

# Smart Rebinning for the Compression of Concentric Mosaic

Yunnan Wu, Cha Zhang, and Jin Li, *Senior Member, IEEE*

**Abstract**—The concentric mosaic offers a quick solution to construct a virtual copy of a real environment, and to navigate in that virtual environment. However, the huge amount of data associated with the concentric mosaic is a heavy burden for its application. A three-dimensional (3-D) wavelet-based compressor has been proposed in a previous work to compress the concentric mosaic. In this paper, we greatly improve the compression efficiency of the 3-D wavelet coder by a data rearrangement mechanism called “smart rebinning”. The proposed scheme first aligns the concentric mosaic image shots along the horizontal direction and then rebins the shots into multiperspective panoramas. Smart rebinning effectively enhances the correlation in the 3-D data volume, translating the data into a representation that is more suitable for 3-D wavelet compression. Experimental results show that the performance of the 3-D wavelet coder is improved by an average of 4.3 dB with the use of the smart rebinning. The proposed coder outperforms MPEG-2 coding of the concentric mosaic by an average of 3.7 dB.

**Index Terms**—Compression, concentric mosaic, image-based rendering (IBR), multiperspective panorama, rebinning, three-dimensional (3-D) wavelet.

## I. INTRODUCTION

**I**MAGE-BASED RENDERING (IBR) techniques have received much attention in the computer graphics realm for realistic scene/object representation. Instead of referring to the complicated geometric and photometric properties as the conventional model-based rendering does, IBR requires only sampled images to generate high-quality novel views. Furthermore, the rendering speed for an IBR scene is independent of the underlying spatial complexity of the scene, which makes IBR attractive for the modeling of highly complex real environments. The concentric mosaic [1] enables quick construction of a virtual copy of a real environment, and navigation in the virtual environment. By rotating a single camera mounted at the end of a level beam, with the camera pointing outward and shooting images as the beam rotates, a concentric mosaic scene can be quickly constructed. At the time of rendering, we just split the

desired view into vertical ray slits, and reconstruct each slit through similar slits captured during the rotation of the camera.

Though it is easy to create a three-dimensional (3-D) walk-through, the amount of data associated with the concentric mosaic is tremendous. As an example, a concentric mosaic scene from [1] includes 1350 RGB images with resolution  $320 \times 240$  and occupies a total of 297 MB. Efficient compression is thus essential for the application of the concentric mosaic. In [1], a vector quantization approach was employed to compress the concentric mosaic scene with a compression ratio of 12:1. However, the size of the compressed bit stream is still 25 MB, far too large for either storage or transmission. Since the captured concentric mosaic shots are highly correlated, much higher compression ratio should be expected.

Since the data structure of the concentric mosaic can be regarded as a video sequence with slowly panning camera motion, it is natural to apply existing still image/video compression technologies to the concentric mosaic. However, the concentric mosaic bears unique characteristics, which have led to new challenges in compression. First, the concentric mosaic is a one-dimensional (1-D) image array, with highly structured camera motion among images. The cross-frame correlation is stronger than that of a typical video sequence. Second, the distortion tolerance of the concentric mosaic is smaller, because each rendered image of the concentric mosaic is viewed statically, and the human visual system (HVS) is much more sensitive to static distortions than to time-variant distortions. Since a rendered view of the concentric mosaic is formed by a combination of the image rays, certain HVS properties such as spatial and temporal masking may not be used, because neighboring pixels in the concentric mosaic data set may not be rendered as neighboring pixels in the final view. Most importantly, a compressed image bit stream is usually decompressed to get back the original image, a compressed video bit stream is played frame-by-frame; however, a compressed concentric mosaic bit stream should not be fully decompressed and then rendered. In fact, the decompressed concentric mosaic is so large that most hardware today has difficulties handling it. It is therefore essential to maintain the concentric mosaic in the compressed form, and decode only the contents needed to render the current view.

We can classify existing concentric mosaic compression algorithms into two categories. The first category is the reference/prediction-based coder. Such coders, e.g., the reference block coder (RBC) proposed in [14], resemble existing video coding standards such as MPEG-2 and H.263, and use motion compensated prediction. The image shots in the concentric mosaic are classified into anchor and predicted frames. The anchor frame is independently encoded, just as an I-frame in

Manuscript received October 26, 2000; revised June 29, 2001. The associate editor coordinating the review of this paper and approving it for publication was Dr. Hong-Yuan Mark Liao.

Y. Wu was with Microsoft Research China, Beijing, China, and also with the University of Science and Technology of China, Hefei China. He is now with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA.

C. Zhang was with Microsoft Research China, Beijing, China, and also with Tsinghua University, Beijing, China. He is now with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

J. Li is with the Microsoft Research, Redmond, WA 98052 USA (e-mail: jinli@microsoft.com).

Digital Object Identifier 10.1109/TMM.2002.802838.

MPEG-2. The predicted frame is motion compensated with regard to one of its neighboring anchor frames, and only the prediction residue is encoded. The process is similar to P-frame coding in MPEG-2, though the predicted frame in RBC only refers to the anchor frame so that the compressed RBC bit stream can be flexibly accessed. A two-level hierarchical index table is also embedded in the RBC bit stream for random data access. Magner *et al.* [16] proposed a model-aided coder for the compression of the Lightfield, which is a two-dimensional (2-D) image array. Five images, one at the polar and four at the equator, are intra-coded. The other images are predictively encoded with reference to the intra-coded frame.

The second category is the high-dimensional wavelet coders, which explore the cross-frame redundancy via wavelet filtering. Luo *et al.* [7] proposed a 3-D wavelet coder for the compression of the concentric mosaic. The mosaic images are aligned and wavelet filtered both within and across the mosaic images. After that, the wavelet coefficients are split into fixed-size blocks, which are then quantized, entropy encoded, and assembled into the compressed bit stream. Magner *et al.* proposed a four-dimensional (4-D) Haar wavelet coder with set partitioning in hierarchical tree (SPIHT) coefficient coding in [17] to compress the Lightfield. One attractive property of the high-dimensional wavelet coder is its spatial, temporal, and quality scalability. Here, the term scalability means that a high-dimensional wavelet coder can compress the scene into a single bit stream, where multiple subsets of the bit stream can be decoded to generate complete scenes of different spatial resolution/temporal resolution/quality commensurate with the proportion of the bit stream decoded. This is particularly useful in the Internet streaming environment where heterogeneous decoder/network settings prevail. Furthermore, since wavelet coders avoid the recursive loop in the predictive coders, they perform better in an error prone environment, such as a wireless network.

