphics APIs like OpenGL or DirectX, these arrays can then be issued for rendering. In the current implementation, points are rendered as screen aligned circles, although more elaborate and high quality splats can be integrated into the system straightforwardly.

If decoding is carried out on the CPU, the vertex and the normal array have to be sent to the GPU, making bus bandwidth a major bottleneck. To overcome this limitation, encoded runs are sent to the GPU in compressed form. Due to the simplicity of the decoding process, runs can be directly decoded on the GPU using parallel streaming computations. Reconstructed point positions are rendered directly without any read back to application memory. In this scenario, the CPU is only used to control which runs are sent to the GPU, i.e. to accommodate view frustum and backface culling (see below).

The GPU decoder exploits functionality on recent graphics cards. On such cards, it is now possible to access texture maps in the vertex units [Mic04] and to allocate memory objects (GL_EXT_framebuffer_object) that can be interpreted as texture maps and vertex arrays alternatively [ATI04]. The second alternative is exploited in this work.

To prepare compressed point runs for GPU processing, they are stored in 2D texture maps. For each run, its start position and normal is stored in a 16 bit RGB texture map. Consecutive points in a run are encoded in 8 bit luminance textures, using the first 3 bits to store adjacency information and the remaining 5 bits to store quantized delta normals.

These normals are decoded from a quantization codebook. The principal layout of all data structures on the GPU is illustrated in figure 8.

In consecutive rendering passes, the GPU first decodes per-point normals and renders these normals into an intermediate memory object, i.e. a normal map. Then, point positions are decoded, the normal map is read to enable lighting computations, and positions as well as computed colors are rendered into a second memory objects. This object is then bound as a vertex array and can thus be rendered without any read-back to the CPU. In the following pass, both intermediate memory objects are read to obtain previous point positions and normals required for incremental decoding. Then, GPU decoding proceeds as described.




Figure 8: This image illustrates the encoding of runs into texture maps on the GPU.

6.1. Culling and LOD

To render the compressed representation, the CPU determines the runs to be rendered, and it sends the textures containing these runs to the GPU. To keep bus transfer and GPU processing as minimal as possible, two different acceleration techniques have been integrated into our approach: First, runs are clustered, i.e. stored in the same texture, according to their cone of normals. This is the primary sorting criterion. Second, within one cone runs are grouped according to their spatial position, i.e. every texture is split into a set of smaller textures for which axis aligned bounding boxes are computed. This is the secondary sorting criterion. At run time, the CPU determines the partitions to be displayed based on the current viewing direction and the size and orientation of the view frustum. Only potentially visible partitions are sent to the GPU, where they are finally decoded and rendered.

The CPU also determines the most appropriate LOD, i.e. grid resolution, to be rendered. We always select the resolution such that grid cells are always projected into an area smaller the size of one pixel under the current viewing parameters. An example of a hierarchical LOD representation with accompanied bounding box hierarchy is shown in Figure 9.




Figure 9: The upper row shows closeups of David's head rendered at 4 different LOD levels. The lower row shows the corresponding image of the David statue as it would be rendered. The bounding boxes enclose parts of the mesh that are tested for frustum culling.

7. Results

The efficiency and effectiveness of the proposed point based rendering system were verified using the large scans from the Digital Michelangelo Project (see Figure 10). Above all, it should be noted the richness in detail that can be seen in these scans, and which is resolved at high fidelity using our approach. Table 1 shows comprehensive results for the three largest meshes of this archive as well as for one smaller reference mesh. Our target architecture is a P4 2.8 GHz CPU equipped with 1GB RAM and an ATI Radeon X800 XT graphics card with 256MB.



Figure 10: The David and Atlas statues from the Digital Michelangelo Project.

In all examples, run optimization resulted in a hit rate below 1.7 and a run efficiency above 70% of a maximum length of 25. As can be seen, even for the atlas mesh the algorithm returns the result in less than 10 hours. Due to the hit rate larger than 1, the original point sets were reduced by a factor of 1.3 to 1.6.

It is obviously clear, that the point clustering approach as described introduces sampling errors. For the presented high resolution examples, these errors are 0.11mm and 0.14mm.

The scanners used for the Digital Michelangelo Project have a minimum sample spacing of $0.25 \times 0.25 \times 0.1 \text{ mm}$ in a plane perpendicular to the laser [Cyb99]. In the worst case, two sample points are as much as $\sqrt{2 \times 0.25^2 + 0.1^2} \text{ mm} \approx 0.367 \text{ mm}$ apart, which is the minimum size of features that can be faithfully reconstructed by the scanning process. Because in all our examples the sampling spacing is significantly higher than the sampling error introduced by our compression method, features present in the original data sets will not be destroyed. If the scanning device has sampled the data above the Nyquist rate of the original signal, our sampling is well above this rate, too, resulting in equal visual quality of the original and the compressed point set (see Figure 1).

