A Flexible Pinhole Camera Model for Coherent Non-Uniform Sampling

Submission ID: papers_0428

Figure 1. Coherent non-uniform sampling (CoNUS) image that allocates more samples to the face regions (left), output frame reconstructed from CoNUS image (middle), and output frame reconstructed from a conventional image (right), for comparison.

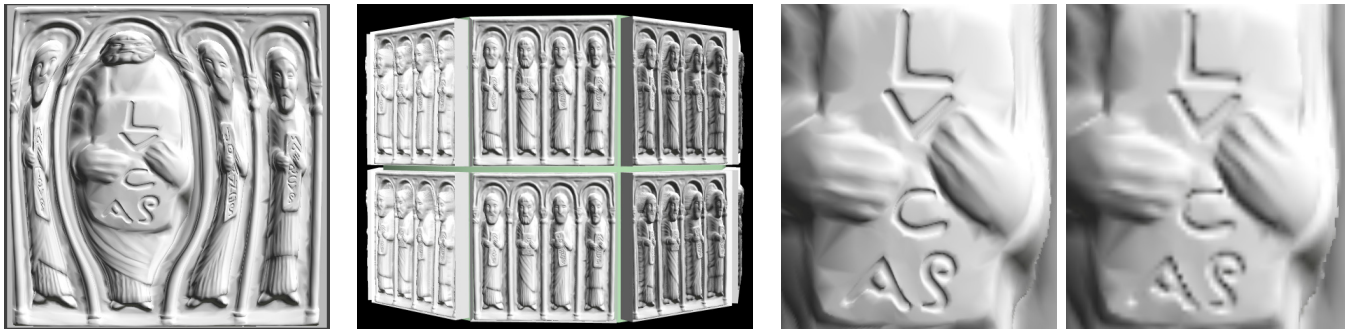

Figure 2. CoNUS relief texture that allocates more samples to a tablet of interest (left), relief texture mapping (middle), and comparison between frames that zoom in on tablet of interest and are rendered with CoNUS and with conventional relief textures of same size (right).

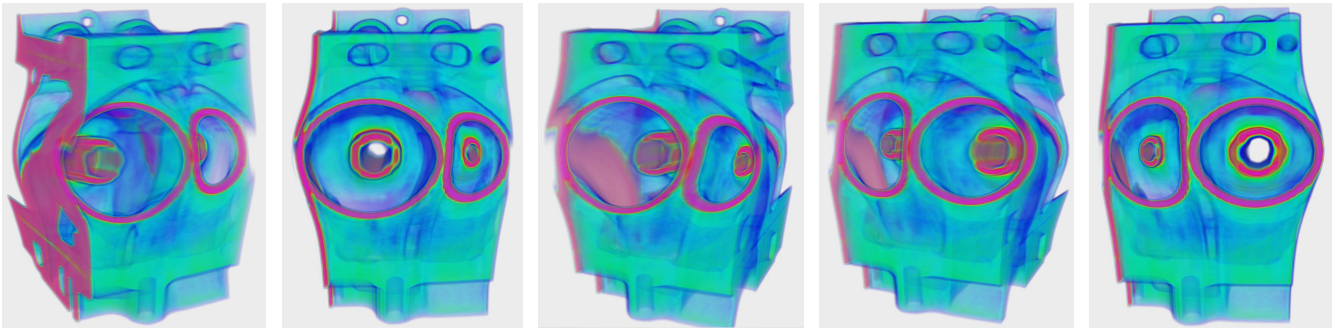

Figure 3. CoNUS focus-plus-context volume rendering visualization emphasizing the left (1-3) and then the right (4-5) cylinder housings.

Abstract

We present a pinhole camera model that allows modulating the sampling rate over the field of view with great flexibility. This flexible pinhole camera or FPC is defined by a viewpoint (i.e. eye) and by a sampling map that specifies the sampling locations on the image plane. The FPC supports many types of data, including image, height field, geometry, and volume data. The FPC provides an inexpensive 3-D projection operation which allows rendering complex datasets quickly, in feed-forward fashion. The resulting image is a coherent non-uniform sampling (CoNUS) of the dataset that matches the local variation of the importance of the dataset. We demonstrate the advantages of CoNUS images in the contexts of remote visualization, of focus-plus-context, and of acceleration of expensive rendering effects such as surface geometric detail and specular reflection rendering.

CR Categories: I.3.3 [Computer Graphics] Picture/Image Generation—Viewing algorithms.

Keywords: camera models, non-uniform sampling, interactive rendering, remote visualization, specular reflection rendering, surface geometric detail rendering, volume rendering, focus-plus-context visualization.

Figure 4. Visualization of sampling map for Figure 1.