One common challenge with the high-dimensional wavelet compression schemes is that the cross-image wavelet filtering does not achieve efficient energy compaction. In other words, the coherence in the cross-image direction is not strong, resulting in many large high-frequency coefficients in that direction. As a result, the 3-D wavelet coder of the concentric mosaic [7] and the 4D wavelet coder of the Lightfield [17] lag behind in compression performance compared to the predictive concentric mosaic coder [14] and the predictive Lightfield coder [16]. The same is the case if the compression performance of the 3-D wavelet video coder [2]–[5] is compared with that of a state-of-the-art video compression standard such as MPEG-4 or H.263++. In a prediction-based video/concentric mosaic coder, local motion can be specified on a per block basis, thus inter-frame correlation due to the moving object/camera can be efficiently explored, which is very beneficial to the coding performance. However, local motion cannot be easily incorporated into the framework of 3-D wavelet compression. Due to the nature of temporal filtering, each pixel has to engage in one and only one transform. Therefore, a pixel without matching correspondence (e.g., a pixel which is newly covered by a foreground object) still has to be filtered with certain other pixels. Furthermore, tricks such as half-pixel motion compensation cannot be used in a 3-D wavelet coder.

There was work to improve the cross-frame correlation (coherence) before the wavelet filtering. In the 3-D wavelet concentric mosaic codec [7], a panorama alignment module was used to rotationally shift the mosaic panorama. This module is similar to the pan motion compensation incorporated in [2]. The aligned concentric mosaic is then transformed and encoded. The 4D wavelet coder of the Lightfield [17] morphed individual images onto a common texture plane, which was then wavelet filtered. For video compression, Wang *et al.* [3] proposed to register and warp all image frames into a common coordinate system and then apply a 3-D wavelet transform with an arbitrary region of support to the warped volume. To make use of the local block motion, Ohm [4] incorporated block matching and carefully handled the covered/uncovered, connected/unconnected regions. By trading off the invertibility requirement, Tham *et al.* [5] employed a block-based motion registration for low motion sequences without filling the holes caused by individual block motion.

In this paper, a smart rebinning operation is proposed as a novel preprocessing technique for the 3-D wavelet compression of the concentric mosaic. Rather than adapting the compression algorithm or the filter structure to the mosaic image array, we modify the data structure for easy compression by the 3-D wavelet. Conceptually speaking, the proposed scheme improves the inter-frame coherence by explicitly clustering similar contents together. The proposed scheme begins with pairwise alignment of the image shots. Then, the original concentric mosaic scene is rebinned to form multiperspective panoramas. The rearranged data maintain the high correlations inside the image shots; meanwhile greatly improve the correlation across the shots. The rebinned panoramas can be compressed very efficiently (with an average improvement of 4.3 dB) by the 3-D wavelet coder.

This paper is organized as follows. The background for the acquisition and display of the concentric mosaic is provided in Section II. The smart-rebinning operation, its rationale and potential benefits for the 3-D wavelet codec are detailed in Section III. The 3-D wavelet coding and rendering of the rebinned concentric mosaic volume, which is no longer of rectangular region of support, is discussed in Section IV. Experimental results are presented in Section V. Finally, we conclude the paper in Section VI.

## II. BACKGROUND: THE CONCENTRIC MOSAIC

A concentric mosaic scene is captured by mounting a camera at the end of a level beam, and shooting images at regular intervals as the beam rotates. The capturing device is shown in Fig. 1. Let the camera shots taken during the rotation be denoted as  $\mathbf{F}_n = \{f(n, h, v) | h, v\}$ , where  $n$  indexes the camera shot;  $h$  indexes the horizontal position within a shot; and  $v$  indexes the vertical position. Let  $N$  be the total number of camera shots; and  $H$  and  $V$  be the horizontal and vertical resolution of each camera shot, respectively. The entire concentric mosaic database can be treated as a series of camera shots  $\mathbf{F}_n$ , or alternatively be interpreted as a series of panorama mosaics  $\mathbf{P}_h = \{f(n, h, v) | n, v\}$  where each individual mosaic consists of vertical slits at position  $h$  of all camera shots. Three rebinned mosaics at different radii are shown in Fig. 2. Mosaic




Fig. 1. Capturing device of concentric mosaic.

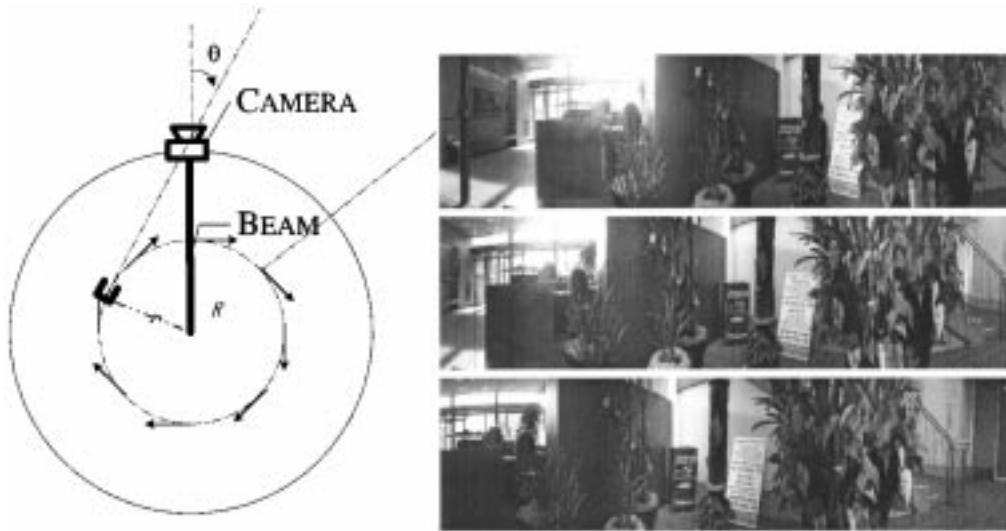


Fig. 2. Concentric mosaic imaging geometry.

