

Radiance Transfer Biclustering for Real-time All-frequency Bi-scale Rendering

Xin Sun^{*,†}, Qiming Hou[‡], Zhong Ren[†], Kun Zhou^{*}, Baining Guo^{†,‡}

^{*} State Key Laboratory of CAD&CG, Zhejiang University

[†] Microsoft Research Asia

[‡] Tsinghua University

Abstract—We present a real-time algorithm to render all-frequency radiance transfer at both macro-scale and meso-scale. At a meso-scale, the shading is computed on a per-pixel basis by integrating the product of the local incident radiance and a bidirectional texture function. While at a macro-scale, the precomputed transfer matrix, which transfers the global incident radiance to the local incident radiance at each vertex, is losslessly compressed by a novel biclustering technique. The biclustering is directly applied on the radiance transfer represented in a pixel basis, on which the BTF is naturally defined. It exploits the coherence in the transfer matrix and a property of matrix element values to reduce both storage and runtime computation cost. Our new algorithm renders at real-time frame rates realistic materials and shadows under all-frequency direct environment lighting. Comparisons show that our algorithm is able to generate images that compare favorably with reference ray tracing results, and has obvious advantages over alternative methods in storage and preprocessing time.

Index Terms—Illumination, Rendering, Shadow Algorithm, Graphics Hardware.

I. INTRODUCTION

Synthesizing realistic images requires faithful simulation of the interactions of light and matter at all relevant scales, which is very challenging due to the large representation gaps between these scales. Many existing rendering techniques are limited to effects at a single scale.

Based on *precomputed radiance transfer* [2], or PRT, Sloan et al. propose bi-scale radiance transfer [3] to model radiance transfer at two scales. The main idea is to decompose the radiance transfer function into a global part, which is represented in a macro-scale and coarsely sampled at each vertex, and a local part, which is represented in a meso-scale, densely sampled at each pixel and mapped over the object. The bi-scale transfer is then combined at run time to render self-shadowing and interreflection effects at both scales. Due to the low-order spherical harmonics(SH) used to represent lighting and

transfer, the method is limited to very blurry shadows and matte materials. Sometimes it is not sufficient to reveal the fine details of the object surface and the relationship between the lighting and different parts of the object, as Fig. 1(a) shows.

PRT has also been extended to all-frequency illumination, by representing the lighting and transfer with wavelets [4]–[7], or spherical radial basis functions [8]. These techniques, however, compute per-vertex shading and would be impractically expensive if the fine details such as in Fig. 1(c) are to be rendered.

We propose a real-time algorithm for all-frequency bi-scale radiance transfer rendering. The algorithm renders hard and soft shadows at both macro-scale and meso-scale(Fig. 1(b)), and supports arbitrary materials.

The main challenge of all-frequency bi-scale radiance transfer is the precomputation and storage of the macro-scale transfer function, which is defined in a 6D space formed by global incident direction, local incident direction, and surface position. To capture all-frequency occlusion and reflection effects, each dimension has to be sampled at a reasonable rate. This makes the transfer function very costly to be precomputed, stored or manipulated at run-time.

We represent the light and transfer in a pixel basis, and compress the macro-scale transfer matrix by a novel biclustering technique, which reduces both the storage and runtime computation cost and enables real-time rendering of all-frequency bi-scale radiance transfer.

A pixel basis is poor in representing smooth functions, as pointed out by Ng et al. [4]. In our case, however, the transfer function is not a smooth one, because the reflectance term is decoupled from the macro-scale transfer function, and the transfer function is determined by visibility, rotation and down-sampling. As a result, the the transfer function is a coarsely discretized spherical function that tends to form blocks(Section IV-A). We exploit this property in biclustering(Section IV-B, IV-C) to losslessly compress the transfer function.

Compared with existing transfer matrix compression




Fig. 1. A Stanford bunny with a bumpy surface illuminated by environment lighting. Note how our all-frequency algorithm more faithfully captures the shadowing effects at both global and local scales, as compared with the low-frequency bi-scale algorithm [1].

techniques such as clustered principal component analysis(CPCA) [1], [6], biclustering takes advantage of specific properties of the transfer matrix, and it also better adapts to different levels of available coherence. As a result, a more efficient compression is achieved, as will be demonstrated in Section VI.

II. RELATED WORK

Local Effects: In computer graphics, local lighting effects had been represented by textures, bidirectional reflectance distribution functions(BRDFs), height fields [9] or bump maps. The bidirectional texture function(BTF) [10], which is a 6D function that encapsulates appearance that varies spatially and with light and view direction, is introduced to represent more general local effects. Hardware based techniques have been developed([11], [12]) for full 6D BTFs rendering. Lower dimensional alternatives have also been introduced([9], [13], [14]) to reduce the storage and rendering cost. BTF capturing([10], [15]), synthesis([12], [16]), and editing tools([17], [18]) have also been proposed, and a wide variety of BTFs are available for representing different kinds of material appearance. See the survey by Müller et al. for a good overview [19].

More general representations include the shell texture function [20], view-dependent displacement map [21], and generalized displacement maps [22]. These representations, though being able to capture more general local effects such as subsurface scattering and fine-scale silhouettes, are also more expensive to store and/or render, and are not as widely used as BTFs.

Recently, Sloan et al. proposed to fit the precomputed transfer function of local features by zonal harmonics, which can then be rotated efficiently at run time for multiplication with the global incident lighting [23], represented in spherical harmonics.

These methods focus on reflectance representation and

rendering of lighting effects caused by local radiance transfer.

Wang et al. render shadows at both scales [24] by introducing a 4D meso-structure distance function(MDF) representation. It is limited to simple lighting sources like point lights, since rendering is based on shadow mapping and accumulation over complex light sources is not affordable for interactive applications.

Global Effects: Realistic rendering of global effects under complex lighting requires integration over the incident directions, which is expensive if conducted directly. Sloan et al. solve the light integration problem by an off-line tabulation of an object's response to low-frequency lighting [2], represented by SH, and turn the run-time integration into a linear combination of the responses to each SH basis lighting. Exploiting the coherence among vertices, compression methods are introduced using PCA [25] or CPCA [1]. We take a different approach for both transfer representation and compression and achieve higher rendering quality at the same storage cost.

Traditional PRT has been extended to fine, repetitive local details. By combining the local and global radiance transfer [3], self-shadowing and interreflection effects at both macro-scale and meso-scale can be rendered in realtime, yet only under low-frequency environment lighting. Our method extends the low-frequency bi-scale transfer to all-frequency, and supports both accurate shadows and complex materials.