1 Introduction

Most computer graphics and visualization applications employ images computed using the planar pinhole camera (PPC) model. The PPC is a good approximation of the human eye which makes it uniquely well suited for applications where the goal is to show users what they would see during an actual exploration of the scene. However, there are applications where the reduced field of view, the single viewpoint, and the uniform sampling rate limitations of the PPC model are a severe disadvantage.

In this paper we address the uniform sampling rate limitation of the PPC model. We introduce the flexible pinhole camera or FPC which allows for adjustments of the sampling rate according to the local importance or complexity of the data imaged. Like the PPC, the FPC is defined by a viewpoint (i.e. center of projection or eye) and an image plane. However, the sampling locations are not defined by a uniform grid but rather by a sampling map that allows shifting sampling locations from one region of the image plane to another. The FPC image provides a coherent non-uniform sampling (CoNUS) of the dataset. The CoNUS image in Figure 1, left, samples the five faces at a higher rate. The sampling map of the FPC used to render it is shown in Figure 4. The sampling map has the topology of a regular rectangular grid but it is distorted to implement the sampling rate modulation.

FPC CoNUS images preserve the advantages of conventional images. A CoNUS image can be computed quickly with the help of GPUs. Data access remains constant time, with the small additional cost of the sampling map indirection. A CoNUS image has good pixel to pixel coherence and conventional image compression algorithms apply. And finally A CoNUS image remains a single-layer 2-D array of samples which defines connectivity implicitly.

The sampling map underlying the FPC can be constructed in a variety of ways. We build sampling maps interactively using a physics-based mass-spring system, by composing multiple sampling maps together, or analytically. The FPC provides fast 3-D projection which allows rendering CoNUS images quickly, in feed-forward fashion, from many types of data. We demonstrate FPC rendering of CoNUS images from image, height field, geometry, and volume data. We explore the use of CoNUS images in the contexts of remote visualization, of focus-plus-context visualization, and of acceleration of expensive effects such as surface geometric detail and specular reflection rendering. We also refer the reader to the accompanying video.

Consider the example of a high resolution portrait photograph that has to be downsampled to be sent over the internet (Figure 1). The recipient is likely to want greater detail on the faces, which, of course, cannot be provided by zooming into the conventional downsampled image. If instead the server sends a CoNUS image of same size but with a higher sampling rate at the known regions

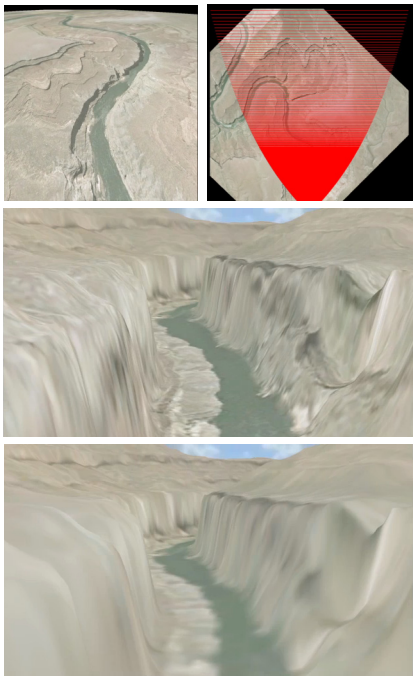

Figure 5. CoNUS height field (top left), sampling pattern of CoNUS height field (top right), output frame rendered from CoNUS (center) and from conventional height field (bottom).

of interest, the user can zoom in with better results. Consider a second example where a height field is visualized remotely. The server sends periodically height field sections corresponding to the current user location. If instead the server sends CoNUS height fields, the fidelity of the output frame increases considerably (Figure 5). Like any height field, the CoNUS height field samples the ground plane orthogonally, avoiding occlusions. However, the sampling pattern is defined analytically to match the sampling rate requested by the output frames.

Many techniques employ depth images in order to accelerate expensive rendering effects. In relief texture mapping a depth image is used to enhance a surface with geometric detail. Quality specular reflections are obtained if objects close to the reflector are modeled with depth images. The main reason depth images accelerate these effects is that one can compute the intersection between a ray and a depth image faster than one can compute the intersection between a ray and the original geometry. A CoNUS depth image brings sampling flexibility (Figure 2, Figure 6), without increasing the cost of the intersection operation.

Whereas in the examples presented so far the need for non-uniform sampling is to improve an auxiliary data representation from which conventional output images are computed, in the case of focus-plus-context visualization the CoNUS image is shown directly to the user. The FPC approach enables a versatile focus-plus-context visualization technique (Figure 3, Figure 7) that can handle any type of data and that provides good control over the focus regions. Note that the CoNUS images are rendered directly from the dataset (e.g. volume data, geometry) using the FPC, and they are **not** obtained by first rendering and then downsampling a high resolution conventional image.

2 Prior Work

We first review prior efforts aimed at removing the uniform sampling rate constraint of conventional images, and then review prior work in the application contexts where our paper examines the benefits of FPC rendered CoNUS images.