Table 1 also shows the excellent compression ratio our method achieves for real-world data sets exhibiting fine scale details. When using ZIP compression, the encoded VRIP version of all of the Digital Michelangelo statues requires about 380MB and can thus be stored twice on an ordinary CD. Plain encoded, it is still small enough to be stored in core of our target architecture. Due to the slice based encoding scheme for point sets, which is at the core of our technique, it is also well suited for streaming processing and progressive transmission of the data [IL04].

While the point scans are considerably compressed, they can still be decoded very efficiently due to the simplicity of the decoding scheme. In the table we give timings for GPU decoding *and* rendering. To measure these timings, acceleration techniques were all switched off, therefore these timings are considerably slower compared to the display times in practice. This can be seen in the accompanied video. If decoding is carried out on the CPU, we observe a loss in performance of about a factor of 13.

Model	Atlas	St. Matthew	David	Dawn
scan resolution	0.25 mm	0.25 mm	1.00 mm	1.00 mm
# Points	254904158	186865425	28184522	3432203
# Samples	158877859	121718168	17190274	2582256
hit rate	1.60	1.53	1.64	1.32
run efficiency	72%	74%	72%	73%
max sampling error	0.11 mm	0.14 mm	0.48 mm	0.44 mm
ply file size	9.94 GB	7.29 GB	1.1 GB	134 MB
DD compressed size	231 MB	182 MB	28.5 MB	4.2 MB
zip compressed file	172 MB	140 MB	21 MB	3.3 MB
encoding time	9.5 hrs	6 hrs	57 min.	5 min. 40 sec.
decode & render time	4.14 sec.	3.05 sec.	0.43 sec.	0.06 sec.

Table 1: Timing and memory statistics for the proposed point based rendering system.

On the GPU, the point rendering system achieves a throughput of about 50 million points per second. This rate includes the decoding of compressed point runs as well as the rendering of decoded points and normals. It is worth noting, that we decode and render about 160M distinct points and normals in roughly 4 seconds on the GPU. Figure 11




Figure 11: Zoom in on the Michelangelo's St. Mathews Statue, note the fine scale details even in the lower rightmost closeup.

shows some more examples that demonstrate the need for a point based rendering system able to handle such an amount of primitives. In all images, the point splat size is automatically set according to the screen space projection of the underlying grid cells.

8. Conclusion

In this paper, we have presented an effective compression scheme for gigantic point scans based on close sphere packing grids. Such grids provide a structure for optimal point clustering, and they establish a spatial relation between points that can be exploited for compression purposes. As our results have shown, the compression scheme achieves an extraordinary compression ratio at very high fidelity. Due to the simplicity of the decoding scheme, point coordinates and normals can be reconstructed on the GPU. As the GPU can also render the decoded primitives without any read-back to the CPU, bandwidth requirements are substantially reduced.

As is demonstrated in the paper, even though GPU rendering includes decoding of point coordinates as well as processing of point geometry, a throughput of 50 million points per second can be achieved. To our best knowledge, this has not been achieved to this day by any other method.

In the future, we will extend the rendering engine about more elaborate space partitioning strategies, which allow for improved culling if the user zooms into the data set or if only a small portion of the data is rendered. The integration of high-quality point rendering techniques, i.e. perspective correct sprites or Phong splats, will be considered as well.

9. Acknowledgement

We would like to thank the people from the Digital Michelangelo Project for scanning the statues and making the mesh data public.