PRT has also been extended to handle all-frequency shadows by using a non-linear wavelet approximation [4], albeit with either fixed lighting or viewpoint. This limitation has been overcome by factoring the transfer function into a visibility term and a reflectance term and computing at run time a wavelet triple product of lighting, visibility and reflectance to yield the outgoing radiance [5]. A similar approach can be difficult to apply to our transfer matrix, since unlike the reflectance

function, the rotation from a global coordinate frame to a local coordinate frame for each vertex is defined in 6D space, and is still too expensive to precompute.

Another dimensionality reduction technique is based on decomposition of the BRDF, which has been used in many all-frequency PRT algorithms([6]–[8], [26], [27]). It is not clear how this technique could be used in our case, since the reflectance function has already been decomposed to meso-scale and the rotation term in the transfer function depends on the surface location, making the decomposition infeasible.

Xu et al. [28] exploit the coherence in directions of both lighting, visibility and BRDF, corresponding to the coherence in the rows of our transfer matrix, and enables object rotation in all frequency rendering of dynamic scenes and on-the-fly BRDF editing. Our method exploits coherence in both rows and columns and can tackle transfer matrix of higher dimensionality, which enables bi-scale radiance transfer.

Garg et al. [29] leverage the symmetry and data-sparseness of the transfer matrix and represent it using hierarchical tensors. The 4th-order tensor representing the light transport is recursively divided into 16 children until the node can be represented using rank-1 approximation. The nodes in that hierarchy are sub-blocks formed by neighboring rows and columns. The characteristic of the transfer matrix in our case is different from theirs, which will be shown in Section IV-A, due to the asymmetry of the input and output of the transfer. Thus we take a different approach to compress the transfer matrix.

Recently, an all-frequency shadow algorithm for dynamic scenes has been proposed [30]. Global incident radiance is approximated by a small number of area light sources, and extended convolution shadow maps [31] are used to render the shadows generated by each of these light sources. No precomputation is needed for the visibility and dynamic scene objects can be supported. The algorithm, however, focuses on shadow rendering instead of local reflectance. Integration of the BRDF across the light source domain is not supported, and the BRDF in the direction of the center of each area light is evaluated to weight the contribution. It is not clear if complex materials such as shown in this paper can be supported.

Integration of meso-scale surface details into the PRT framework has recently been shown possible [32] by exploiting a clustered piecewise constant representation for the visibility and lighting. Precomputation of this representation for the 4D BRDF, however, is already prohibitively expensive so that it is evaluated dynamically at the cluster centers. In our approach, the global radiance

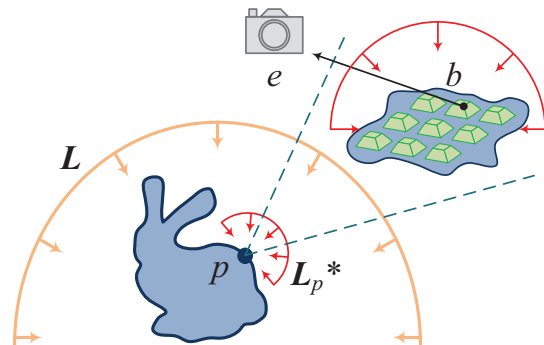


Fig. 2. Transfer the global incident radiance L to the local coordinate frame of p , yielding the local incident radiance L_p^* . Per-pixel shading is then computed by integrating the product of BTF and L_p^* .

transfer matrix are compressed instead of approximated, and the local response encoded in the 6D BTF is precisely rendered.

Biclustering: Biclustering is a technique which has been extensively used in biological data analysis to find submatrices where the genes exhibit highly correlated activities for every condition. For more details, see a recent survey by Madeira and Oliveira [33]. We apply the idea to transfer matrix compression and develop a novel algorithm for bicluster construction and real-time rendering.

III. OVERVIEW AND TERMINOLOGY

We use math italics for scalars, 3d points or vectors (e.g., x , p), boldface italics for higher dimensional vectors (e.g., L), and sans serif for matrices or biclusters(e.g., T , B). A submatrix is noted as $T[j_1, \dots, j_r; k_1, \dots, k_l]$, where j_1, \dots, j_r are the row indices and k_1, \dots, k_l are the column indices.

As in low-frequency bi-scale rendering [3], we decouple the coarse (PRT) and fine (BTF) radiance transfer in precomputation and combine the transfer at both scales at run-time. In a preprocess, we compute and store the matrix that transfers the source lighting, represented in the global coordinate frame, into the local frame of each vertex. Both visibility and rotation are incorporated into the transfer matrix, and macro-scale shadowing effects are captured.

As depicted in Fig. 2, the global incident radiance function L is transferred to each vertex by a multiplication with the transfer matrix:

$$L_p^* = T_p L \quad (1)$$

We use the vector form(such as L) to indicate a spherical function(such as $L(\omega_i)$) represented in any basis defined on the sphere. And “*” is used to emphasize that L_p^* is defined in the local coordinate frame of p . T_p is the

transfer matrix sampled at p , which will be explained in more detail below.

It is sometimes more convenient to use the compact form of Eq. (1), which simply packs the L_p^* vectors for all vertices in the scene into a single vector L^* :

$$L^* = \mathbb{T}L \quad (2)$$

The meso-scale transfer is computed per pixel to capture the finer variations of the meso-structures on object’s surface, as shown in Fig. 2. More precisely, for each pixel x , shading is computed as

$$B(x, e) = \sum_d b(u, e, d)L^*(x, d) = \mathbf{b}(u, e) \cdot L^*(x) \quad (3)$$


where $B(x, e)$ is the outgoing radiance, b is the 6D BTF, u is a 2D texture coordinate and e is the view direction, d is the direction of the incident radiance. Note that e and d are both defined in the local coordinate frame of the point corresponding to x . $L^*(x, d)$ is the local incident radiance at pixel x . The last part of Eq. (3) is a reformulation in vector form. The per-pixel local incident radiance vector $L^*(x)$ can be obtained by interpolating L_p^* defined at each vertex.

We choose the pixel basis to represent both lighting and transfer. In other words, each component of L is simply the sampled incident radiance value of the corresponding pixel on the direction cubemap.