Figure 6. CoNUS depth image emphasizing all 4 engraved tablets (top left), scene setup (top right), and reflection details rendered with CoNUS (bottom left) and conv. (bottom right) depth image.

Non-uniform sampling rate

Hierarchical spatial partitioning schemes such as kd-trees do improve representation efficiency by stopping subdivision in regions where data is sampled accurately. One could define the FPC sampling map with such a partitioning scheme. Compared to our distorted grid approach (Figure 4), the hierarchical approach has the advantage of supporting a wider range of sampling rates, but it has the important disadvantages of sampling rate discontinuity, of lack of contiguity, and of greater construction and usage complexity.

Images with a non-uniform sampling rate were first obtained in the quest of removing the field of view limitation of conventional images. For spherical panoramas the sampling rate variation was an unwanted side effect and they were replaced by cube maps with a more uniform sampling rate. Recently single image panoramas received attention again, as the programmability of modern graphics hardware enables sampling patterns that avoid the earlier undersampling problems [1]. Researchers have also been addressing the single viewpoint limitation of conventional images with innovations at the camera model level such as the general linear camera [2] and the occlusion camera [3]. The FPC CoNUS is complementary to these approaches, providing sampling rate flexibility to panoramic and non-pinhole cameras.

Irregular sampling patterns have also occurred in the contexts of image-based rendering by 3-D warping [4] and of shadow antialiasing [5, 6]. In both cases depth images are reprojected to novel views where the forward mapped samples are irregular. The granularity with which sampling is controlled in the FPC CoNUS approach is insufficient to sample the shadow map precisely at the locations where the output image does, as needed to completely eliminate shadow aliasing. However, shadow aliasing could be reduced by using a CoNUS shadow map with a higher sampling rate in regions that are magnified in the output image.

The general pinhole camera is defined by a center of projection and an image plane with a set of arbitrary sampling locations [7]. This broadly defined theoretical camera model allows for any sampling pattern. The practical implementation however is restricted to a cubic or quadratic polynomial perturbation of the sampling locations in a rectangular region of the image plane. The FPC amounts to a different specialization of the abstract general pinhole camera model. The advantage of the FPC CoNUS approach is greater flexibility in defining the sampling locations. A 32x32 CoNUS sampling map defines 1,000 degrees of freedom compared to the 8 provided by the general pinhole camera's rectangular region and its border.

In texture mapping non-uniform sampling has been pursued through compression, atlasing, enhancement with explicitly modeled high frequency features, and distortion. We only discuss the last two approaches as they are closest to our work. Textures enhanced with edges modeling shadow silhouettes [11] or abrupt changes in color [9, 10] are more robust to magnification. The approach is compatible with CoNUS textures. When edges are derived from vector graphics primitives the edges have to undergo the sampling map distortion, and long edges have to be split. For edges derived from the texture, the CoNUS texture can be used directly. Space-optimized textures [8] distort textures with a similar mechanism to our sampling map. Our FPC work contributes a camera model that allows rendering CoNUS images efficiently from many data types, including 3-D geometry, contributes additional methods for sampling map construction, and applies CoNUS images to 3 different application domains.

Remote visualization

As the size of acquired and computed datasets continues to increase, so will the importance of remote visualization which is called upon to provide access to remote datasets for clients with no high-end storage and visualization capabilities. One approach is to reduce the dataset on the server to a size that can be transmitted to and visualized by the client. Many techniques can

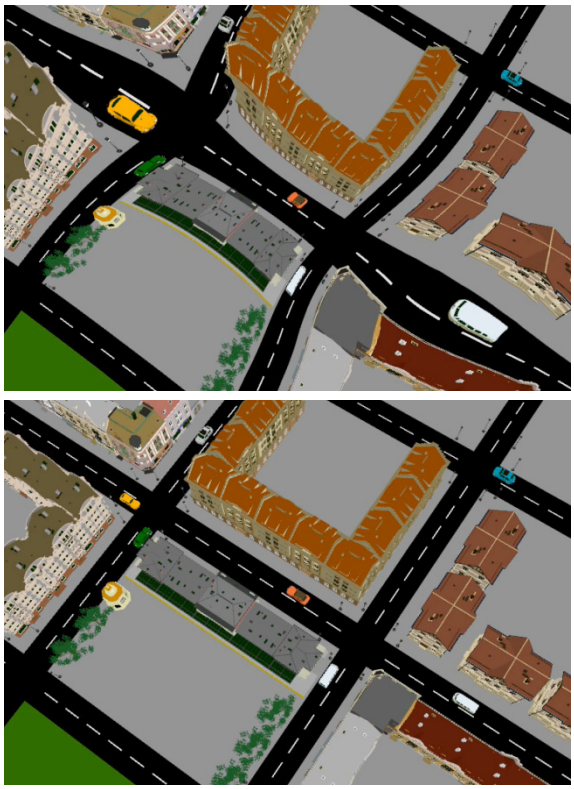

Figure 7. CoNUS focus-plus-context visualization emphasizing the yellow and white cars (top), and conventional image (bottom).

be used to this effect including data compression (e.g. [12, 13]), feature extraction (e.g. [14, 15]), and level of detail (e.g. [16]). A second approach is to compute the visualization at the server and send images to the client (e.g. [17]). No visualization capabilities are required at the client whatsoever, but the limited network bandwidth limits visualization resolution and frame rate.

