

Spherical Q^2 -tree for Sampling Dynamic Environment Sequences

Liang Wan[†] Tien-Tsin Wong[†] Chi-Sing Leung[‡]

The Chinese University of Hong Kong[†] City University of Hong Kong[‡]

Abstract

Previous methods in environment map sampling seldom consider a sequence of dynamic environment maps. The generated sampling patterns of the sequence may not maintain the temporal illumination consistency and result in choppy animation. In this paper, we propose a novel approach, spherical Q^2 -tree, to address this consistency problem. The local adaptive nature of the proposed method suppresses the abrupt change in the generated sampling patterns over time, hence ensures a smooth and consistent illumination. By partitioning the spherical surface with simple curvilinear equations, we construct a quadrilateral-based quadtree over the sphere. This Q^2 -tree allows us to adaptively sample the environment based on an importance metric and generates low-discrepancy sampling patterns. No time-consuming relaxation is required. The sampling patterns of a dynamic sequence are rapidly generated by making use of the summed area table and exploiting the coherence of consecutive frames. From our experiments, the rendering quality of our sampling pattern for a static environment map is comparable to previous methods. However, our method produces smooth and consistent animation for a sequence of dynamic environment maps, even the number of samples is kept constant over time.

1. Introduction

Realistic rendering of objects illuminated by distant environment maps is useful in many applications. For objects with Lambertian and semi-glossy BRDFs, efficient rendering can be achieved by approximating the illumination with a set of directional lights obtained by sampling the environment map [CD01, KK03, ARBJ03, ODJ04, PH04]. Such approximation can account for shadowing and highlights. However, existing methods mostly assume the environment map is static. With the recent advancement in video acquisition and display technologies [KUWS03, SHS*04], dynamic HDR panoramic video becomes more accessible. Dynamic change (e.g., an area light source moves around, the outdoor illumination changes from dawn to midday, or the scene is set on fire) between the consecutive environment maps requires the regeneration of samples. Regeneration is not just time-consuming, but more importantly, it may not ensure the consistency between consecutive sets of samples.

Figure 10 demonstrates the temporal illumination inconsistency of consecutive sampling patterns and the corresponding frames.

In this paper, we propose a novel hierarchical quadrilateral-based sampling method for dynamic sequences of distant environment maps. Figure 1 shows four snapshots from the illumination-consistent rendering sequence obtained by the proposed method. We first introduce a sphere-to-rectangle mapping and partitioning scheme that effectively maps the sphere to a rectangular structure with equal solid-angle properties (Section 3). With such partitioning, we can construct a quadrilateral-based quadtree on sphere, instead of triangle-based quadtree on icosahedron subdivision [Fek90, SS95]. The quadrilateral, other than the triangular element, allows us to use summed area table [Cro84] for fast integration over a region. We call this spherical representation the *spherical Q^2 -tree*, with one Q referring to the quadrilateral shape of all elements and the other Q referring to the quadruple topology in the tree structure. It adaptively and deterministically subdivides, hence samples, the environment maps, and no relaxation is needed (Section 4).

[†] {lw, ttwong}@cse.cuhk.edu.hk

[‡] eeleungc@cityu.edu.hk




Figure 1: Snapshots from the illumination-consistent rendering sequence obtained by the proposed spherical \mathcal{Q} -tree sampling. A honey bee is illuminated by a dynamic sequence of distant environment maps. The sequence shows the Grace Cathedral environment map on fire. Note the shadow underneath the honey bee is unchanged as it is mainly contributed by the unchanged illumination from the window area in the Grace environment map. The proposed method also preserves the vigorous change of shadow behind the honey bee, caused by the illumination variation of the fire. The fire circle (shown as the unrolled cylindrical panorama) below each image shows the progress of firing.

The hierarchical structure of the method allows efficient sampling of a sequence of dynamic environment maps by exploiting the frame correlation (Section 5). More importantly, its local adaptive nature suppresses the abrupt change in the sampling pattern (both spatial and intensity distributions) due to the variation of total illumination over time. Such suppression is achieved even the number of samples is kept constant throughout the whole sequence. In Section 6, we compare our approach with previous methods for both static and dynamic environments in terms of rendering quality, temporal illumination consistency, and sampling speed. Finally, conclusions are drawn in Section 7.

2. Related Work

Environment mapping [BN76, Gre86] has been intensively used in many real-time applications due to its simplicity and ability in approximating complex illumination. Recent works [CON99] [KM00] [KVHS00] [RH01] [RH02] further improve its realism and efficiency. To realistically and efficiently compute the reflected radiance from an object in a distant environment, Monte-Carlo integration [Rub81] can be employed. Assuming the object is non-emissive and non-refractive, the reflected radiance is computed as,

$$I(x, \vec{s}) = \int_{\Omega} L_{in}(\vec{\omega}) \rho(x, \vec{\omega}, \vec{s}) v(x, \vec{\omega}) (\vec{\omega} \cdot \vec{n}) d\vec{\omega}, \quad (1)$$

where the integration integrates over the hemisphere Ω of directions above x ; $\vec{\omega}$, \vec{s} and \vec{n} refer to the incident direction, viewing direction, and surface normal respectively; ρ is the BRDF; L_{in} is the incident radiance referring to the environment map; and v is the visibility, modeled as a binary function. The key is to efficiently sample the spherical functions inside the integral.

It is possible to use multiple importance sampling [VG95] for the purpose of variance reduction. Alternatively, we can sample the environment map, $L_{in}(\vec{\omega})$ [KK03, ARBJ03, ODJ04, MSV03, PH04, GM00]. This environment map sam-

pling can be regarded as approximating the illumination of a distant environment with N directional lights.