Note that Eq. (3) requires $L^*(x, d)$ to be sampled at the same set of directions as $b(u, e, d)$. Assuming the BTF is sampled at M incident radiance directions, $L^*(x)$ is M -dimensional, and so is L_p^* . The light source $L(\omega_i)$, on the other hand, is often discretized in a higher resolution (N sample directions, $N > M$). The vertex transfer matrix T_p is a $M \times N$ transfer matrix, with its element (j, k) representing the contribution of the global incident radiance at direction k to the local incident radiance at direction j . For a scene with n vertices, the T is $Mn \times N$.

The transfer matrix T is very large, due to the high-dimensional nature of the transfer function. For example, for a moderate scene configuration with 15k vertices, $6 \times 32 \times 32$ source lighting directions and 8×8 BTF lighting sample directions, T has nearly 6 billion entries. Though sparse, it is still quite cumbersome to store or relight with.

To compress T and accelerate run-time multiplication, we decompose it into a number of submatrices, in a way that minimizes the computation cost of Eq. (2). The decomposition, which will be described in detail in the next section, is conducted as a preprocess and is referred to as the biclustering of the transfer matrix.



(a) A small patch and a sample vertex, together with a particular local incident direction



(b) Transfer matrix and the corresponding row

Fig. 3. An example of transfer matrix construction. A small patch of 76 vertices is used to show the process.

IV. PRECOMPUTATION

In a preprocess, we construct the transfer matrix T_p at each vertex to obtain the overall transfer matrix T , on which a greedy iterative search is performed to yield a compact bicluster representation.

A. Transfer Matrix Construction

For a given local incident direction d_j at vertex p , the transfer matrix element $(T_p)_{jk}$ represents the contribution of the global incident direction d_k . In other words, the element indicates how much of the global incident radiance at d_k would reach the local incident radiance direction d_j .

The construction of the transfer matrix is depicted in Fig. 3. We rasterize a high resolution visibility hemi-cube at each vertex p , and down-sample it to the same resolution as the environment lighting. Each pixel of the cubemap corresponds to an incident radiance direction. These directions are iterated to evaluate their contribution to each local incident radiance direction d_j .

At each incident direction d_k , we first check the visibility, and totally blocked directions are bypassed. We find the nearest local incident radiance direction d_j and set $(T_p)_{jk}$ to the visibility value at d_k .

Note that the local incident radiance directions d_j are defined in p ’s local coordinate frame and to find the

nearest d_j for global incident radiance direction d_k we need to rotate all the local incident radiance directions to the global coordinate frame.

Fig. 3(a) shows a particular local incident radiance direction d_j at vertex p with the contributing cubemap directions shown in red. These contributing directions correspond to elements in the j^{th} row of the transfer matrix. After all vertices are processed, we pack the transfer matrices T_p to obtain the overall transfer matrix. We show the transfer matrix for this example in Fig. 3(b), and the corresponding row for the direction in Fig. 3(a) is highlighted in red.

Due to the down-sampling, the visibility at each pixel can have values other than 0 and 1. Generally, if we perform $2^m \times 2^m$ down-sampling, the number of possible visibility values is $2^{2^m} + 1$.

One alternative representation for the radiance transfer would be to store the down-sampled visibility vectors and the rotation matrix as reshuffling vectors. This, however, swaps the order of down-sampling and rotation, and the low sampling rate of the rotation will introduce noticeable rendering artifacts.

From the small example in Fig. 3, we see two properties of the transfer matrix. First, the transfer matrix is very sparse. Actually, only 0.69% of the elements are non-zero. This can be easily understood: for a given local incident radiance direction, only a small number of directions on the source lighting cubemap may contribute. The sparsity alone, however, still does not make the transfer matrix tractable, since there are still tens of millions of entries to be stored and multiplied with the incident radiance at runtime.

Second, we see regularity in the distribution of the non-zero entries - they tend to form blocks. This regularity is due to the structure in visibility and the distribution of normal directions of vertices. Note that unlike in the transfer matrix between symmetric incident and exitant light field, where the matrix entries tend to form continuous sub-blocks that can be well approximated using rank-1 factors [29], the "blocks" here might be discontinuous, since the transfer is between the incident 2D environment lighting and exitant 4D surface light field and we do not take inter-reflections into account. This kind of regularity inspires us to use a biclustering technique to exploit the coherence in arbitrary rows and columns. Unlike CPCA, which exploits the coherence between elements of fixed size (the rows), we use biclusters, whose size can adapt to the available coherence. This, joined with exploiting the limited number of possible matrix entries, makes it possible for us to develop a technique that provides higher compression rates than existing techniques such as CPCA and wavelet.

B. Bicluster

The most straightforward way to exploit the coherence in the transfer matrix would be to extract sub-matrices having constant entries. This simple clustering approach however tends to yield a large number of very small clusters and we instead seek a more general representation for the coherence, and we found the biclustering technique that is extensively used in biological data analysis is very suitable for our purpose.

We base our algorithm on a very simple observation. Suppose we have a submatrix of T formed by selecting the j_1, j_2, \dots, j_r rows and the k_1, k_2, \dots, k_l columns from T , noted as $T[j_1, j_2, \dots, j_r; k_1, k_2, \dots, k_l]$. If each column of the submatrix is composed of elements of a constant value:

$$T_{j_1 k_m} = T_{j_2 k_m} = \dots = T_{j_r k_m} \triangleq t_m, \forall m \in \{1, 2, \dots, l\} \quad (4)$$

then we have

$$T[j_1, \dots, j_r; k_1, \dots, k_l] L[k_1, \dots, k_l] = \overbrace{\begin{bmatrix} a & \dots & a \end{bmatrix}}^{r \text{ scalars}}^T, \\ a = t_1 L_{k_1} + t_2 L_{k_2} + \dots + t_l L_{k_l}. \quad (5)$$

The submatrix that satisfies Eq. (4) can be represented as a *constant column bicluster*, or simply *bicluster*. And we call such a submatrix a *bicluster submatrix*. A bicluster is determined by the row indices and the column indices of the submatrix, plus the values of each column t_1, t_2, \dots, t_l .

Given a bicluster $B[j_1, \dots, j_r; k_1, \dots, k_l; t_1, \dots, t_l]$ belongs to a $Mn \times N$ transfer matrix, we can define its multiplication with a N vector L

$$(BL)_j = \begin{cases} \sum_{m=1}^l t_m L_{k_m} & \text{if } j \in \{j_1, j_2, \dots, j_r\} \\ 0 & \text{otherwise} \end{cases} \\ j = 1, 2, \dots, Mn, \quad (6)$$

yielding a $M \times n$ vector. Note we only need l floating point multiplications and $l - 1$ additions, as well as r additions to accumulate the resulting vector of Eq. (6). We define the *cost* of a bicluster as

$$W(B) = r + 2l - 1, \quad (7)$$

reflecting the number of FLOPs involved in the bicluster multiplication with B .