$P_h$  can be considered as taken by a virtual slit camera rotating along a circle cocentered with the original beam with a radius  $d = R \sin \theta$ , where  $R$  is the radius of the rotation beam;  $d$  is the equivalent radius of the slit camera; and  $\theta$  is the angle between the tangent of the mosaic  $P_h$  and the camera norm. Since the entire data volume  $P_h, h = 0, \dots, H - 1$  can be considered as a stack of cocentered mosaics with different radii, it is called the *concentric mosaic* [1].

The concentric mosaic is able to capture a realistic environment and render arbitrary views within an inner circle of radius  $r = R \sin(\text{FOV}/2)$ , where FOV is the horizontal field of view of the capturing camera. Rendering the concentric mosaic involves reassembling slits from the captured data set. Let  $P$  be a novel viewpoint and  $AB$  be the field of view to be rendered, as shown in Fig. 3. We split the view into multiple vertical slits, and render each slit independently. A basic hypothesis behind the concentric mosaic rendering is that the intensity of any ray does not change along a straight line unless blocked. Thus, when a slit  $PV$  is rendered, we simply search for the slit  $P'V$  in the captured data set, i.e., either in the captured image set  $F_n$  or equivalently in the mosaic set  $P_h$ , where  $P'$  is the intersection

point between the direction of the ray and the camera track. Because of the discrete sampling, the exact slit  $P'V$  might not be found in the captured data set. The four sampled slits closest to  $P'V$  may be  $P_1V_{11}, P_1V_{12}, P_2V_{21}$ , and  $P_2V_{22}$ , where  $P_1$  and  $P_2$  are the two nearest captured shots;  $P_1V_{11}$  and  $P_1V_{12}$  are the slits closest to  $P_1V$  in the shot  $P_1$ ; and  $P_2V_{21}$  and  $P_2V_{22}$  are closest to  $P_2V$  in the shot  $P_2$ . We may choose only the slit that is closest to  $P'V$  from the above four to approximate the intensity of  $PV$ . Such a scheme is called point sampling interpolation. However, a better approach is to use bilinear interpolation, where all four slits are employed to interpolate the rendered slit  $PV$ . The environmental depth information may assist the search for the best approximating slits and alleviate the vertical distortion. More detailed description of the concentric mosaic rendering may be found in [1].

### III. SMART REBINNING: A CROSS-SHOT DECORRELATION APPROACH

The framework of the smartly rebinned 3-D wavelet concentric mosaic coder can be shown in Fig. 4. Our proposed




Fig. 3. Rendering with concentric mosaic.

smart rebinning technique serves as a preprocessing stage which rearranges the original data volume. The rearranged data volume is then decomposed by 3-D wavelet transform. After that, the wavelet coefficients are cut into cubes, and each cube is quantized and compressed independently into an embedded bit stream. Finally, a global rate-distortion optimizer is used to assemble the bit stream.

This framework is similar to the 3-D wavelet concentric mosaic compression scheme we have proposed in [7]. A key contribution of this paper is an efficient data reorganization strategy that greatly improves the efficiency of cross-shot filtering (equivalent to the temporal direction in video, as there is no time domain in the concentric mosaic), while maintaining about the same filtering efficiency along the other directions.

In recognition of the significant role of motion compensation in the 3-D wavelet compression, we look for an efficient de-correlation scheme along the cross-shot direction. In the previous work [7], the concentric mosaic is aligned through global panning of the mosaic  $P_h$ . However, this approach does not improve the compression efficiency much because the discontinuity inside each mosaic remains there. In this paper, we begin by aligning the image shots and forming a new data volume of nonrectangular region of support. Since the concentric mosaic assumes static scenery and the camera is slowly swinging along a planar circle, the motion between two successive images is predominantly horizontal translation, with little or none vertical motion. We can easily calculate the horizontal translation vector between each pair of consecutive shots. Let  $x_n$  denote the horizontal displacement between shot  $F_n$  and  $F_{n+1}$ . Since the shots are circularly captured, shot  $F_0$  is right next to shot  $F_{N-1}$ . We thus denote  $x_{N-1}$  as the displacement vector between frame  $F_{N-1}$  and  $F_0$ . Note that the horizontal displacement vectors may not be equal for all frames. They are inverse proportional to the distance of the object, i.e., larger for shots with a close-by object, and smaller for shots with the far away background. We can maximize the correlation between neighboring shots by horizontally aligning them according to the calculated displacement vector, as shown in Fig. 5. We term this approach *horizontal shot alignment*. We use seven concentric mosaic image shots  $F_0, F_1, \dots, F_6$  as an example. Each black

region on the horizontal line in Fig. 5 corresponds to one captured image. An additional virtual image  $F_0$  is drawn right after the last image  $F_6$  to show the circular capturing activity of the camera. The gray area contains no data, thus the aligned image shots form a data set of nonrectangular region of support.