A hybrid approach is to transfer images that have more data than what is needed for the current frame at the client. Such an enhanced image should be sufficient for a quality reconstruction of a sequence of frames at the client, without any additional data from the server. Images enhanced with per pixel depth [18] and with additional samples at the center [7] have been used to allow translating and zooming in at the client. Remote visualization based on transferring CoNUS images falls in this third—hybrid—category. A CoNUS image that samples known regions of interest in greater detail anticipates the user's intention to zoom in on those regions. A CoNUS height field that samples the ground plane orthogonally yet at a higher rate close to the user supports 6 degrees of freedom navigation at the client in the neighborhood of the current view.

Rendering acceleration using depth images

Depth images are powerful geometry approximations used for acceleration in many contexts including rendering complex geometric surface details, specular reflections, refractions, and ambient occlusion. We limit the discussion of prior work to the first two contexts which are used in this paper to illustrate the advantages of FPC-rendered CoNUS depth images.

Relief texture mapping is a technique for adding geometric detail to surfaces. The technique produces correct silhouettes and correct interactions between relief and other relief and non-relief geometry (e.g. intersections, casting and receiving shadows) [19, 20]. The relief texture is a depth image attached to a base box. Rendering the box triggers intersecting the eye ray with the depth image at every pixel. The intersection computation is performed by projecting the ray onto the depth image and following the ray projection until the first intersection is found.

Specular reflections are challenging for the feed-forward 3-D graphics pipeline because one cannot easily compute the image plane projection of reflected vertices. We group specular reflection rendering techniques into four categories: ray tracing [21], approximations of the projection of reflected vertices (e.g. [22]), image-based rendering (e.g. [26, 27]), and approximations of the reflected scene. We only discuss the fourth category since it is the category where the CoNUS specular reflection rendering method falls. The most drastic approximation is undertaken by environment mapping [25], where the reflected scene is assumed to be infinitely far away from the reflector. Environment mapped reflections are incorrect for objects close to the reflector. Approximating these objects with billboards or depth images [26, 27] improves reflection accuracy. Using CoNUS depth images as relief textures or to approximate reflected objects brings sampling flexibility without a considerable increase of the cost of ray / depth image intersection.

Focus-plus-context visualization

Focus-plus-context visualization allocates more pixels to dataset regions of greater interest to the user. An important challenge stems from the fact that displays have a uniform pixel resolution (with the exception of special focus-plus-context screens [31]). Consequently the focus-plus-context image cannot be displayed directly and has to be mapped to displays with uniform resolution by introducing distortions between the focus and context regions.

Focus-plus-context visualization is typically applied to 2-D data (e.g. to hierarchies [28], graphs [29], and maps [30]). Applying the technique to 3-D data can be done either by distorting the dataset and then visualizing it with a conventional camera [32, 33], or by distorting the camera model [7, 34]. FPC focus-plus-context visualization falls in the second category. The advantages of the FPC approach are flexibility in defining the regions of interest which allows multiple, dynamic, and overlapping foci.

Figure 8. Mass-spring system for sampling map definition. The user loads the system interactively with a circular brush (yellow).

3 FPC Coherent Non-Uniform Sampling

The FPC CoNUS pipeline comprises three main stages:

- Flexible pinhole camera model construction (Section 3.1).
- FPC CoNUS image rendering (Section 3.2).
- CoNUS image use (Sections 4, 5, and 6).

3. 1 FPC model construction

The FPC model is defined by a viewpoint (i.e. eye), and image plane, and a sampling map. What is needed is a sampling map that provides a flexible, compact, and fast encoding of the variable sampling rate over the image plane. These requirements are met by starting with a regular 2-D mesh and distorting it such that image regions where higher sampling rates are desired are covered by larger quads. The sampling map is a 2-D array of 2-D locations that define the perturbed mesh nodes.

The sampling map defines a straight forward distortion operation that maps conventional image locations—i.e. undistorted locations—to CoNUS image locations—i.e. distorted locations. Given an undistorted location (u, v) and a sampling map SM the corresponding distorted location (u_d, v_d) is found by looking up the sampling map at location (u, v) using bilinear interpolation.

$$(u_d, v_d) = BilinearInterpolationLookUp(SM, u, v) \quad (1)$$

The sampling rate variation induced by the sampling map is piecewise bilinear. Four neighboring sampling map locations distort all the samples that land within the rectangle they define. The resolution of the sampling map defines the granularity with which the sampling rate variation is controlled. The sampling map has a low resolution compared to the resulting CoNUS image, thus it only implies a small additional storage cost.