By decomposing the transfer matrix into bicluster submatrices, we can obtain its *bicluster representation*. Notice that an 1×1 submatrix of T can be represented by a smallest bicluster described by three scalars. This guarantees the existence of a bicluster representation for any given T . What we need though is an optimal biclustering of the transfer matrix that minimizes the runtime computation cost.

C. Transfer Matrix Biclustering

Given the transfer matrix T , the biclustering of T is a set of biclusters B_i that satisfies

$$TL \equiv \sum_i B_i L, \forall L \quad (8)$$

We want to find the biclustering that minimizes $\sum_i W(B_i)$.

It has been shown that such a biclustering of a matrix is NP-complete [33]. In this section, we describe a greedy iterative algorithm for transfer matrix biclustering.

Qualitatively, when possible, we want to look for larger biclusters, which provide higher reduction in computation and storage. However, looking for submatrices that strictly satisfy Eq. (4) tends to yield very small biclusters, weakening the effectiveness of biclustering. An important observation is that we can decompose any submatrix into a bicluster submatrix and a residual submatrix. Below is an example.

$$\begin{bmatrix} 1 & \frac{1}{2} & \frac{3}{4} & 1 \\ 1 & \frac{3}{4} & \frac{3}{4} & 1 \\ 1 & \frac{1}{2} & \frac{3}{4} & 0 \end{bmatrix} = \begin{bmatrix} 1 & \frac{1}{2} & \frac{3}{4} & 1 \\ 1 & \frac{1}{2} & \frac{3}{4} & 1 \\ 1 & \frac{1}{2} & \frac{3}{4} & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \frac{1}{4} & 0 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

This allows us to find a larger bicluster submatrix, paying the price that the additional non-zero elements in the residual submatrix should be handled in later iterations. But that is a reasonable trade-off as long as we can keep the number of non-zero elements in the residual submatrix small.

We start from a randomly chosen element, which is a 1×1 bicluster submatrix itself, then iteratively try to insert(delete) rows or columns into(from) the current bicluster submatrix. Every insert or delete operation produces a new bicluster submatrix that changes the total computation required by the matrix-vector multiplication(Eq. (8)). We evaluate the changes and select the operation that reduces the computation most.

Operations are performed iteratively to improve the current bicluster submatrix until no more computation reduction is possible. As a bicluster is produced, the corresponding submatrix is subtracted from the transfer matrix. The algorithm terminates when all the remaining elements in the transfer matrix become zeroes.

The most important issue is how to evaluate the computational change that is related with a particular submatrix and its bicluster representation. For a given $r \times l$ submatrix M and its bicluster representation B , the floating point operations related to the submatrix is $2\text{nnz}(M)$, where the $\text{nnz}(M)$ is the number of non-zero elements of M . This accounts for a per-element multiplication and addition that are involved in the

Algorithm 1 Pseudo code for transfer matrix biclustering

Input: transfer matrix T
Output: $\{B_i\}$, a biclustering of T

- 1: $T_c \leftarrow T$.
- 2: **if** $\text{nnz}(T_c) = 0$ **then**
- 3: **terminate**
- 4: **end if**
- 5: Randomly choose a non-zero element $(T_c)_{jk}$
- 6: $T_b \leftarrow T_c[j; k]$ {initialed as an 1×1 matrix}
- 7: $\Delta W_c \leftarrow \Delta W(T_b)$
- 8: **for all** possible operation f **do**
- 9: **if** $\Delta W(f(T_b)) > \Delta W_c$ **then**
- 10: $\Delta W_c \leftarrow \Delta W(f(T_b))$
- 11: $f_c \leftarrow f$
- 12: **end if**
- 13: **end for**
- 14: **if** $\Delta W_c = \Delta W(T_b)$ **then**
- 15: Output T_b 's corresponding bicluster
- 16: $T_c \leftarrow T_c - T_b$
- 17: go to 2
- 18: **else**
- 19: $T_b \leftarrow f_c(T_b)$
- 20: go to 7
- 21: **end if**

matrix-vector multiplication. The operations involved in the corresponding bicluster multiplication is $W(B) = r + 2l - 1$. If a residual submatrix exists, then every non-zero element in the residual submatrix M_R requires another addition and multiplication. Thus, the change in computation cost can be evaluated as

$$\Delta W(M) = 2\text{nnz}(M) - (r + 2l - 1) - 2\text{nnz}(M_R) \quad (9)$$

$\Delta W(M)$ is a measurement of the change of the required FLOPs due to representation transfer. It can be either negative or positive.

We list the pseudo code for transfer matrix biclustering in Algorithm 1. Note that for ‘‘possible operation’’ in line 8, we refer to the insertion of a new row/column to the current bicluster matrix or the deletion of an existing row/column from it. For deletions, the corresponding non-zero elements in the residual submatrix are absorbed. Note that T_b is usually smaller than T_c and the matrix subtraction in line 16 means subtracting the elements of T_b from the corresponding elements of T_c .

V. RENDERING

Rendering is divided into per vertex computation and per pixel computation, corresponding to macro-scale

transfer and meso-scale transfer, respectively.

We take the cubemap representation of the environment lighting as the input light vector \mathbf{L} , perform the bicluster multiplications and accumulate the results to calculate $\sum_i \mathbf{B}_i \mathbf{L}$. Each multiplication is computed via Eq. (6). The resulting incident radiance vector, \mathbf{L}^* , is packed into the vertex buffer.

Per pixel computation is conducted on the GPU according to Eq. (3). In a pixel shader, we take the interpolated vertex data as the per-pixel incident radiance vector \mathbf{L}_i^* . We also obtain the view direction and texture coordinate, which are then used to lookup the BTF vector $\mathbf{b}(u, e)$. Then a simple dot product $\mathbf{b}(u, e) \cdot \mathbf{L}_i^*$ is performed to yield the shading at the corresponding pixel.

Current graphics hardware has a limitation on the number of available pixel input registers, and we cannot process all the vertex data in a single pass. Multi-pass rendering has to be conducted and alpha-blending is used to add up the results of each pass.