After the horizontal shot alignment, the correlation across image shots is expected to improve; however, since the resultant data volume is highly sparse and not rectangular, the compression efficiency may be compromised. Our proposed *smart rebinning* goes beyond the horizontal shot alignment one step further. The idea is to cut and paste (i.e., to rebin) the skewed data set into panoramas by *pushing* the skewed data volume downward in Fig. 5, and form smartly rebinned panoramas. The details of the smart rebinning operation are shown in Fig. 6. The seven image shots are rebinned into five panoramas. The aligned frame boundaries are shown with dashed lines. We divide the original shots into groups of aligned slits according to the horizontal displacement vectors, which are called *stripes*. The stripe is the smallest integral unit in the smart rebinning. Let the stripe be denoted as  $s_{n,j}$ , where  $n$  is the image shot  $F_n$  to which the stripe belongs, and  $j$  indexes the stripe within  $F_n$ . The length of the first stripe  $s_{n,0}$  is  $x_n$ , the horizontal displacement vector between frame  $F_n$  and  $F_{n+1}$ . The length of the  $j$ th stripe  $s_{n,j}$  is  $x_{(n+j) \bmod N}$ , correspondingly. The number of stripes is not constant for all frames; it is inversely proportional to the horizontal displacement vector. Therefore, there are few stripes for the frame with a close-by object (fast moving), and more stripes for that with the faraway background. We then downward stack the stripes and form the rebinned panorama set. We warp the right part of the data volume to the left due to the circular nature of the camera shots. Let the maximum number of stripes for all frames be  $S$ . A total of  $S$  panoramas are obtained with equal horizontal length  $x_0 + x_1 + \dots + x_{N-1}$ . The first rebinned panorama  $P_0$  is constructed by concatenating the first stripes of all frames, or by stripes  $s_{0,0}, s_{1,0}, \dots, s_{6,0}$ . The second panorama  $P_1$  is consisted of the second stripes from all shots. To align the first and the second panoramas in the cross-panorama direction, the second panorama is rotationally shifted so that the stripe from frame  $F_{N-1}$  is at the head. In general, a smartly rebinned panorama  $P_i$  consists of the  $i$ th stripes of all frames cut and pasted sequentially, with the  $i$ th stripe of frame  $F_0$  at the  $i$ th slot

$$P_i = \{s_{(-i) \bmod N, i}, s_{(-i+1) \bmod N, i}, \dots, s_{(-i+N-1) \bmod N, i}\} \\ i = 0, 1, \dots, S-1.$$

An illustration of the resultant rebinned panorama is shown in Fig. 6. Some portions of the stripes in panorama  $P_4$  contain no data, as the corresponding image shot do not have a full fifth stripe. In this specific example, every stripe is at least partially filled. However, for an actual concentric mosaic scene with a larger variation of scene depth, several end stripes may be completely empty. The smartly rebinned panoramas are thus not of rectangular region of support. Special handling for coding those empty regions will be addressed in the next section. Note that in order to reverse such a rearrangement, only the horizontal displacement between each pair of neighboring shots needs to be recorded, adding minor overhead to the compressed bit stream.

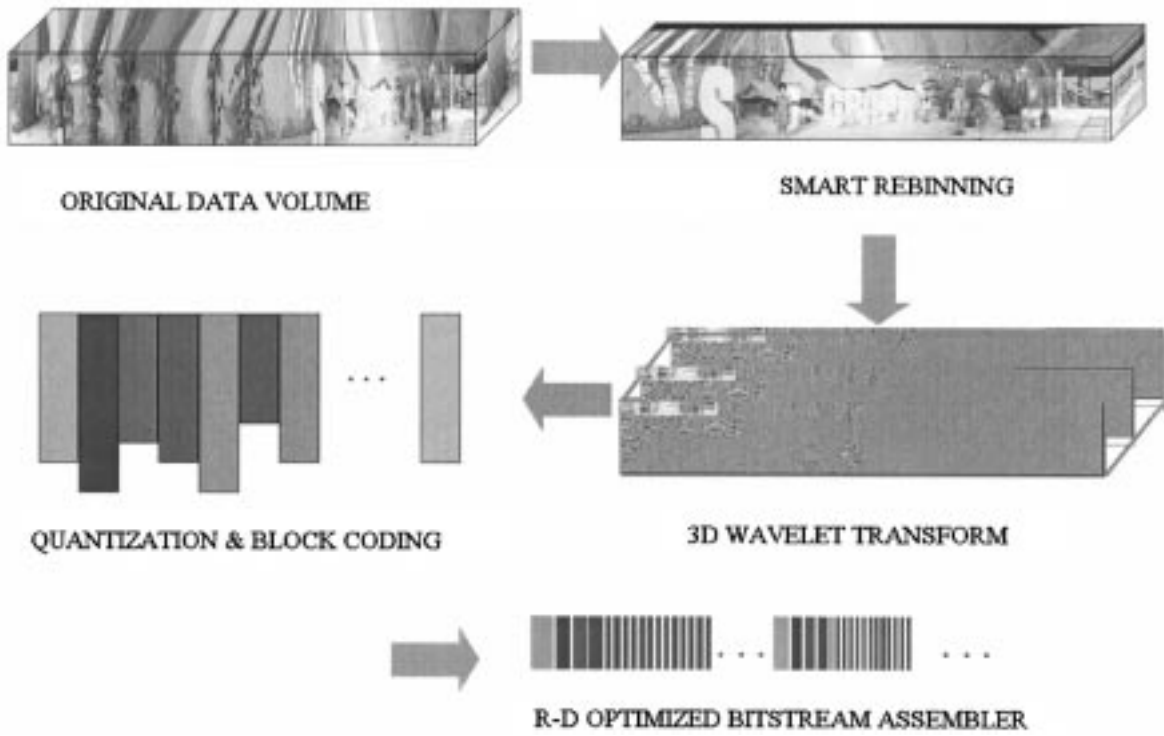


Fig. 4. Framework of smartly rebinned 3-D wavelet coder.




Fig. 5. Horizontal shot alignment of concentric mosaic image shots.