Naïve Monte-Carlo sampling generates uniformly distributed random samples with a convergence rate of $1/\sqrt{N}$, where N is the number of samples. To improve the convergence rate by variance reduction, we can use importance sampling or stratified sampling. Importance sampling samples the function according to its probability density function [ODJ04]. However, importance sampling may oversample the bright region and may not preserve the overall energy contribution of the distant environment. On the other hand, stratified sampling [KK03, CD01] divides the illumination sphere into strata and positions a sample in each stratum. Energy conservation can be done by collapsing the stratum to a directional source. However, stratified techniques usually employ relaxation or iterative schemes, like k-means clustering [CD01] or Lloyd’s relaxation [KK03]. Hence, generation of samples could be time-consuming. Agarwal et al. [ARBJ03] proposed a hybrid metric that combines the importance sampling and stratified sampling.

Existing methods mainly focus on sampling a static environment map. When they are applied to a sequence of dynamic environment maps, temporal illumination consistency may not be maintained as some of their computation rely on the total illumination of the environment which may change over time.

3. Sphere-to-Rectangle Mapping and Partitioning

Astrophysicists map allsky images to rectangular images for signal processing, analysis, and compression. Among these mapping schemes, Hierarchical Equal Area isoLatitude Pixelisation (HEALPix) [GHW98] has been designed to represent the skydome for cosmology study. It subdivides the spherical surface into 12 base spherical quadrilaterals (“quad” from now on) as shown in Figure 2, each with the equal surface area, i.e., equal solid-angle. Recursive subdivision of quads generates smaller sub-quads while maintain-

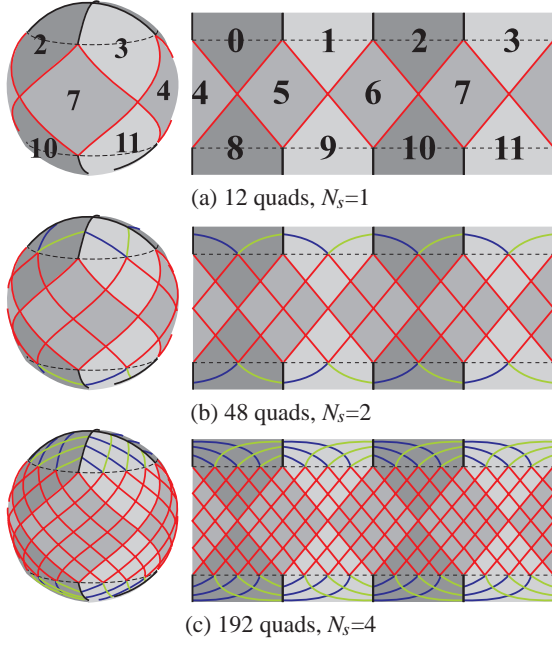


Figure 2: Recursive subdivision of HEALPix. The right column unrolls the sphere into the space of $[\cos\theta, \phi]$, where (θ, ϕ) is the spherical coordinate.

ing the solid-angle equality. We found that its associated sampling pattern, formed by positioning a sample in each quad, has a nice property in spherical sampling, namely low-discrepancy.

Several spherical representations [Fek90, SS95] make use of spherical tessellation that recursively tessellates a base polyhedron such as an octahedron or icosahedron [Kle56], and projects vertices onto the spherical surface. The base faces of these polyhedra are usually triangular. Instead of tessellating a base polyhedron, HEALPix directly performs partitioning on the spherical surface. The goal is to partition the sphere into multiple equal-area quads. The first-level partitioning produces 12 quads (Figures 2(a) and 3(c)). Recursive subdivision of the quads guarantees to obtain equal-area sub-quads. To do so, the sphere is first divided into equatorial and polar zones (Figures 3(a) & (b)). Zone enclosed by “arctic” and “antarctic” circles at $z = \pm 2/3$ is called the equatorial zone while the rest are the polar zones.

Within the equatorial zone ($|\cos\theta| \leq 2/3$), the sphere is partitioned by simple curvilinear equations in the form of

$$-\frac{1}{2} + t + \frac{3}{4}z = \frac{k}{N_s}, \quad k = 0, 1, \dots, (4N_s - 1) \quad (2)$$

and

$$-\frac{1}{2} + t - \frac{3}{4}z = \frac{l}{N_s}, \quad l = 0, 1, \dots, (4N_s - 1) \quad (3)$$

where $z = \cos\theta$ and $t = \phi/\frac{\pi}{2}$ for shorthand; N_s is the number of intervals along the side of each base quad. Parameter N_s




Figure 3: HEALPix partitioning. (a) & (b): equatorial and polar zones in 2D side and 3D views. θ and ϕ are the polar and azimuth angles respectively. (c): the 12 equal solid-angle base quads produced from first-level partitioning.