VI. IMPLEMENTATION AND RESULT

In this section, we will first discuss some implementation details, and then present the experimental results.

A. Implementation Details

Data Organization: It is important to guarantee that the local incident radiance directions are distributed evenly on the hemisphere. We use a low distortion area preserving parameterization [34] for the incident radiance direction and view direction hemispheres.




Fig. 4. Packing a 6D BTF into a 3D RGBA texture.

As depicted in Fig. 4, we pack BTFs into 3D RGBA textures, each layer of which corresponds to a sample view direction. For each layer, the data is organized with a number of blocks, each of which corresponds

to a sample incident radiance direction. Thus, each of these blocks is a sample of the patch rendered under a particular incident directional light and view direction combination. The interpolation across the block uses built-in hardware, and the interpolation in the view direction is performed manually in the pixel shader. Instead of using simple bilinear interpolation, which leads to ghosting artifacts, we map the view direction back to the direction sphere and retrieve the four adjacent samples for the interpolation. No interpolation is needed for the incident radiance directions, since we directly transfer the global incident radiance to the local sample incident directions.

Precomputation Optimization: Biclustering is the main bottleneck of precomputation. As a simple optimization for Algorithm 1, we maintain an influence value for each row and each column, defined as the influence of the operations on the change of cost in Eq. (9). We organize all row and column operations into a priority queue according to the influence value. At every step we select the operation at the front of the operation queue. And after an operation is conducted, we update the affected operations in the operation queue, and refresh the queue accordingly.

This select-conduct-update process is more efficient than the straightforward implementation, since we do not need to iterate all the operations to select the one that reduces the cost most. This is achieved at the cost of updating the queue after each operation. Fortunately, the updating process is highly local, since each operation only affects a very small number of operations in the queue.

The above process is repeated until the operation at the front of the queue has negative influence on ΔW , when all operations in the queue cannot reduce the overall cost anymore. Note that during the whole process the size of the queue remains as $Mn + N$ - the sum of column number and row number.

For example, for a new column k_0 which is not in the current submatrix $\mathbf{T}_b[j_1, j_2, \dots, j_r; k_1, k_2, \dots, k_l]$. The constant value t_0 is chosen to be the non-zero value that most frequently appears in the rows j_1, j_2, \dots, j_r of column k_0 . Then the influence of the new column to ΔW consists of the cost of the column itself, which corresponds to the $r + 2l - 1$ term in Eq. (9), plus the sum of the influence of each of the elements in the rows j_1, j_2, \dots, j_r , which correspond to the $2\text{nnz}(\mathbf{T}_b) - 2\text{nnz}((\mathbf{T}_b)_R)$ term. For this new column, the cost of the column is -2 , since it causes l to increase by 1. To compute the influence of a particular element, three different cases need to be considered:

- if the element's value is t_0 , it causes $\text{nnz}(\mathbf{T}_b)$ to

increase by 1, and the influence is 2;

- if the element's value is 0, it causes $\text{nnz}((\mathbb{T}_b)_R)$ to increase by 1, and the influence is -2 ;
- if the element's value is a non-zero value other than t_0 , it causes both $\text{nnz}(\mathbb{T}_b)$ and $\text{nnz}((\mathbb{T}_b)_R)$ to increase by 1, and the influence is 0.

If the insertion of this new column is conducted, first, the influence value of the column itself should be changed, since the corresponding operation with the column has changed from insertion to deletion. Second, influence values of all the rows that has a non-zero element in this column, plus the rows that the residuals in the column are located should be changed, according to the above element influence evaluation.

For other operations, including the deletion of an existing column from the current submatrix and the deletion/insertion of rows, the influence value can be derived similarly.

Rendering Optimization: Biclustering multiplication is mostly bounded by memory access, since the computation itself is very simple. To improve the memory access coherence, we first perform all the bicluster multiplications $B_i L$. A runtime bicluster buffer P_B is maintained, each slot of which corresponds to a bicluster. We iterate all the biclusters and calculate a single scalar $\sum_{m=1}^l t_m L_{k_m}$ (Eq. (6)) for each of them. At the end of this iteration we have the result of all required bicluster multiplications in the buffer.

Then we need to accumulate the products to appropriate local incident radiance directions. For each direction(row) j , we store the indices of all the related biclusters in I_j , which are then used to access the bicluster buffer and sum up the products. The pseudo code is listed in Algorithm 2.

We implement the above computation by CUDA [35], and pass the result to the final BTF shading by CUDA-OpenGL interoperability routines.

Another easy optimization is to reuse the local incident radiance vector if only the view direction is changed. In the results presented in Section VI, we will list the fps numbers for lighting and view changes separately.

B. Results and Comparisons

Platform: We implemented our system on a workstation with a 3.00 GHz Quad Xeon CPU, 4GB RAM, and a NVidia GTX 280 graphics card. We use CUDA to compute the incident radiance of all vertices and render the shading of BTF with OpenGL.

Comparison with Low-frequency Bi-scale Rendering: In Fig. 1, we compare our rendering result with low-frequency bi-scale rendering [3]. In the global scale,

Algorithm 2 Pseudo code for computing $\sum_i B_i L$

Input:

$\{B_i\}$, transfer matrix biclustering

L , lighting vector

Output:

L^* , local incident radiance vec-

tor

```

1: for all bicluster index  $i$  do
2:    $(P_B)_i \leftarrow \sum_{m=1}^l t_m L_{k_m} \{ \{t_m\}, \{k_m\} \text{ in } B_i \}$ 
3: end for
4: for  $j = 1$  to  $Mn$  do
5:    $L_j^* = 0$ 
6:   for  $k = 1$  to  $K_j$  do
7:      $\{K_j\}$ : number of biclusters for row  $j$ 
8:      $i \leftarrow (I_j)_k \{I_j\}$ : related bicluster's indices
9:      $L_j^* \leftarrow L_j^* + (P_B)_i$ 
10:  end for
11: end for

```



(a) SH

(b) Biclustering

(c) Reference

Fig. 5. Our result better captures the appearance of glossy materials than its low-frequency counterpart.

we can see a more accurate shadowing effect, especially for the bunny's ears, while in the local scale, a better perception of the bumpy bunny surface is achieved by our result, thanks to the more accurate capturing of the self-shadowing effects of the surface details.

Fig. 5 shows the advantages of our algorithm over low-frequency bi-scale rendering in rendering glossy materials. The material of the chains is modeled by an anisotropic analytic representation([36], [37]). The specularly of the material is better captured in our result, producing a higher shading contrast. In addition, our algorithm better captures the thickness of the armor, thanks to the more faithful rendering of the self-shadowing effects at the meso-scale.