We construct sampling maps in one of three ways. One way is through the use of an interactive physics-based 2-D mass-spring system. The image is covered with regularly distributed particles connected with springs to form a quadrilateral mesh. All particles have the same mass and all springs have the same resting length. The resting length is set to 10% of the initial particle distance. The user perturbs the system interactively by adding repulsive forces between particles using a circular brush (Figure 8). The force magnitude decreases from the center towards the periphery of the

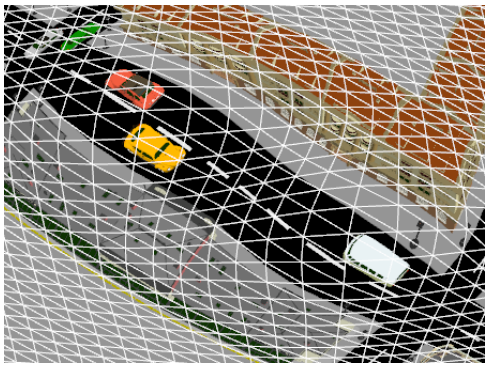

Figure 9. Sampling map with 2 overlapping focus regions obtained by composing two instances of a sampling map.

brush exponentially. The differential equations are solved using Euler's method. A mesh of 256x256 particles is updated at 30 fps and a stable state is reached in less than 2s. The sampling map is defined by the final position of the particles, and it can have a lower resolution than the particle mesh.

Sampling maps can be also generated by translating, scaling, and adding existing sampling maps. The output sampling map location is set by combining the distortion vectors of each input distortion map (Figure 9). A third approach is to do away with the discrete representation and define the distortion analytically (Figure 5, top right).

3. 2 FPC CoNUS image rendering

We discuss the FPC rendering of CoNUS images for various types of data.

Image and height field data

Given an image I, an FPC with a sampling map SM of resolution n x m, and a desired CoNUS image resolution w_c x h_c, a CoNUS image is constructed as follows.

1. Construct a 2-D mesh QM with n x m vertices.
    a. Vertex $v_{ij} = SM(i, j)$.
    b. Texture coordinates $(s, t)_{ij} = (i/n, j/m)$.
2. Render QM at resolution w_c x h_c
    a. Let (s, t) be the texture coordinates at current pixel.
    b. Set color at current pixel by looking up I at (s, t).

QM is defined by the distorted locations stored in SM. The vertices of QM carry the undistorted coordinates as texture coordinates. As QM is rasterized and the texture coordinates are interpolated at each pixel, the undistorted texture coordinates are used to look up the original image.

A CoNUS image should have a lower resolution than the original image. Also the highest sampling rate of the CoNUS image should not exceed the sampling rate of the original image, since at that point data is lost in low sampling rate regions without any gain. A CoNUS height field sampled orthogonally to the base plane is constructed similarly with the exception that the pixel is setup by looking up the depth in the original height field instead of (or in addition to) looking up the color.

Geometry data (i.e. triangle meshes)

Given a triangle mesh T, an FPC with sampling map SM, and a desired view modeled with a conventional planar pinhole camera PPC, a CoNUS image of the triangle mesh is rendered with the FPC by projecting the vertices of T with PPC, by distorting the vertex projections using SM (Equation 1), and by rasterizing the distorted projected mesh conventionally. The projected triangles have to be small enough such that conventional rasterization provides a good approximation of the nonlinear projection induced by the sampling map. Most detailed datasets have small triangles and conventional rasterization is acceptable without further subdivision. When subdivision is needed, an offline approach is preferred in order to circumvent the performance bottleneck of having to issue a large number of primitives in the geometry shader.

Volume data

A CoNUS image is volume rendered by casting CoNUS rays through the volume. Compared to conventional volume rendering the only difference is the non-uniform pattern of rays. In order to determine the CoNUS ray at the current pixel, volume rendering is initiated by rendering the sampling mesh QM as described above in the image data subsection. At each pixel, texture coordinates (s, t) provide undistorted coordinates, which are used to find the image plane point P. The eye and point P define the CoNUS ray at the current pixel, which is traced conventionally.

4 Remote Visualization

The resolution of digital cameras continues to increase faster than network bandwidth. It is also the case that workstation displays now have a lower resolution than the simplest digital cameras attached to cellular phones. For example Apple's 30" LCD has a resolution of 4 mega pixels (MPs), whereas Apple's iPhone 4 has a camera with 5MP resolution. Consequently, even if the image makes it across the network at full resolution, it is most likely going to be downsized for viewing. Often not all pixels in a digital image have the same relevance for the application. For example faces in a portrait photograph are more important than the furnishings in the room. Moreover faces are found automatically by digital cameras for focusing purposes. In the context of an online jewelry store, the pixels capturing the ring, earring, and medallion showcased by a model have higher importance. In the context of an online geographic atlas, pixels sampling famous locations or locations marked by other users as interesting have higher relevance. In the context of remote scientific visualization, some image regions might be of higher interest to scientists.