controls the level of subdivision and must be a positive integer. With this setting, indices k and l indicate the coordinate of a sub-quad. These partitioning curves are color-coded in red in Figures 2 and 3(c). They are designed to be linear functions of z and ϕ , hence they form straight lines when projected onto the lateral surface of an enclosing cylinder ($[\cos\theta, \phi]$ space) as shown on the right hand side of Figure 2. As these partition curves are equally spaced, they ensure equal-area property in the space of $[\cos\theta, \phi]$ as well as on the spherical surface, thanks to the Archimedes’ Hat-Box Theorem [CR89]. On the other hand, the polar zones ($(|\cos\theta| > 2/3)$) are partitioned by a set of curves:

$$\phi = \frac{\pi}{2}k, \quad k = 0, 1, 2, 3 \quad (4)$$

indicated as black curves in Figure 3(c),

$$\sqrt{3(1-z)} \cdot t = \frac{k}{N_s}, \quad k = 1, \dots, (N_s - 1) \quad (5)$$

indicated as green curves in Figure 2, and

$$\sqrt{3(1-z)} \cdot (1-t) = \frac{l}{N_s}, \quad l = 1, \dots, (N_s - 1) \quad (6)$$

indicated as blue curves in Figure 2. Note that the first level of partitioning in the polar zone involves Equation (4) only. Equations (5) and (6) take place when the tessellation continues to the next level. The solid-angle equality in these zones can be proved by integrating the area of quads bounded by Equations (5) and (6). Further note that the boundaries are *not geodesic*.

The base quads can be further partitioned into sub-quads as illustrated in Figure 2 by increasing N_s . If N_s is chosen to be 2^i where i is a non-negative integer, we obtain a hierarchical partitioning, while high-resolution partitioning contains the partitioning boundaries of the low-resolution ones. The total number of quads on sphere is $12N_s^2$.

3.1. Discrepancy

Since there is no way to tell which part of the environment contributes the most during rendering, Monte Carlo sampling requires the sampling pattern to be unbiased. Therefore, if the environment map is an image of constant radiance, samples should be uniformly distributed over the

sphere. Mathematicians measure the uniformity using discrepancy. Shirley [Shi91] and Dobkin et al. [DEM96] applied the discrepancy to analyze sampling techniques used in rendering. Early discrepancy measurements are mainly designed for planar structure. For uniformity of a point set on sphere, generalized discrepancy [CF97] can be employed. It is defined as,

$$D(N) = \frac{1}{2\sqrt{\pi}N} \left[\sum_{i,j=1}^N \left(1 - 2\ln \left(1 + \sqrt{\frac{1 - \vec{\eta}_i \cdot \vec{\eta}_j}{2}} \right) \right) \right]^{\frac{1}{2}} \quad (7)$$

where $\{\vec{\eta}_1, \dots, \vec{\eta}_N\}$ is a N -point sequence and each $\vec{\eta}_i$ is a point on sphere. The lower the D is, the more uniformly distributed the sampling pattern is.




Figure 4: Discrepancy of various sampling patterns. In general, the discrepancy decreases as the number of samples, N , increases.

Figure 4 shows the discrepancy $D(N)$ of the five sampling patterns, where N is the total number of samples. Among the tested methods, equal lattice (in longitude-latitude domain) have much higher discrepancies than the others. The discrepancy of equal lattice keeps almost the same when N increases. Although the cubemap representation is much better than the equal lattice, it is still worse than the rest. Its discrepancy decreases very slowly and soon saturates as N increases. It can be shown that HEALPix is the best as its discrepancy is the lowest (the most uniformly distributed) for all N among all methods. It is better than icosahedron triangulation [Kle56], and even better than the quasi-Monte Carlo point set, Hammersley (base 2) [HS64, WLH97], that has been referred by mathematicians as low-discrepancy [CF97].

4. Hierarchical Quadrilaterals for Structured Importance Sampling

If we subdivide the base quad with $N_s = 2^i$, we naturally construct a spherical quadtree with i as the level of subdivision. Unlike the previous icosahedron-based spherical quadtree [Fek90] with each element being triangle, our spherical quadtree is constructed by a hierarchy of quads on sphere. This is why we called our representation the *spherical Q^2 -tree*. Our Q^2 -tree is more similar to the classical quadtree on planar structure [War69]. Techniques that are applicable on planar quadtree can be easily adapted to our spherical Q^2 -tree without much modification. Moreover, the

HEALPix partitioning guarantees that quads on the same level are equal in solid angle. Hence, the solid angle of quad in any level can be trivially calculated given i and significantly simplifies our computation in stratification. Note that icosahedron subdivision does not guarantee the solid-angle equality.

The adaptive subdivision, and hence sampling, starts with a two-level tree containing 12 leaf nodes on the second level, one for each base quad (Figure 5(a)). Strictly speaking, the tree is actually a forest of 12 Q^2 -trees. The quad (leaf node) with the highest importance is then selected and subdivided into four equal solid-angle sub-quads (four child nodes). The subdivision process continues on leaf nodes until the desired number of quads, or samples, is reached. Figure 5(a) shows one example subdivision with quads on different levels color-coded differently. The spherical surface is “unrolled” to a minimally-distorted *fish-view*, named due to its shape, on the left hand side of Figure 5(a) instead of the highly distorted $[\cos \theta, \phi]$ or $[\theta, \phi]$ domains.

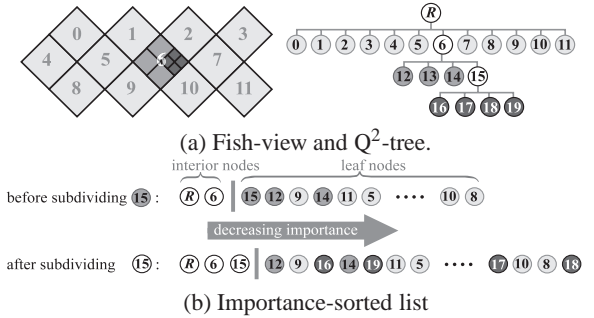


Figure 5: (a) The adaptive subdivision and its corresponding Q^2 -tree. The numbers on the base quads correspond to that in Figure 2(a). (b) This sorted list offers a fast location of the most important (leftmost) leaf node. In this example, node 15 is located from the sorted list and then subdivided. Its children, nodes 16-19, are created and inserted into the Q^2 -tree and the list according to their importances.