Comparisons with Alternative Methods: One important advantage of our algorithm is the storage efficiency. By exploiting coherence in both rows and columns of the transfer matrix, as well as exploiting the property that the non-zero matrix entries only evaluate to




Fig. 6. Comparisons of different compression methods of global radiance transfer. The original transfer matrix is obtained by pre-computing for each of the 16K vertices a transfer matrix from the 6,144 global incident directions to the 64 BTF incident directions. For CPCA, the rows of the transfer matrix are divided into 32 clusters, and 8 eigen-vectors are generated for each cluster, yielding 36.2MB final data. For wavelet, each row of the matrix is non-linearly approximated by area-weighted selecting the first 12 wavelet basis lights, yielding 37.1MB final data.

a limited number of values, our method provide much higher quality than alternative methods at comparable storage cost.

In Fig. 6 we compared the result generated by our method and CPCA [1] and non-linear wavelet approximation [4], [5]. We choose the parameters of these alternative methods in a way that their storage costs are roughly the same as ours(36.4MB). Severe rendering artifacts can be observed for both of these methods, as seen in Fig. 6(c)(d), while ours(Fig. 6(c)) is almost visually indistinguishable from the reference(Fig. 6(d)).

A comparison of the storage costs of these methods at the same quality would make the comparison more complete. The high cost of CPCA compression(more than 10 hours), however, made such a search of parameters infeasible.

Parameters and Performance: Environment source lighting is represented as a cubemap of resolution $6 \times 32 \times 32$ and the dimension of the light vector is 6,144. For visibility sampling, 2×2 down sampling is performed to yield $6 \times 32 \times 32$ cubemaps.

The resolution of incident directions of BTFs is 8×8 , except for the “Armor” scene, where the resolution of




Fig. 7. A temple scene with a five different BTFs. The eaves molding uses a patch size of 64×64 , while all other parts use patches of 32×32 .

TABLE I
TEST SCENE STATISTICS

Scene	Bunny	Armor	Chair	Temple
Figure	Fig. 1	Fig. 5	Fig. 6	Fig. 7
Resolution	640×480	480×640	640×640	640×480
#Vertices	3,286	10,652	16,101	32,409
#BTFs	1	1	3	5
BTFs size(MB)	113.2	255.2	169.9	226.5
nnz(T) (M)	10.6	21.9	40.5	62.8
MFLOPs	21.2	43.9	81.0	125.6
Reduced to	5.8	14.9	9.0	7.9
Computation%	27.4	33.9	11.1	6.3
Data size(MB)	1,400	9,500	6,300	12,600
Reduced to(MB)	24.2	59.0	36.4	32.9
Data size%	1.7	0.62	0.58	0.26
Precomp.(min)	12	46	33	59
Avg. FPS	97.0/56.1	42.8/24.9	51.8/33.7	61.1/34.3

* Statistics for the 4 test scenes used in our paper. For rendering performance, we give the average frame rates for view/lighting changes, respectively.

12×12 is used since the BTF used for that scene has a larger high frequency component. The resolution of viewing directions is 12×12 . The patch sizes used for the BTFs of all our demos are 64×64 or 32×32 . Other detailed settings of parameters and performance are included in Table I(Note that two performance numbers are measured for each scene, the first one is for fixed lighting and dynamic view, from which the cost of global transfer is excluded, and the second is for dynamic lighting and view). The scene named “Temple”(Fig. 7) is constructed to show our algorithm’s ability to render multiple BTFs. Five BTFs are used for different parts of that scene, while the transfer matrix is defined and compressed for the entire scene.

From the numbers listed in Table I, we can see that the efficiency of biclustering is dependent on the properties of the scene geometry. For scenes mainly consisting of “flat” geometries, such as “Temple”, the coherence in

the transfer matrix is very high, leading to more efficient compression. For curved geometry, as the “Bunny”, the compression is less effective. The cost of transfer matrix multiplication and BTF shading are roughly comparable, and the former is dominated by the gathering phase, in which the products in the runtime bicluster buffer are accumulated to the corresponding incident radiance directions. For BTF shading, the important factors are the number of pixels to be shaded and the coherence of the BTF fetch. Denser BTF tiling leads to lower coherence between adjacent pixels, and in turn lower shading performance.

The biclustering algorithm scales well with the resolution of local incident directions, as shown in Fig. 8. The storage cost of the compressed transfer matrix at the resolution of 32×32 is 58.5MB, which is less than 2.5 times larger than the storage at 8×8 . While this is partially due to the sparsity increase of the transfer matrix, the biclustering algorithm did achieve practical storage cost at higher resolution. As for the rendering performance, the interpolation and BTF shading cost scales roughly quadratically with the local incident direction resolution, and dominates the rendering cost at higher resolution. This can be seen from the diminishing gap between the performance of fixed and moving lighting. Note that from (a) to (b) we observe superquadratic performance drop, mainly due to the large BTF size which lead to worse cache hit rate in BTF data fetches.

We have also tested the scalability with regard to scene complexity by using the same bunny model with different vertex number, and the result are shown in Table II. We obtain steady compression ratio for different vertex numbers. The rendering performance, on the other hand, scales roughly linearly with the vertex number.

It should be noted that our technique is orthogonal to BTF compression techniques ([38]–[40]), since we focus on the compression of the global transfer matrix.

VII. CONCLUSION AND FUTURE WORK

In this paper, we address the problem of rendering all-frequency bi-scale radiance transfer. The main challenge is the large transfer matrix that needs to be stored and manipulated. We propose a lossless compression algorithm based on transfer matrix biclustering. We introduce a novel algorithm for finding the optimized biclustering of the transfer matrix and minimizing runtime computation cost. Our algorithm is able to reduce storage and computational complexity down to 5%-30%, enabling real-time rendering.