In such contexts, the CoNUS capability of the FPC could help reduce bandwidth requirements and improve interactivity as follows. The server FPC renders a CoNUS image that samples the regions of interest at a higher rate. Then the CoNUS image is transferred to the client, where a decoding application resamples it to a conventional PPC output frame. The application tours the CoNUS image, showing the regions of interest in detail. The resampling is done by distorting the output image pixel location (Equation 1) and then looking up the CoNUS image at the distorted location.

We have also investigated the use of the FPC CoNUS approach in the context of remote terrain visualization. We have designed a CoNUS height field with a sampling rate that decreases with the distance from the user. The CoNUS height field sampling pattern is shown in Figure 5, top right. The current user location is at the middle of the bottom edge of the CoNUS height field. The CoNUS height field rows cover wider regions and the distance between consecutive rows is larger as the distance from the user increases. The sampling map is defined analytically rather than discreetly. The distortion function we use is the conventional perspective projection *applied to the ground plane*. The perspective projection perturbs sampling locations in the ground plane and *does not* project the 3-D mesh defined by the height field. The distortion function is invertible which allows

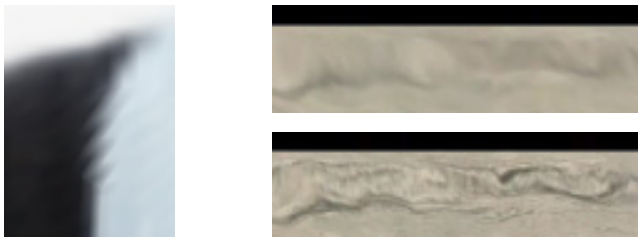

Figure 10. Sampling artifact outside the regions of interest in a frame reconstructed from the CoNUS image in Figure 1 (left), and undersampling of distant mountain by CoNUS height field compared to original height field (right).

undistorting the CoNUS height field at the client. Undistortion creates the 3-D mesh which is rendered for each user view.

Quality The CoNUS image shown in Figure 1 allows rendering all five faces in great detail. The CoNUS height field produces frames that are comparable to frames rendered from the original high-resolution height field (Figure 5).

Performance For Figure 1, once the FPC model is known, rendering the CoNUS image takes negligible time; the FPC sampling map was designed interactively using the spring-mass system. For the example in Figure 5, we use a CoNUS height field of 1k x 1k resolution, which is rendered (at the server) and used (at the client) at over 400 and 100 frames per second, respectively.

Limitations The FPC CoNUS approach increases the sampling rate of the regions of interest at the expense of the rest of the image. When high frequencies are present outside the regions of interest, the undersampling can become noticeable (Figure 10). The approach does not address occlusions. Whereas occlusions do not occur for images or orthogonally sampled height fields, the FPC CoNUS approach will have to be integrated with an occlusion alleviation scheme such as a non-pinhole camera to support remote visualization of general 3-D data.

5 Depth Image Rendering Acceleration

A depth image is a powerful method for approximating geometry: the depth image is computed quickly with the help of graphics hardware, and a depth image can be quickly intersected with a ray. Because of these important advantages depth images have been used to accelerate the rendering of complex effects such as specular reflections, refractions, ambient occlusion, and relief texture mapping. Eliminating the uniform sampling rate constraint of conventional depth images using the FPC CoNUS approach could benefit all these techniques provided that the efficiency of depth image construction and of ray intersection is preserved. CoNUS depth images can be rendered efficiently from height field or geometry data using the FPC as discussed in Section 3.2.

A conventional depth image *DI* is intersected with a ray *r* with the following steps.

- If *r* does not intersect the 3-D bounding box of *DI*, return no intersection; else clip *r* with *DI* to obtain r_c

- Project r_c on the image plane of *DI* to segment r_p

- Rasterize r_p from near to far; at each step check for intersection between the current sub-segment of r_p and the depth image; return first intersection and return no intersection if none is found.

In the case of a CoNUS depth image, the projection r_p of the clipped ray r_c is no longer a line segment, but rather a curve segment. The fundamental advantage of depth images of 1-D intersection with a ray is preserved, at the cost of a slightly more complicated projection of r_c. The ray r_c cannot longer be projected solely by projecting its endpoints. Instead, r_c is first projected conventionally, its conventional projection is subdivided with intermediate (2-D) points, and each intermediate point is distorted using Equation 1. We integrated CoNUS depth images into relief texture mapping and into specular reflection rendering; the only difference is that once the CoNUS depth image is intersected with eye rays and once with reflected rays.