During the subdivision, we also maintain a list of indices, each points to one interior/leaf node in the Q^2 -tree, to efficiently locate the desired leaf node. It is sorted according to the quad importance with the more important nodes on the left. Therefore, at any time, the list contains all interior nodes on the left and all leaf nodes on the right. At each iteration, we first pick the leftmost leaf node l from the sorted list, subdivide its corresponding quad into four sub-quads, and create 4 new leaf nodes under the node l . Next, the importances of these four sub-quads are computed and their indices are inserted into the sorted list according to their importances. The node l now becomes the rightmost interior node in the list as illustrated by an example in Figure 5(b).

4.1. Importance Metric

During the sampling of distant environment maps, we want to sample more at brighter regions while sample less at




Figure 6: Three subdivision steps of the spherical Q^2 -tree based on the importance metric.

darker regions. The proposed Q^2 -tree subdivision copes with the previous importance metrics used in previous stratification approaches [CD01, ARBJ03].

Agarwal et al. [ARBJ03] combined the stratified sampling and illumination-based importance sampling by proposing a hybrid importance metric $L^a \Delta\omega^b$ to describe a region Ω' , where L is the integrated illumination, given by

$$L = \int_{\Omega'} L_{in}(\vec{\omega}) d\vec{\omega}; \quad (8)$$

$\Delta\omega$ is the solid angle of Ω' ; and parameters a and b are non-negative constants. A large value of a favors the illumination component while a large value of b favors the solid angle component. In this paper, we adopt their setting, given by

$$L\Delta\omega^{1/4}. \quad (9)$$

Figure 6 shows 3 consecutive steps of a Q^2 -tree subdivision based on this importance metric. It shows how our subdivision scheme copes with the importance metric, as it can effectively locate small and medium bright spots.

With the proposed method, we can partition an environment map into quadrilateral strata without any time-consuming relaxation. The subdivision is locally adaptive and deterministic. It seems that the computation of the importance metric is time-consuming. However, the importance computation of any quad can be computed in constant time. Due to the equal solid-angle property, a quad on the i -th level (base quad is level 0) can be directly computed as,

$$\Delta\omega = \frac{\pi}{3 \cdot 4^i} \quad (10)$$

Moreover, thanks to the rectangular structure of

HEALPix, we compute the sum of illumination, L , in constant time with the aid of the classical texture prefiltering technique, summed area tables [Cro84]. We can trivially apply the original summed area table on 2D plane to the unrolled spherical domain. By preparing a summed area table for the sum of radiances T , the sum of illumination of a quad can be computed as

$$L = T(i_1, j_1) - T(i_0, j_1) - T(i_1, j_0) + T(i_0, j_0) \quad (11)$$

where $T(i, j)$ looks up the element (i, j) in T that holds the sum of illumination of a rectangular region spanning from the lower-left corner of the base quad to coordinate (i, j) ; (i_0, j_0) and (i_1, j_1) are the lower-left and upper-right corners of the quad respectively. Note that the application of summed area table is not available for icosahedron-based spherical quadtree [Fek90], nor the fast computation of solid angle, because the triangular elements on the same subdivision level may have different solid angles.

Once the Q^2 -tree subdivision is completed, we can collapse all the pixels within a quad, or stratum, to a directional light source. This approach preserves the overall illumination contribution of the distant environment. Instead of positioning the light at the centroid of the strata, we jitter deterministically the light position around the centroid to avoid regularity. Figure 7 shows the sampling patterns of four environment maps obtained by our method. Note how the method adaptively samples the environment according to the radiance distribution.

5. Dynamic Sampling

So far, we have discussed how to sample a static environment map. For a dynamic sequence of environment maps obtained from HDR video capturing devices, one can sample each frame individually to obtain a sequence of sampling patterns. However, the difference between two consecutive environment maps is usually small. Individual sampling is wasteful as the computation of consecutive frames are mostly the same. More seriously, if the temporal consistency is not maintained, the sampling patterns of consecutive frames may be substantially different and may cause to abrupt change in the animation of rendered results. Figure 10 illustrates such inconsistency. The existing non-deterministic sampling method, such as Lloyd's relaxation, is undesirable for dynamic sequence as the sampling pattern may not be reproducible even for the same environment.

Due to the local adaptive nature of quadtree, the proposed Q^2 -tree sampling can be used for exploiting the correlation and suppressing the inconsistency between consecutive environment frames. Recall that a Q^2 -tree and a sorted list are obtained after partitioning an environment map. To exploit the correlation between two consecutive frames, we construct the Q^2 -tree of the latter frame from that of the former. Firstly, importances of all interior and leaf nodes are