References

- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Point set surfaces. In *Proceedings of Visualization '01* (2001), pp. 21–28.
- [AG92] ALLEN GERSHO R. M. G.: *Vector Quantization and Signal Compression*. Kluwer International Series in Engineering and Computer Science, 1992.
- [ATI04] ATI: Superbuffers OpenGL Extension. www.ati.com/developer/gdc/SuperBuffers.pdf, 2004.
- [BK03] BOTSCH M., KOBBELT L.: High-quality point-based rendering on modern gpus. In *Proceedings of Pacific Graphics 2003* (2003), p. 335.
- [BWK02] BOTSCH M., WIRATANAYA A., KOBBELT L.: Efficient high quality rendering of point sampled geometry. In *Proceedings of the 13th Eurographics workshop on Rendering* (2002), pp. 53–64.
- [CSB87] CONWAY J. H., SLOANE N. J. A., BANNAI E.: *Sphere-packings, Lattices, and Groups*. Springer-Verlag New York, Inc., 1987.
- [Cyb99] CYBERWARE: 3D Scanner Designed To Scrutinize Works Of Michelangelo. www.cyberware.com/news/pressReleases/pr023.html, 1999.
- [DVS03] DACHSBACHER C., VOGELGSANG C., STAMMINGER M.: Sequential point trees. *ACM Computer Graphics (Proc. SIGGRAPH '03)* 22, 3 (2003), 657–662.
- [GD98] GROSSMANN J., DALLY W.: Point sampled rendering. In *Proceedings Eurographics Rendering Workshop* (1998), pp. 181–192.
- [GM04] GOBBETTI E., MARTON F.: Layered point clouds. In *Eurographics Symposium on Point Based Graphics* (2004), pp. 113–120, 227.
- [GPA*] GROSS M., PFISTER H., ALEXA M., PAULY M., STAMMINGER M., ZWICKER M.: Point-based computer graphics. SIGGRAPH '04 Course Note.
- [IL04] ISENBURG M., LINDSTROM P.: *Streaming meshes*. Technical Report UCRL-CONF-201992, LLNL, 2004.
- [Joh66] JOHNSON N. W.: Convex Polyhedra with Regular Faces. *Canadian Journal of Mathematics* 18 (1966), 169–200.
- [KB04] KOBBELT L., BOTSCH M.: A survey of point-based techniques in computer graphics. *Computers & Graphics* 28, 6 (2004), 801–814.
- [Lev00] LEVOY M.: Digitizing the forma urbis romae. In *Proceedings of the ACM SIGGRAPH and Eurographics Campfire Workshop on on Computers and Archeology* (Snowbird, Utah, USA, 2000).
- [LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital michelangelo project: 3D scanning of large statues. In *ACM Computer Graphics (Proc. SIGGRAPH '00)* (2000), pp. 131–144.
- [LW85] LEVOY M., WHITTED T.: *The use of points as a display primitive*. Technical Report 85-022, University of North Carolina at Chapel Hill, 1985.
- [Mat04] MATTAUSCH O.: Practical reconstruction and hardware-accelerated direct volume rendering on body-centered cubic grids. In *CESCG 2004* (2004).
- [MH01] MROZ L., HAUSER H.: Space-efficient boundary representation of volumetric objects. In *Proceedings IEEE/TVCG Symposium on Visualization 2001* (2001), pp. 180–188.
- [Mic04] MICROSOFT: Shader Model 3 Specification. <http://msdn.microsoft.com>, 2004.
- [NM02] NEOPHYTOU N., MUELLER K.: Space-time points 4d splatting on efficient grids. In *IEEE Symposium on Volume Visualization '02* (2002), pp. 97–106.
- [OS04] OCHOTTA T., SAUPE D.: Compression of point-based 3d models by shape-adaptive wavelet coding of multi-height fields. In *Eurographics Symposium on Point Based Graphics* (2004).
- [PZvBG00] PFISTER H., ZWICKER M., VAN BAAR J., GROSS M.: Surfels: surface elements as rendering primitives. In *ACM Computer Graphics (Proc. SIGGRAPH '00)* (2000), pp. 335–342.
- [RL00] RUSINKIEWICZ S., LEVOY M.: Qsplat: a multiresolution point rendering system for large meshes. In *ACM Computer Graphics (Proc. SIGGRAPH '00)* (2000), pp. 343–352.
- [RPZ02] REN L., PFISTER H., ZWICKER M.: Object space ewa splatting: A hardware accelerated approach to high quality point rendering. In *Proceedings Eurographics 2002* (2002), pp. 371–378.
- [SK01] SAUPE D., KUSKA J.: Compression of iso-surfaces. In *Proceedings of Vision, Modeling and Visualization '01* (2001), pp. 330–340.
- [WWL*04] WASCHBÜSCH M., WÜRMLIN S., LAMBORAY E. C., EBERHARD F., GROSS M.: Progressive compression of point-sampled models. In *Eurographics Symposium on Point Based Graphics* (2004).
- [ZPvBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *ACM Computer Graphics (Proc. SIGGRAPH '01)* (2001), pp. 371–378.
- [ZRB*04] ZWICKER M., RASANEN J. J., BOTSCH M., DACHSBACHER C., PAULY M.: Perspective accurate splatting. In *Proceedings of Graphics Interface 2004* (2004), pp. 247–254.