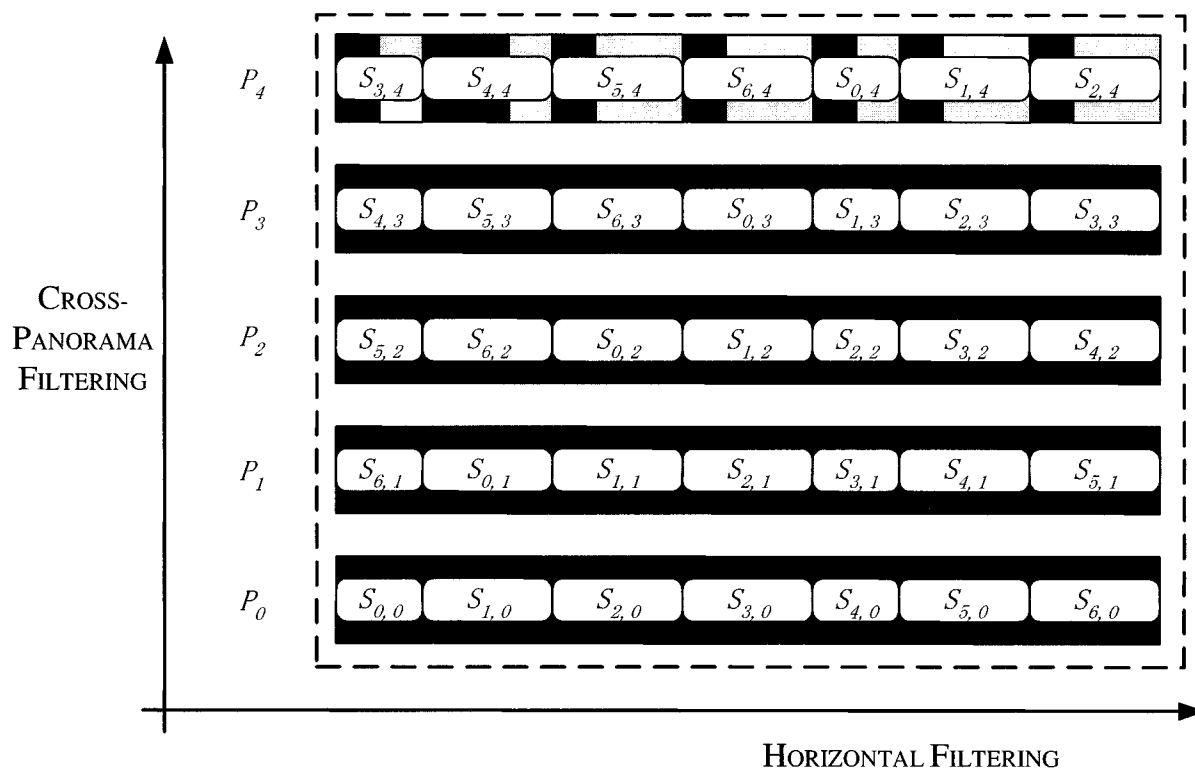


Fig. 6. Smartly rebinned data volume.

For the smartly rebinned panorama, the unfilled regions of the skewed data set are largely reduced, which makes the compression much more efficient and the implementation much more convenient. The 3-D wavelet filtering of the rebinned panorama is still very efficient. Filtering across the panorama is exactly equivalent to filtering across the image shots in the horizontal shot alignment approach shown in Fig. 5. However, filtering within an image shot is replaced by filtering within the rebinned panorama. The newly generated panorama  $P_i$  is highly correlated internally, because each stripe consists of successive slits in one original shot image, and two neighbor stripes are smoothly connected because they are from the matching stripes in neighboring concentric mosaic image shots. Consequently, horizontal filtering is still very efficient.

A degenerated approach of the smart rebinning is to restrict all horizontal translation vectors to be exactly the same

$$x_0 = x_1 = \dots = x_{S-1} = x.$$

We call this approach *simple rebinning*. All image shots now have the same number of stripes. If there are unfilled slits at the last stripe, we simply fill them by repeating the last slit. Rebinning the stripes into panoramas, a set of panoramas with a rectangular region of support is formed. The approach is similar to the formation of the concentric mosaic  $P_h = \{f(n, h, v) | n, v\}$  in [1]. The difference lies in that multiple slits are obtained from each shot to generate the rebinned panorama.

We show the volume of the original concentric mosaic in Fig. 7. The rebinned concentric mosaic forms a cube, with the front view showing a concentric mosaic panorama, the side view a camera shot, and the top view a cross section slice at a certain

height. We then show the smartly rebinned panorama volume in Fig. 8 as a comparison. The smartly rebinned panorama forms a volume of nonrectangular support, and the black region in Fig. 8 identifies the unsupported region. We note that the area with a smaller region of support is closer to the capturing camera, because it has a larger horizontal displacement vector, and thus contains a smaller number of stripes. In comparison with the concentric mosaic, the smartly rebinned panorama appears to be smoother and more natural looking, as it adjusts its sampling density according to the distance of the shot to the object, and maintains a relative uniform object size as seen by the camera. The smartly rebinned panoramas have strong correlation across the panoramas. A set of rebinned panoramas at the same horizontal location is extracted and shown in Fig. 9. We observe that most objects in the rebinned panoramas are well aligned. Only a few objects, such as the lightbulb at the upper-left corner and the balloons behind the girl, show differences due to the gradual parallax transition among the rebinned panoramas. Such a well-aligned data volume can be efficiently compressed by a 3-D wavelet transform.

In fact, the smartly rebinned panorama belongs to a general category of *multiperspective* panoramas that have recently become popular in the computer graphics realm, such as manifold mosaics [8], multiple-center-of-projection image [9], and circular projection [10]. Multiperspective panorama extends the conventional panorama by relaxing the requirement of having one common optical center and allows several camera viewpoints within a panorama. The idea of multiperspective panorama construction via cutting and pasting stripes was first introduced in [8]. It has also been extended to enable stereo viewing in [10], where the stripes taken from the left




Fig. 7. Volume of the concentric mosaic.



Fig. 8. Part of the volume of the rebinned multiperspective panorama set.



Fig. 9. Set of smartly rebinned panoramas at the same horizontal location (note the parallax shown by the lightbulb and the balloon behind the leg of the girl).

side of each image shot generate the right eye panorama and those from the right generate the left eye view. However, in contrast to the work of [8], [9], and [10], where only one or two panoramas are generated for their specific graphic application, we generate a whole set of rebinned panoramas to provide a dense representation of the environment, and to efficiently compress the concentric mosaic data set.