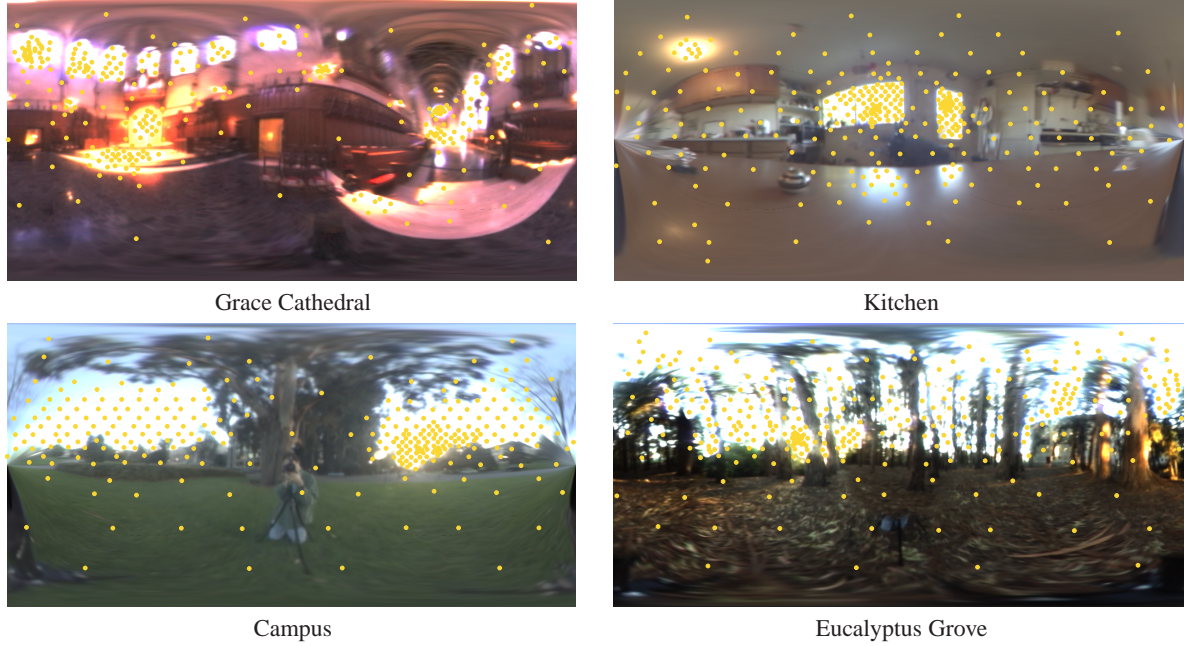


Figure 7: Sampling patterns on various environment maps. All patterns contain 300 points.

re-computed with the new radiance values. Secondly, the interior and leaf nodes in the list are sorted separately. Obviously, the resulting list does not ensure nodes are sorted accordingly. At the boundary of interior and leaf nodes, some interior nodes may have a smaller importance than that of leaf nodes. Such “incorrectness” tells us exactly how the Q^2 -tree should be adjusted in order to reflect the illumination contribution of the new frame.

To adjust the Q^2 -tree and correct the sorting order of the list, we perform two primitive operations, *merge* and *split*. An operation *merge* picks the rightmost interior node and its corresponding children, combine them into one single leaf node, and insert the new leaf node into the leaf part of the list accordingly (the middle row in Figure 8). Hence, each merge operation reduces three leaf nodes. On the other hand, operation *split* converts the leftmost leaf node to an interior node and creates four child leaf nodes (Figure 8, bottom row). The new interior and leaf nodes are inserted into the interior and leaf parts of the list according to their current importances. Each *split* operation creates three more leaf nodes. By performing one *merge* and one *split*, we maintain the total number of leaf nodes as constant. The *merge-and-split* iteration continues until the sorting order is recovered. Note that the resulting pattern converges to the pattern produced by sampling from scratch.

We may, on the other hand, terminate the iteration before completion. Then some temporal variations are not taken into account, which may cause “delay” or “temporal blurs” in the illumination changes. Such blurring can also somehow reduce the inconsistency. A simple termination criterion is to




Figure 8: Operations *merge* and *split* are performed iteratively until the sorting order is corrected.

set a tolerance on the importance difference. When the importance difference between the leftmost leaf node and the rightmost interior node is smaller than the tolerance, we can stop the iteration.

6. Results and Discussion

To evaluate the proposed method, we compare it to previous methods. We divide the comparison into two parts, one for static environment map and the other for a dynamic sequence of environment maps. Figure 9 compares four methods in terms of sampling pattern and rendering effects. These methods include LightGen [CD01], structured importance sampling [ARBJ03], Penrose-based importance sampling [ODJ04], and the proposed spherical Q^2 -tree sampling. To generate the control image (bottommost row of Figure 9), we use 100k uniformly distributed samples. The peak signal-to-ratio (PSNR) of the rendered images, with re-




Figure 9: Static environment comparison of various methods. Three hundreds samples are used for all tested cases.

spect to the control images, are then measured and listed in the images accordingly.

In terms of sampling patterns, our pattern is comparable to that of structured importance sampling as both methods combine stratification and importance-based sampling. In terms of visual quality of the rendered images, our results are very close to the control images and comparable to that of structured and Penrose-based importance sampling methods. In terms of PSNR, our method is the first runner-up.

The strength of our method is its ability in maintaining illumination consistency over time when a sequence of HDR environment maps is sampled. This sequence can be either synthesized or captured by the latest HDR video camera. Such consistency is important in smooth video animation. Otherwise, there may be noticeable “jumping” of highlight and/or shadow in the rendered sequence.

To setup the experiment, we construct a synthetic fire sequence with the Grace Cathedral environment map as the

background. The fire starts at a point and propagates anti-clockwise, as viewed from the top, until a fire ring is formed (top-left of Figure 10). The fire (bottom parts in Figure 1) has continuous flame jumping within a small extent. Three sampling methods (structured, Penrose-based, and ours) are tested to generate the 300-sample pattern for each frame in the fire sequence of environment maps. We compare the consistency among the consecutive frames in these three sequences (Figure 10). [Readers are referred to the companion [video](#) for a much apparent comparison as static pictures may not be obvious to illustrate the dynamic problem.]