Quality The sampling flexibility afforded by CoNUS depth images allowed improving the clarity of the engraved tablets (Figure 2) and of their reflection (Figure 6).

Performance For both conventional and CoNUS depth images, the performance bottleneck for relief texture mapping and specular reflection rendering is the depth image / ray intersection computation. Intersecting a ray with a CoNUS depth image brings the additional cost of distorting a 2-D point at every step along the ray. However, CoNUS distortion is fast, and we measured an average frame rate penalty of only 5%.

Limitations CoNUS depth images inherit the occlusion limitations of conventional depth images. The sampling tradeoff discussed earlier can lead to visual artifacts outside regions of interest.

6 Focus-Plus-Context Visualization

The FPC CoNUS approach is well suited for focus-plus-context visualization because it offers good control over the sampling rate, which allows precisely designing one or multiple focus regions, and because CoNUS images can be rendered quickly, which supports the interactive change of focus region parameters and dynamic scenes. The CoNUS image is shown directly to the user thus no decoding is needed. The CoNUS image can be rendered efficiently from a variety of data as described in Section 3.2. The only remaining challenge is sampling map construction.

Unlike for the previous applications of CoNUS images, in focus plus-context-visualization the sampling map has to be constructed online, once for every output frame, which precludes the use of the mass-spring approach. We construct sampling maps by composing canonical circular sampling maps, one for every focus region. We demonstrate the approach in the context of volume rendering (Figure 3), where the user manipulates focus region and view parameters interactively to examine a volume dataset, and in the context of a city scene modeled with triangle meshes (Figure 7), where focus regions track moving cars. The focus region location is computed by projecting the center of the tracked car in the output view.

Quality The CoNUS approach enables high quality focus-plus-context visualization for a variety of data types. The focus regions have strong magnification and low distortion. Focus region parameters can change and focus regions can merge and then separate again without abrupt changes in the output visualization. Focus plus context visualization is particularly robust to undersampling outside the focus region—users are likely to focus on the region that they themselves selected as important, and focus regions can be shifted interactively to visualize any region in more detail.

Performance In our experiments FPC volume rendering was on average 7% slower than conventional volume rendering. The cost of volume rendering by ray casting is dominated by the traversal of the volume, thus computing the perturbed rays for the CoNUS approach has no impact on performance. We attribute the slight performance decrease to a larger output image footprint for the distorted volume, and to more rays being focused on the center of the dataset where volume traversal distances are longer. The vertex distortion performed when rendering CoNUS images from triangle meshes had no measurable performance impact.

Limitations Since the CoNUS approach does not alleviate occlusions, tracked objects of interest can become hidden and the user has to change the view to reveal the object. As future work we will examine changing the view automatically to keep the tracked object visible.

7 Conclusions and future work

We have presented a general method for removing the uniform sampling rate constraint of conventional images. CoNUS images can be rendered efficiently from image, height field, geometric, and volume data. Like a conventional image, a CoNUS image has a single layer and good pixel to pixel coherence, thus conventional image compression algorithms can be readily applied. The underlying sampling map can be constructed in a variety of ways including using a mass spring system, by composing multiple input sampling maps, and analytically.

Possible directions for future work include exploring other uses of CoNUS images (e.g. geometric simplification, acceleration of additional rendering effects), investigating the benefit/cost tradeoff of higher order interpolation of the sampling map to achieve C^1 sampling rate continuity, and developing automatic sampling map constructors. This paper describes *how* to sample at a non-uniform rate. We are particularly interested in tightly coupling the FPC CoNUS approach with automatic techniques for determining *what* to sample in more detail, such as automatic geometric complexity analysis, object recognition, eye tracking, and saliency maps.

We foresee that FPC-rendered CoNUS images will have wide applicability as they are compatible with virtually all contexts where images are used.

References

[1] J.-D. Gascuel et al. Fast Non-Linear Projections using Graphics Hardware. *ACM I3D*, pp. 107-114, 2008.

[2] J. Yu, L McMillan. General Linear Cameras. *ECCV'04,* 14-27.

[3] C. Mei, V. Popescu, and E. Sacks. The Occlusion Camera. Eurographics '05, CG Forum. vol. 24, issue 3, 2005.

[4] L. McMillan and G. Bishop. Plenoptic Modeling: an Image-Based Rendering System. SIGGRAPH '95, pp. 39-46, 1995.

[5] G. Johnson, J. Lee, C. A. Burns, W. R. Mark. The Irregular Z-Buffer. ACM ToG, Volume 24, Issue 4, pp. 1462-1482, 2005.

[6] T. Aila and S. Laine. Alias-Free Shadow Maps. Eurographics Symposium on Rendering, pp. 161-166, 2004.

[7] Popescu, V., et al. The General Pinhole Camera. IEEE TVCG, vol. 16, no. 5, 2010, pp. 777-790.