#### IV. THREE-DIMENSIONAL WAVELET CODING OF REBINNED PANORAMAS

We further encode the rebinned panoramas with a 3-D wavelet coder. Although other coders, such as the RBC in [14] can also be applied, 3-D wavelet coding is ideal because good alignment across multiple image shots can be more efficiently explored by the 3-D wavelet coder. For the simple rebinning, straightforward 3-D wavelet encoding may be adopted. We

use the 3-D wavelet codec with block arithmetic coding, as proposed in our previous paper [7]. The data volume of the concentric mosaic is decomposed by a multiresolution 3-D wavelet transform. The wavelet coefficients are then cut into fixed-size blocks, embedded encoded, and assembled with a rate-distortion optimization criterion. For details of the 3-D wavelet coding algorithm, we refer the reader to [7].

For the smartly rebinned panoramas, a 3-D wavelet coding algorithm that handles a data volume with an arbitrary region of support must be developed. Fortunately, there are wavelet algorithms designed to encode arbitrary shaped objects in the literature, most developed in the standardization process of MPEG-4 [11]. There are two categories of approaches:

- 1) padding the data to a rectangular volume and then using a rectangular codec;
- 2) using an arbitrary shape wavelet transform and coder.

We investigate both approaches.

A simple approach is to pad the arbitrary shaped region of support to the tightest rectangular volume containing it and apply the rectangular 3-D wavelet transform and coding algorithm to the padded data volume. In order to minimize the overhead incurred, padding should be designed such that the transform coefficients corresponding to the padded region are mostly zero or very small, because small coefficients consume fewer bits in the subsequent entropy coder. In this work, we extend the low-pass extrapolation (LPE) adopted in MPEG-4 for the padding. The unsupported regions are first filled with the average pixel value of the boundary of the supported/unsupported region, and then a low-pass filter is applied to the unsupported region several times. Since in the unsupported region, all pixel values are initialized with the same average value, the effect of the low-pass filter is primarily at the boundary, where a gradual transition is built up. After the wavelet transform, coefficients in the unsupported regions will be mostly zeros, except at the boundary. The padded data volume is then compressed with the 3-D wavelet codec described in [7]. Since the number of wavelet coefficients after padding is still more than the number of pixels in the supported region, the padding increases the coding rate, and therefore the compression performance is affected. The advantage is that the padding involves the least change in the 3-D wavelet codec, and is very easy to implement. Moreover, although the padding operation adds complexity in the encoder, it does not affect the decoder, which decodes the entire data volume and simply ignores the decoded pixels in the unsupported region.

Another feasible solution is to use an arbitrary shape wavelet transform [12] directly on the arbitrary shaped region of support. For wavelet transform in each direction, a set of straight lines parallel to the axis intersects the supported region and creates several segments. Each segment is then decomposed separately using a bi-orthogonal symmetric filter with symmetric boundary extension. We then store the coefficients in the wavelet domain, and record the region of support for the wavelet coefficients. The process can be recursively applied for multiresolution decomposition, and can transform the arbitrarily shaped concentric mosaic volume into an exact number of wavelet coefficients as that of the original data. For details of the scheme, we refer the reader to [12] and [13]. A block arithmetic coder with an arbitrary shaped region of support in the wavelet domain is then used to encode the transform coefficients. We call this codec the *3-D arbitrary shape wavelet codec*. It is observed that the arbitrary shape wavelet transform and coding is slightly superior in compression performance to padding the unsupported region. However, it slows down the decompression speed. It is also more complex to implement, as we need to add the support of the arbitrary shape region to both the transform and entropy coding modules.

The smartly rebinned and 3-D wavelet compressed concentric mosaic can be efficiently rendered as well. The rendering engine is very similar to the progressive inverse wavelet synthesis (PIWS) engine that we have proposed in [15]. Instead of decompressing the entire compressed concentric mosaic and then rendering it, the entire concentric mosaic can be kept in the compressed form, and only the part of the bit stream necessary to render the current view is decoded and rendered. This not only

reduces the memory requirement of the renderer, but also avoids a long decoding delay at the beginning. The working flow of the selective rendering engine can be shown in Fig. 10. According to the current viewing point and direction of the user, the rendering engine generates a set of slits that need to be accessed from the concentric mosaic data set. It then identifies the position of the accessed slits in the rebinned panorama set. Note that the position of a slit in the rebinned panorama is only related to the horizontal translation vectors between image shots, and can be easily calculated. After that, the PIWS engine is used to locate the positions of the slits in the wavelet domain. Only the wavelet coefficient blocks containing accessed slits are decoded from the compressed bit stream. With a mixed cache transition among the wavelet coefficients, the intermediate results, and the reconstructed pixels, PIWS ensures minimal computation is performed to recover the accessed slits.

Because smart rebinning can be considered as a preprocessing step of the 3-D wavelet coder, the only extra step in rendering the smartly rebinned concentric mosaic is to locate the slits in the rebinned panorama. The computational complexity of rendering the smartly rebinned concentric mosaic is thus similar to the rendering of 3-D wavelet compressed concentric mosaic. With the PIWS engine, a rendering rate of 12 frames/s is achievable, which is fast enough for the real-time rendering applications.

## V. EXPERIMENTAL RESULTS

The performance of the 3-D wavelet concentric mosaic compression with smart rebinning is demonstrated with extensive experimental results. The test scenes are *Lobby* and *Kids*. The scene *Lobby* has 1350 frames at resolution  $320 \times 240$ , and the total data amount is 297 MB. The scene *Kids* has 1462 frames at resolution  $352 \times 288$ , and the total data amount is 424 MB. The *Kids* scene contains more details, and is thus more difficult to compress than the *Lobby* scene. The scenes are first converted from RGB to YUV color-space with 4:2:0 subsampling, and then compressed by different coders. We compress the *Lobby* scene at ratio 200:1 (0.12 b/p, 1.48 MB) and 120:1 (0.2 b/p, 2.47 MB), and the *Kids* scene at 100:1 (0.24 b/p, 4.24 MB) and 60:1 (0.4 b/p, 7.07 MB). The peak signal-to-noise-ratio (PSNR) between the original and decompressed scene is used as the objective measure of the compression quality. We report the PSNRs of all three color components ( $Y$ ,  $U$ , and  $V$ ) in Table I; however, it is the PSNR result of the  $Y$  component that matters most. Therefore, we comment only on the  $Y$  component PSNR in the discussion.