Figure 10 shows the structured importance sampling on the second row, Penrose-based importance sampling on the third row, and the spherical Q^2 -tree method on the bottom-most row. We pick 3 consecutive frames, 89, 90 & 91, from the rendered sequences for detailed comparison. Moreover, we blow up (orange boxes in the environment map) the sam-

pling patterns and place them besides each frame for better visualization. In these 3 frames, the curtain behind the honey bee blocks most of the fire. The total illumination contributed by the visible fire is also quite steady. Hence, the shadow and highlight should not change abruptly.

The rendering result of the structured important sampling (the second row) exhibits choppy change in shadowing underneath the honey bee. This change can be explained by the sudden shift of samples in upper part of the environment map, window area of the Cathedral. As these samples own high radiance values, their move unavoidably introduces “jumping” of shadows. For the Penrose-based importance sampling (the third row), there seems to be sudden changes in illumination between two consecutive frames. In fact, there is a steady increase in radiance contribution due to the fire, especially as viewed from the honey bee. The sudden changes in appearance can be accounted by shift of samples at the contributing regions from frame 89 to 90, and sudden appearance of samples in right part of the environment map from frame 90 to 91. On the other hand, the proposed method maintains a smooth and consistent illumination. There is no sudden change in shadow nor highlight. This can be explained as the sampling pattern (both spatial distribution and intensity distribution of samples) on the upper part of the environment map is not affected by the fire in the lower part of the environment map. The number of samples within the fire region is gradually increased while the illumination variation due to the “dancing” of fire is still preserved in the supplementary video.

We analyze the three methods and figure out the reasons that may cause temporal inconsistency. Structured importance sampling computes the hierarchical threshold based on the total illumination contribution. As the fire introduces different amount of extra energy into the environment, it inevitably affects the computed thresholds over time and hence the sampling patterns. Similarly, the change of total illumination over time also affects the global information computation in the Penrose-based importance sampling. A local change in the environment map may *globally* affect both the spatial and intensity distributions of the generated pattern. On the other hand, the local adaptive nature of spherical Q^2 -tree suppresses the change in sampling pattern (both spatial and intensity distributions) when the environment gradually changes over the time.

We implemented the spherical Q^2 -tree sampling on a PC with Pentium IV 2.6GHz CPU. For generating 300 samples from individual environment maps (no correlation among these maps), it takes 0.70 second on average. For an environment map sequence, the average time to generate one 300-sample pattern from a frame is only 0.41 second as we exploit the frame correlation. Note that our code is not yet optimized. On the same machine, Penrose-based importance sampling (implementation of Ostromoukhov et al.) takes 0.07 second to generate one 300-sample pattern if the scal-

ing factor is given. If the scaling factor is not provided, it takes 1.53 seconds on average.

Note that we currently fix the number of samples for the sequence of environment maps. For the Penrose-based importance sampling, the abrupt change may be reduced by fixing the scaling factor. However, the total number of samples may vary when the total illumination due to the environment varies. It is also possible to adapt some existing methods to the sequence of environment maps, but some may have the difficulty due to their dependence on the global information.

Limitations One limitation of the proposed method is that the number of samples used must not be too small. Since our method starts with 12 base quads, the minimum number of samples used must be at least 12. Small amount of samples, say less than 20, may not allow our method to adequately represent the environment map containing several small but bright luminaries. Moreover, the final number of samples generated by our method may slightly deviate from the desired number of samples specified by users as each subdivision must generate three more samples.

7. Conclusion

Previous environment map sampling methods lay samples based on the total radiance contribution of the environment. However, such global nature of computation may harm the illumination consistency in a dynamic environment. The local adaptive nature of the proposed spherical Q^2 -tree sampling scheme effectively suppresses the abrupt change in sampling pattern when there is variation in the overall illumination contribution. Hence, a smooth and consistent rendering sequence can be obtained. The sampling pattern generated by the proposed method is low-discrepancy (uniformly distributed). It achieves comparable visual rendering quality, high PSNR statistics, and high-speed performance.

Currently, we fix the number of samples for every frame in the sequence. Whether a variable sampling number should be used needs further investigation. When a small number of samples is used for an environment map with scattered bright spots, it may result in the loss of some shadow. It could also be worth for exploring our sampling method in Monte Carlo rendering system when there is dynamic geometry change. Readers can find more information about the work at the following homepage

<http://www.cse.cuhk.edu.hk/~ttwong/papers/q2tree/q2tree.html>

8. Acknowledgment

The authors would like to thank all anonymous reviewers for their constructive comments and suggestions. Thanks to Gary Ho, Guang-Yu Wang, and Lin Shi for helping out in the experiments, Louis Wang for designing 3D scenes and preparing the video, and Dr. Peter Clark for the narration in the video. We would also like to thank other people in graphics group of CUHK for their valuable discussions and suggestions. The honey bee model is from The

Singularity and Sparky. This project is supported by the Research Grants Council of the Hong Kong Special Administrative Region, under RGC Earmarked Grants (Project No. CUHK 4189/03E) and a research grant from City University of Hong Kong (Project No. 7001703).