[8] L Balmelli, G. Taubin, F. Bernardini. Space-Optimized Texture Maps. Eurographics 2002, pp 411-420.

[9] Ramanarayanan, G., Bala, K., and Walter, B. Feature-Based Textures. EGSR 2004, pp 265-274.

[10] Sen., P. Silhouette Maps for Improved Texture Magnification. Eurographics Graphics Hardware, 2004.

[11] Sen, P, Cammarano, M., Hanrahan, P. Shadow Silhouette Maps. ACM Transactions on Graphics, Vol. 22, No. 3, July 2003.

[12] L. Lippert, M. H. Gross, and C. Kurmann. Compression Domain Volume Rendering for Distributed Environments. Computer Graphics Forum, 16(3):C95-C107, 1997.

[13] M. Isenburg, et al. Streaming Compression of Triangle Meshes. Symposium on Geometry Processing, pp. 111-118, 2005.

[14] Y. Livnat, S.G. Parker, C.R. Johnson. Fast Isosurface Extraction Methods for Large Image Data Sets. In Handbook of Medical Imaging, Academic Press, pp. 731--745, 2000.

[15] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. Topology-Based Simplification for Feature Extraction from 3-D Scalar Fields. IEEE Vis. '05, pp. 272-280, 2005.

[16] S. P. Callahan, J. L. D. Comba, P. Shirley, and C. T. Silva. Interactive Rendering of Large Unstructured Grids using Dynamic Level-of-Detail. IEEE Vis.'05, pp. 199-206, 2005.

[17] S. Stegmaier, M. Magallón, and T. Ertl. A Generic Solution for Hardware-Accelerated Remote Visualization. In EG/IEEE TCVG Symp. on Data Visualization'02, pp. 87-94, 2002.

[18] E. J. Luke, C. D. Hansen. Semotus Visum: a Flexible Remote Visualization Network. In Proceedings of the IEEE Conference on Visualization '02, pp. 61-68, 2002.

[19] F. Policarpo and M. Oliveira. Relief Mapping of Non-Height-Field Surface Details. In ACM I3D 2006, pp. 55-62.

[20] F. Policarpo, M. Oliveira, and J. Comba. Real-Time Relief Mapping on Arbitrary Polygonal Surfaces. In ACM Symp. on Interactive 3-D Graphics and Games 2005, pp. 155-162.

[21] T. Whitted. An improved illumination model for shaded display. Comm. of the ACM (1980), 23, 6, pp. 343-349.

[22] E. Ofek and A. Rappoport. Interactive reflections on curved objects. In Proc. of SIGGRAPH '98, ACM Press, 333-342.

[23] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The Lumigraph. In Proc. of SIGG. 96, pp.43-54.

[24] M. Levoy, and P. Hanrahan. Light Field Rendering. Proc. of SIGGRAPH 96, pp. 31-42.

[25] J. Blinn and M. Newell. Texture and Reflection in Computer Generated Images. CACM 19:10, pp. 542-547, 1976.

[26] L. Szirmay-Kalos et al. Approximate Ray-Tracing on the GPU with Distance Impostors. Comp. Graphics Forum, 24(3), 2005. pp. 171-176.

[27] V. Popescu, C. Mei, J. Dauble, and E. Sacks. Reflected-Scene Impostors for Realistic Reflections at Interactive Rates. Computer Graphics Forum, (25):3 (EG 2006), pp. 313-322.

[28] J. Lamping and R. Rao. The Hyperbolic Browser: A focus + context technique for visualizing large hierarchies. Journal of Visual Languages and Computing, 7(1):33–35, 1996.

[29] N. Wong, M. S. T. Carpendale, and S. Greenberg. EdgeLens: An interactive method for managing edge congestion in graphs. IEEE InfoVIS 2003, pp. 51–58, 2003.

[30] E. Pietriga and C. Appert. Sigma Lenses: Focus-Context Transitions Combining Space, Time, and Translucence. In Proceedings of 26th CHI'08, pp. 1343-1352, 2008.

[31] P. Baudisch, N. Good, and P. Stewart. Focus Plus Context Screens: Combining Display Technology with Visualization Techniques. In Proceedings of UIST 2001, pp. 31-40, 2001.

[32] M. S. T. Carpendale, D. J. Cowperthwaite, and D. F. Fracchia. Distortion Viewing Techniques for 3-Dimensional Data. IEEE Symposium on Info. Visualization '96, pp. 46-, 1996.

[33] N. Elmqvist. BalloonProbe: Reducing Occlusion in 3D using Interactive Space Distortion. ACM Symposium on Virtual Reality Software and Technology 2005, pp. 134-137, 2005.

[34] L. Wang, Y. Zhao, K. Mueller, and A. Kaufman. The Magic Volume Lens: An Interactive Focus+Context Technique for Volume Rendering. IEEE Visualization '05, pp. 367-374, 2005.