For the meso-scale representation of the surface details, though we only focus on BTFs in this paper,

TABLE II
SCALABILITY TEST RESULTS

#Vertices	3,286	12,380	47,960
nnz(T) (M)	10.6	39.2	151.3
MFLOPs	21.2	78.4	302.7
Reduced to	5.8	21.2	80.7
Computation%	27.4	27.0	26.7
Data size(MB)	1,400	4,868	18,859
Reduced to(MB)	24.2	84.1	316.9
Data size%	1.7	1.7	1.7
Precomp.(min)	12	49	279
Avg. FPS	277.9/86.8	115.8/32.3	48.5/11.9

* Bunny models with different numbers of vertices are used to test our algorithm. Note the consistent compression ratios. The rendering performance scales linearly with the vertex number.

our algorithm is not limited to BTF rendering. Other surface detail representations, such as VDM, normal map, spatially invariant or variant BRDF, can all be used. For simpler representations such as normal map or spatially invariant BRDF, our algorithm can handle dynamic surface details, since they are not involved in the precomputation. Fully dynamic environment lighting can also be supported, since we directly use the cubemap representation for the lighting.

As a limitation, our algorithm only handles distant lighting. Also, it is difficult to incorporate global scale inter-reflections, since many possible values will be introduced into the transfer matrix, breaking the prerequisite of our biclustering algorithm. Another limitation is that, as other per-vertex PRT algorithms, we require the models to be reasonably tessellated since the per-pixel local incident lighting is interpolated linearly from vertex data. Significant error can also be introduced by interpolating across vertices that pointing in very different directions, which could be alleviated by applying a proper cease angle as in shading interpolation.

In future work, we are interested in developing mechanisms that enable trade-offs between quality and storage/performance. Currently our method is based on a compression of the radiance transfer matrix, instead of an approximation. A possible way of approximating the transfer matrix is to ignore the residual matrices according to a carefully designed error metric. It is also interesting to improve the algorithm by making use of the knowledge of lightings and BTFs.

REFERENCES

- [1] P.-P. Sloan, J. Hall, J. Hart, and J. Snyder, “Clustered principal components for precomputed radiance transfer,” *ACM Trans.*





Fig. 8. Glossy bumpy Stanford bunny with different resolution of local incident direction. The material used for the bunny is a Phong model with the exponential equals to 128. As in Table I, performance for both fixed lighting and moving lighting is shown. The transfer matrices are compressed to 24.2MB, 38.3MB and 58.5MB, at the cost of 12, 32 and 93 minutes of precomputation, respectively. The cost of computing macro-scale transfer is 7.9ms, 22.4ms and 37.0ms, respectively. The patch size of the BTF is 16×16 .


- Graph.*, vol. 22, no. 3, pp. 382–391, 2003.
- [2] P.-P. Sloan, J. Kautz, and J. Snyder, “Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments,” in *ACM SIGGRAPH*, 2002, pp. 527–536.
- [3] P.-P. Sloan, X. Liu, H.-Y. Shum, and J. Snyder, “Bi-scale radiance transfer,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 370–375, 2003.
- [4] R. Ng, R. Ramamoorthi, and P. Hanrahan, “All-frequency shadows using non-linear wavelet lighting approximation,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 376–381, 2003.
- [5] R. Ng, R. Ramamoorthi, and P. Hanrahan, “Triple product wavelet integrals for all-frequency relighting,” *ACM Trans. Graph.*, vol. 23, no. 3, pp. 477–487, 2004.
- [6] X. Liu, P.-P. Sloan, H.-Y. Shum, and J. Snyder, “All-Frequency Precomputed Radiance Transfer for Glossy Objects,” in *Eurographics Symposium on Rendering*. Norrköping, Sweden: Eurographics Association, 2004, pp. 337–344.
- [7] R. Wang, J. Tran, and D. Luebke, “All-frequency relighting of glossy objects,” *TOG*, vol. 25, no. 2, pp. 293–318, 2006.
- [8] Y.-T. Tsai and Z.-C. Shih, “All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation,” *ACM Trans. Graph.*, vol. 25, no. 3, pp. 967–976, 2006.
- [9] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel, “Illuminating micro geometry based on precomputed visibility,” in *ACM SIGGRAPH*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 455–464.
- [10] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, “Reflectance and texture of real-world surfaces,” *ACM Trans. Graph.*, vol. 18, no. 1, pp. 1–34, 1999.
- [11] K. Daubert, H. P. A. Lensch, W. Heidrich, and H.-P. Seidel, “Efficient cloth modeling and rendering,” in *Eurographics Workshop on Rendering*, London, UK, 2001, pp. 63–70.
- [12] X. Liu, Y. Hu, J. Zhang, X. Tong, B. Guo, and H.-Y. Shum, “Synthesis and rendering of bidirectional texture functions on arbitrary surfaces,” *IEEE Trans. Vis. Comp. Graph.*, vol. 10, no. 3, pp. 278–289, 2004.
- [13] T. Malzbender, D. Gelb, and H. Wolters, “Polynomial texture maps,” in *ACM SIGGRAPH*. New York, NY, USA: ACM, 2001, pp. 519–528.
- [14] M. Ashikhmin and P. Shirley, “Steerable illumination textures,” *ACM Trans. Graph.*, vol. 21, no. 1, pp. 1–19, 2002.
- [15] M. A. O. Vasilescu and D. Terzopoulos, “Tensortextures: Multilinear image-based rendering,” *TOG*, vol. 23, no. 3, pp. 336–342, 2004.
- [16] X. Tong, J. Zhang, L. Liu, X. Wang, B. Guo, and H.-Y. Shum, “Synthesis of bidirectional texture functions on arbitrary surfaces,” in *ACM SIGGRAPH*. New York, NY, USA: ACM, 2002, pp. 665–672.
- [17] K. Zhou, P. Du, L. Wang, J. Shi, B. Guo, and H.-Y. Shum, “Decorating surfaces with bidirectional texture functions,” *IEEE Trans. Vis. Comp. Graph.*, vol. 11, no. 5, pp. 519–528, 2005, member-Yasuyuki Matsushita.
- [18] J. Kautz, S. Boulos, and F. Durand, “Interactive editing and modeling of bidirectional texture functions,” *ACM Trans. Graph.*, vol. 26, no. 3, p. 53, 2007.
- [19] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein, “Acquisition, synthesis, and rendering of bidirectional texture functions,” *Comput. Graph. Forum*, vol. 24, no. 1, pp. 83–109, 2005.
- [20] Y. Chen, X. Tong, J. Wang, S. Lin, B. Guo, and H.-Y. Shum, “Shell texture functions,” *ACM Trans. Graph.*, vol. 23, no. 3, pp. 343–353, 2004.
- [21] L. Wang, X. Wang, X. Tong, S. Lin, S. Hu, B. Guo, and H.-Y. Shum, “View-dependent displacement mapping,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 334–339, 2003.
- [22] X. Wang, X. Tong, S. Lin, S. Hu, B. Guo, and H.-Y. Shum, “Generalized Displacement Maps,” in *Eurographics Symposium on Rendering*. Norrköping, Sweden: Eurographics Association, 2004, pp. 227–233.
- [23] P.-P. Sloan, B. Luna, and J. Snyder, “Local, deformable precomputed radiance transfer,” *ACM Trans. Graph.*, vol. 24, no. 3, pp. 1216–1224, 2005.
- [24] J. Wang, X. Tong, J. Snyder, Y. Chen, B. Guo, and H.-Y. Shum, “Capturing and rendering geometry details for btf-mapped surfaces,” *The Visual Computer*, vol. 21, no. 8-10, pp. 559–568, 2005.
- [25] D. L. James and K. Fatahalian, “Precomputing interactive dynamic deformable scenes,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 879–887, 2003.
- [26] R. Wang, J. Tran, and D. Luebke, “All-Frequency Relighting of Non-Diffuse Objects using Separable BRDF Approximation,” in *Eurographics Symposium on Rendering*. Norrköping, Sweden: Eurographics Association, 2004, pp. 345–354.
- [27] X. Sun, K. Zhou, Y. Chen, S. Lin, J. Shi, and B. Guo, “Interactive relighting with dynamic brdfs,” *ACM Trans. Graph.*, vol. 26, no. 3, p. 27, 2007.
- [28] K. Xu, Y.-T. Jia, H. Fu, S. Hu, and C.-L. Tai, “Spherical piece-