References

- [ARBJ03] AGARWAL S., RAMAMOORTHY R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. *ACM Trans. Graph. (SIGGRAPH 2003)* 22, 3 (2003), 605–612. 1, 2, 5, 6
- [BN76] BLINN J. F., NEWELL M. E.: Texture and reflection in computer generated images. *Communications of the ACM* 19, 10 (1976), 542–546. 2
- [CD01] COHEN J., DEBEVEC P.: LightGen HDRShop plugin, 2001. 1, 2, 5, 6
- [CF97] CUI J. J., FREEDEN W.: Equidistribution on the sphere. *SIAM Journal on Scientific Computing* 18, 2 (1997), 595–609. 4
- [CON99] CABRAL B., OLANO M., NEMEC P.: Reflection space image based rendering. In *Proceedings of SIGGRAPH 1999* (1999), pp. 165–170. 2
- [CR89] CUNDY H., ROLLETT A.: Sphere and cylinder–Archimedes’ theorem. In *Mathematical Models*, 3rd ed. Tarquin Pub., 1989, ch. 4.3.4, pp. 172–173. 3
- [Cro84] CROW F. C.: Summed-area tables for texture mapping. In *Proceedings of SIGGRAPH 1984* (1984), pp. 207–212. 1, 5
- [DEM96] DOBKIN D. P., EPPSTEIN D., MITCHELL D. P.: Computing the discrepancy with applications to supersampling patterns. *ACM Transaction on Graphics* 15, 4 (1996), 354–376. 4
- [Fek90] FEKETE G.: Rendering and managing spherical data with sphere quadtrees. In *Proceedings of the 1st conference on Visualization '90* (1990), pp. 176–186. 1, 3, 4, 5
- [GHW98] GÓRSKI K. M., HIVON E., WANDEL T. B. D.: Analysis issues for large CMB data sets. In *Proceedings of the MPA/ESO Conference on Evolution of Large-Scale Structure: from Recombination to Garching* (August 1998). 2
- [GM00] GIBSON S., MURTA A.: Interactive rendering with real-world illumination. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, 2000), Springer-Verlag, pp. 365–376. 2
- [Gre86] GREENE N.: Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications* 6, 11 (1986). 2
- [HS64] HALTON J. H., SMITH G. B.: Radical-inverse quasi-random point. *Communications of the ACM* 7, 12 (1964), 701–702. 4
- [KK03] KOLLIG T., KELLER A.: Efficient illumination by high dynamic range images. In *In Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering* (2003), pp. 45–51. 1, 2
- [Kle56] KLEIN F.: *Lectures on the Icosahedron and the Solution of Equations of the Fifth Degree*. Dover, New York, 1956. 3, 4
- [KM00] KAUTZ J., MCCOOL M.: Approximation of glossy reflection with prefiltered environment maps. In *Graphics Interface* (2000), pp. 119–126. 2
- [KUWS03] KANG S. B., UYTENDAELE M., WINDER S., SZELISKI R.: High dynamic range video. *ACM Trans. Graph. (SIGGRAPH 2003)* 22, 3 (2003), 319–325. 1
- [KVHS00] KAUTZ J., VÁZQUEZ P.-P., HEIDRICH W., SEIDEL H.-P.: Unified approach to prefiltered environment maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (2000), pp. 185–196. 2
- [MSV03] MADSEN C. B., SØRENSEN M. K. D., VITTRUP M.: Estimating positions and radiances of a small number of light sources for real-time image-based lighting. In *Proceedings: Annual Conference of the European Association for Computer Graphics, EUROGRAPHICS 2003, Granada, Spain* (September 2003), pp. 37–44. 2
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P. M.: Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph. (SIGGRAPH 2004)* 23, 3 (2004), 488–495. 1, 2, 6
- [PH04] PHARR M., HUMPHREYS G.: *Physically-based rendering: from theory to implementation*. Morgan Kaufmann, 2004. 1, 2
- [RH01] RAMAMOORTHY R., HANRAHAN P.: An efficient representation for irradiance environment maps. In *Proceedings of SIGGRAPH 2001* (2001), pp. 497–500. 2
- [RH02] RAMAMOORTHY R., HANRAHAN P.: Frequency space environment map rendering. In *Proceedings of SIGGRAPH 2002* (2002), pp. 517–526. 2
- [Rub81] RUBINSTEIN R. Y.: *Simulation and the Monte Carlo Method*. John Wiley & Sons, New York, 1981. 2
- [Shi91] SHIRLEY P.: Discrepancy as a quality measure for sample distributions. In *Proceedings of Eurographics '91* (1991), pp. 183–193. 4
- [SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: High dynamic range display systems. *ACM Trans. Graph. (SIGGRAPH 2004)* 23, 3 (2004), 760–768. 1
- [SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of SIGGRAPH 1995* (1995), pp. 161–172. 1, 3
- [VG95] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of SIGGRAPH 1995* (1995), pp. 419–428. 2
- [War69] WARNOCK J. E.: *A hidden surface algorithm for computer generated halftone pictures*. PhD thesis, University of Utah, 1969. 4
- [WLH97] WONG T.-T., LUK W.-S., HENG P.-A.: Sampling with Hammersley and Halton points. *Journal of Graphics Tools* 2, 2 (1997), 9–24. 4