We compare the proposed smartly rebinned 3-D wavelet coder with three benchmark algorithms. The first algorithm (A) compresses the entire concentric mosaic as a video sequence using a MPEG-2 video codec. The MPEG-2 software is downloaded from [www.mpeg.org](http://www.mpeg.org). In the MPEG-2 codec, the first frame is independently encoded as I-frame, and the rest of the frames are predictively encoded as P-frames. The second algorithm (B) is the direct 3-D wavelet codec reported in [7], where we convert the concentric mosaic image shots into mosaic panoramas, align the panoramas, and encode them with the 3-D wavelet and block arithmetic coding. The third



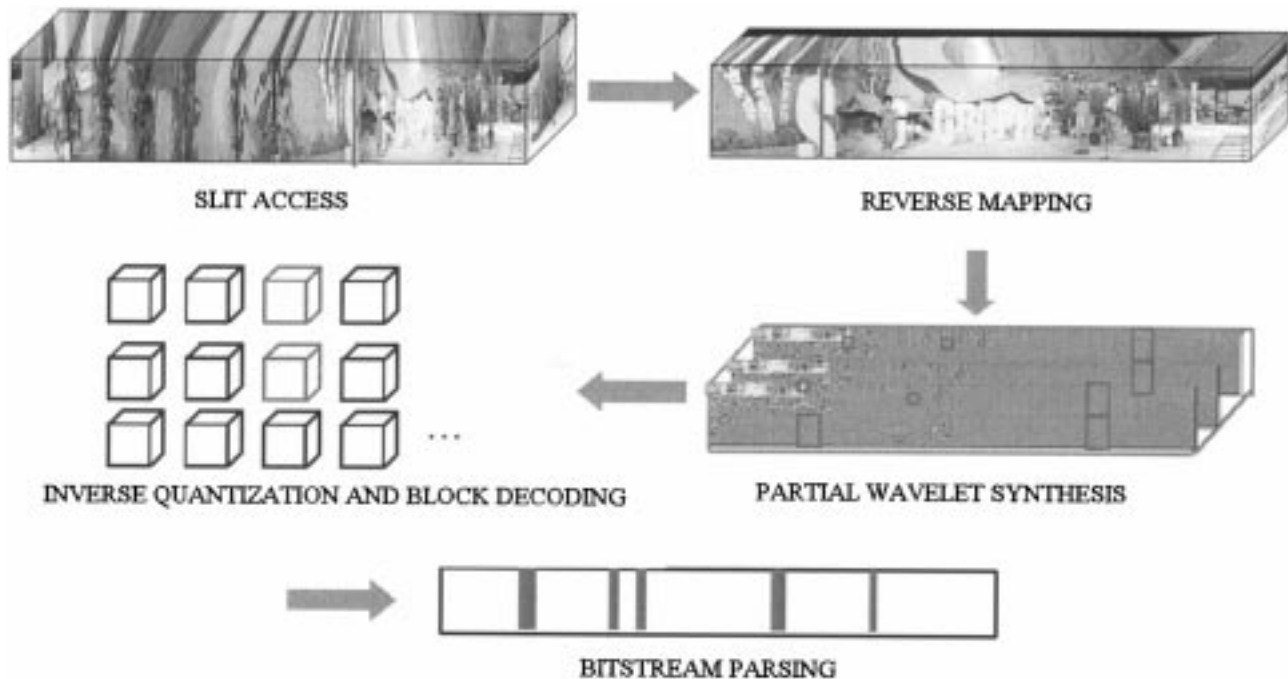


Fig. 10. Framework of the selective rendering engine.

TABLE I  
COMPRESSION RESULTS FOR THE CONCENTRIC MOSAIC SCENES

Test Dataset \ Algorithm		LOBBY	LOBBY	KIDS	KIDS
		(0.2 bpp)	(0.12 bpp)	(0.4 bpp)	(0.24 bpp)
A	<i>MPEG-2</i>	Y: 32.2	Y: 30.4	Y: 30.1	Y: 28.3
		U: 38.7	U: 37.4	U: 36.6	U: 34.8
		V: 38.1	V: 36.9	V: 36.7	V: 34.9
B	<i>3D Wavelet</i>	Y: 31.9	Y: 30.0	Y: 29.4	Y: 27.3
		U: 40.3	U: 39.3	U: 36.5	U: 34.9
		V: 39.9	V: 38.9	V: 37.2	V: 35.7
C	<i>RBC</i>	Y: 32.8	Y: 29.8	Y: 31.5	Y: 28.7
		U: 39.7	U: 38.4	U: 39.3	U: 37.3
		V: 40.5	V: 39.0	V: 38.9	V: 36.6
D	<i>Simple rebinning</i>	Y: 35.5	Y: 33.6	Y: 32.8	Y: 30.5
		U: 41.5	U: 40.7	U: 39.3	U: 37.7
		V: 40.9	V: 40.2	V: 40.1	V: 38.5
E	<i>Smart rebinning + padding</i>	Y: 36.0	Y: 34.0	Y: 33.4	Y: 31.1
		U: 41.6	U: 40.9	U: 39.9	U: 38.4
		V: 41.0	V: 40.2	V: 41.1	V: 39.6
F	<i>Smart rebinning + arbitrary shape wavelet codec</i>	Y: 36.3	Y: 34.3	Y: 33.8	Y: 31.3
		U: 43.9	U: 42.9	U: 41.1	U: 39.5
		V: 42.8	V: 42.0	V: 41.2	V: 39.6

benchmark algorithm (C) is the RBC reported in [14]. It is a prediction-based codec tuned for the compression of concentric mosaic. We observe that the direct 3-D wavelet coding of the concentric mosaic scene (Algorithm B) is not very efficient; it is 0.3 to 1.0 dB inferior to MPEG-2 with an average of 0.6 dB, and is inferior to the RBC codec with an average of 1.1 dB.