wise constant basis functions for all-frequency precomputed radiance transfer,” *IEEE Trans. Vis. Comp. Graph.*, vol. 14, no. 2, pp. 454–467, 2008.

- [29] G. Garg, E.-V. Talvala, M. Levoy, and H. P. A. Lensch, “Symmetric photography: Exploiting data-sparseness in reflectance fields,” in *Eurographics Symposium on Rendering*. Eurographics Association, 2006, pp. 251–262.
- [30] T. Annen, Z. Dong, T. Mertens, P. Bekaert, H.-P. Seidel, and J. Kautz, “Real-time, all-frequency shadows in dynamic scenes,” *ACM Trans. Graph.*, vol. 27, no. 3, pp. 34:1–34:8, 2008.
- [31] T. Annen, T. Mertens, P. Bekaert, H.-P. Seidel, and J. Kautz, “Convolution shadow maps,” in *Eurographics Symposium on Rendering*, J. Kautz and S. Pattanaik, Eds., vol. 18. Grenoble, France: Eurographics Association, 2007, pp. 51–60.
- [32] E. Cheslack-Postava, R. Wang, O. Akerlund, and F. Pellacini, “Fast, realistic lighting and material design using nonlinear cut approximation,” *ACM Trans. Graph.*, vol. 27, no. 5, pp. 128:1–128:10, 2008.
- [33] S. C. Madeira and A. L. Oliveira, “Biclustering algorithms for biological data analysis: A survey,” *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 01, no. 1, pp. 24–45, 2004.
- [34] P. Shirley and K. Chiu, “A low distortion map between disk and square,” *J. Graph. Tools*, vol. 2, no. 3, pp. 45–52, 1997.
- [35] NVIDIA, “CUDA homepage,” 2007, <http://developer.nvidia.com/object/cuda.html>.
- [36] M. Ashikhmin and P. Shirley, “An anisotropic phong brdf model,” *J. Graph. Tools*, vol. 5, no. 2, pp. 25–32, 2000.
- [37] M. Ashikhmin, S. Premoze, and P. Shirley, “A microfacet-based brdf generator,” in *ACM SIGGRAPH*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 65–74.
- [38] F. Suykens, K. vom Berge, A. Lagae, and P. Dutr, “Interactive rendering with bidirectional texture functions,” *Comput. Graph. Forum*, vol. 22, no. 3, pp. 463–472, 2003.
- [39] M. Sattler, R. Sarlette, and R. Klein, “Efficient and Realistic Visualization of Cloth,” in *Eurographics Symposium on Rendering*. Leuven, Belgium: Eurographics Association, 2003, pp. 167–177.
- [40] J. Meseth, G. Müller, and R. Klein, “Reflectance field based real-time, high-quality rendering of bidirectional texture functions,” *Computers & Graphics*, vol. 28, no. 1, pp. 105–112, 2004.



Xin Sun graduated from Zhejiang University, Hangzhou, China, and received the Bachelor’s degree and Ph.D degree in computer science in 2002 and 2008 respectively. After that, he joined Internet Graphics Group in Microsoft Research Asia as an associate researcher. His research interests lies in real-time global illumination rendering and GPU-based photorealistic rendering.



Qiming Hou received the BS degree in the Academic Talent Program of Tsinghua University (Mainland China) in 2006 and is currently pursuing the PhD degree in computer science at Tsinghua University (Mainland China). His research interests include general purpose processing using graphics processors, compiler techniques, realistic rendering, and interactive rendering. For the past three years, he has been an intern consultant in Microsoft Research Asia.



Zhong Ren received his Ph.D degree in computer science in 2007 from Zhejiang University, Hangzhou, China. He also holds a Master’s degree in mechanical engineering and a Bachelor’s degree in information engineering, both from Zhejiang University. In 2007, he joined Microsoft Research Asia, where he is currently an associate researcher in the Internet Graphics group. His research interests include real-time rendering of soft shadows and participating media, spherical harmonics for real-time rendering, and also GPU-based photorealistic rendering.



Kun Zhou is currently a Cheung Kong Distinguished Professor in the Computer Science Department of Zhejiang University, and a member of the State Key Lab of CAD&CG. Before joining Zhejiang University, he was a Leader Researcher of the graphics group at Microsoft Research Asia. He received his BS degree and PhD degree in computer science from Zhejiang University in 1997 and 2002, respectively. His research interests include shape modelling/editing, texture mapping/synthesis, real-time rendering and GPU parallel computing.



Baining Guo is the Assistant Managing Director of Microsoft Research Asia, where he also serves as the head of the graphics lab. Prior to joining Microsoft in 1999, Dr. Guo was a senior staff researcher with the Microcomputer Research Labs of Intel Corporation in Santa Clara, California. Dr. Guo received Ph.D. and M.S. from Cornell University and B.S. from Beijing University. His research interests include computer graphics and visualization, in the areas of texture and reflectance modeling, texture mapping, translucent surface appearance, real-time rendering, and geometry modeling.