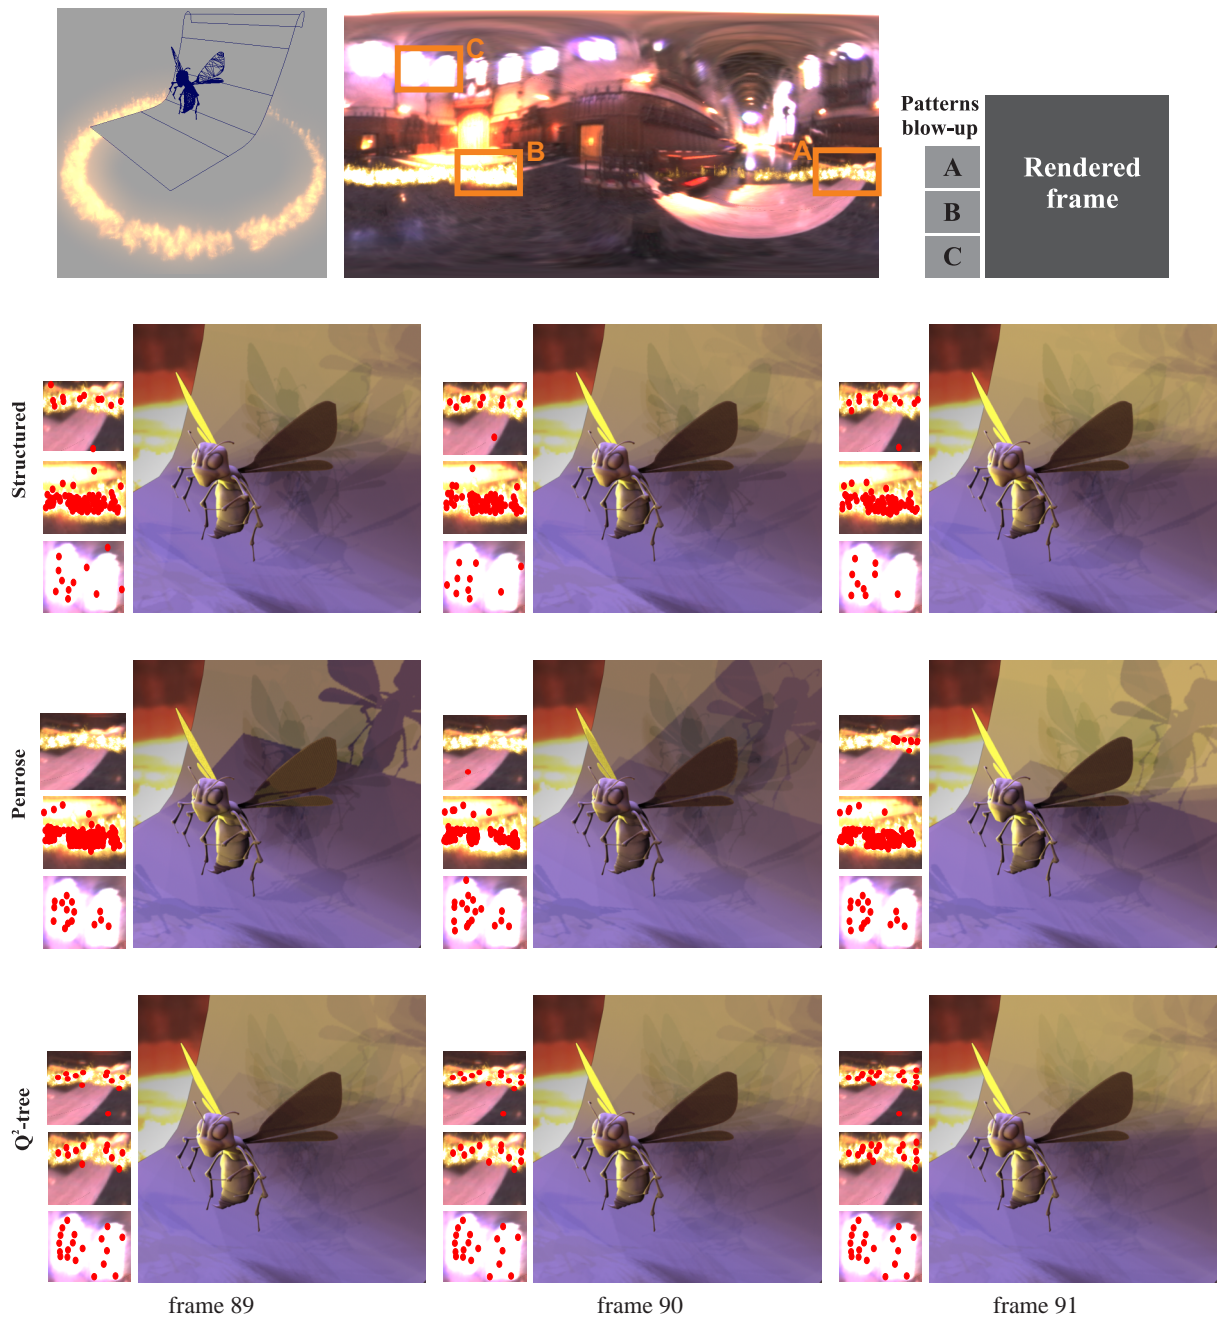



Figure 10: The first row: An illustration of the setup in the fire sequence is shown on the left. In the middle, the boxed regions in the HDR environment Grace Cathedral are blown up and placed beside each rendered frame. The second row: The sampling patterns and rendering result of frames 89, 90 and 91 are generated by structured importance sampling. Note the choppy “jump” of shadow underneath the honey. The third row: Results from the Penrose-based importance sampling. The sudden change in appearance between successive frames is obvious. The bottommost row: Results from the proposed Q^2 -tree. The illumination of the honey bee is only gradually changed across the three consecutive frames. See also the color plate for the color version. [Readers are referred to the companion video for more apparent comparison of three animated sequences and extra examples.]

COLOR PLATE: 1/2 page



Color Figure 10: Comparison of structured, Penrose and Q^2 -tree sampling methods in terms of temporal consistency.