We test three different configurations of the 3-D wavelet codec with smart rebinning. In the first configuration (Algorithm D), we restrict the horizontal displacement vector between frames to be constant, i.e., the simple rebinning is used. The actual displacement vector is 2 and 3 pixels for the *Lobby* and *Kids* scenes, respectively. The resultant rebinned

concentric mosaic forms a rectangular panorama volume and is compressed by exactly the same 3-D wavelet and block arithmetic coder as Algorithm B. It is observed that by simply rebinning multiple slits into the panorama, a large compression gain can be achieved. In fact, compared with the direct 3-D wavelet codec, the PSNR improves between 3.2 and 3.6 dB, with an average gain of 3.5 dB. The 3-D wavelet coder with simple rebinning outperforms the MPEG-2 concentric mosaic codec by 2.9 dB, and outperforms the RBC codec by 2.4 dB.

We then apply the full-fledged smart rebinning algorithm. The horizontal displacement vectors are calculated by matching neighborhood concentric mosaic image shots. In this paper, this

is done by first calculating the average intensity of each vertical line, and then correlating the vertical average intensity of the two shots

$$x_n = \arg \max_d \frac{\sum_h f_{av}(n, h) f_{av}(n+1, h-d)}{\sqrt{\sum_h f_{av}^2(n, h)} \sqrt{\sum_h f_{av}^2(n+1, h-d)}}$$

where

$$f_{av}(n, h) = \frac{1}{V} \sum_v f(n, h, v).$$

The vectors are then stored in the compressed bit stream. After the rebinning operation, the bounding volume for the rebinned panoramas is  $2832 \times 162 \times 240$  for the *Lobby* scene and  $5390 \times 149 \times 288$  for the *Kids* scene. In the *Lobby* scene, object is of relatively constant depth to the camera, and the unsupported regions occupy only 6% of the bounding volume. However, in the *Kids* scene which has a larger depth variation, 36% of the bounding volume is unsupported. The smartly rebinned panoramas are compressed through two approaches. In the first approach, we compress the rebinned panoramas through padding the data volume and applying a rectangular 3-D wavelet codec used in the Algorithms B and D (denoted as Algorithm E). Alternatively, we use an arbitrary shape wavelet transform and coefficient coding algorithm developed in [13] (denoted as Algorithm F). According to the results shown for Algorithm F, the smart rebinning further improves the compression performance over simple rebinning by 0.7 to 1.0 dB, with an average gain of 0.8 dB. The average gain of the arbitrary shape wavelet codec (F) over the padding approach (E) is 0.3 dB. Note that the system of Algorithm E is very close in complexity to that of the simple rebinning (Algorithm D), because both systems use the rebinning, rectangular 3-D wavelet transform, and block arithmetic coding. The only difference is that the Algorithm D rebins a fixed number of slits into the panorama, while Algorithm E rebins a variable number of slits into the panorama, which is then padded before coding. In terms of PSNR performance, Algorithm E outperforms Algorithm D by 0.5 dB on average. Therefore, general smart rebinning with calculated horizontal translation vectors does have an advantage over simple rebinning, where a fixed translation vector is used for all image shots.

Overall, smart rebinning with arbitrary shape wavelet transform and coding is the best performer of the proposed approaches. It outperforms the MPEG-2 concentric mosaic codec by an average of 3.7 dB, outperforms the direct 3-D wavelet video encoder by 4.3 dB, and outperforms the RBC by 3.2 dB. The PSNR of the smart rebinning compressed *Lobby* at 0.12 b/p is even superior to prior concentric mosaic coders operated at 0.2 b/p. More specifically, it is 2.1 dB superior to the MPEG-2, 2.4 dB superior to the direct 3-D wavelet, and 1.5 dB superior to the RBC compressed scene at 0.2 b/p. Since the PSNR of the *Lobby* scene compressed at 0.2 b/p is on average 2.1 dB higher than the PSNR of the same scene compressed at 0.12 b/p, the smart rebinning almost quadruples the compression ratio for the *Lobby* scene over prior approaches. In the same way, we observe that the smart rebinning nearly doubles the compression ratio for the *Kids* scene over prior approaches. Considering the huge amount of

data of the concentric mosaic, and the relatively large bit stream even after a high ratio compression has been applied (1.48–7.07 MB), smart rebinning is a very effective tool to greatly reduce the amount of data of the concentric mosaic.

## VI. CONCLUSION AND EXTENSION

A technology termed *smart rebinning* is proposed in this paper to improve the 3-D wavelet compression of the concentric mosaic. Through cutting and pasting stripes into a set of multiperspective panoramas, smart rebinning greatly improves the performance of cross-shot filtering, and thus improves the transform and coding efficiency of the 3-D wavelet coder. The region of support after smart rebinning may cease to be rectangular, and a padding scheme and an arbitrary shape wavelet coding scheme have been used to encode the resultant data volume of smart rebinning. With the arbitrary shape wavelet codec, smart rebinning outperforms MPEG-2 by 3.7 dB, outperforms a direct 3-D wavelet coder by 4.3 dB, and outperforms the RBC by 3.2 dB on the tested concentric mosaic scenes. It nearly quadruples the compression ratio for the *Lobby* scene, and doubles the compression ratio for the *Kids* scene.

It will be interesting to extend the smart rebinning technology to general 3-D wavelet coding of video sequence, and see if similar performance gain can be achieved. However, such extension is far from straightforward. The motion is more complicated in a typical video sequence. Moreover, video needs to be decoded frame-by-frame, which does not match well for a data rearrangement scheme such as rebinning. We are investigating along this direction.

## ACKNOWLEDGMENT

The authors would like to thank H. Shum, H. Sun, and M. Wu for the raw concentric mosaic data and the concentric mosaic render; and J. Xu, S. Li, and Y. Zhang for providing the arbitrary shape wavelet transform and coder in [13].

## REFERENCES

- [1] H.-Y. Shum and L.-W. He, "Rendering with concentric mosaic," in *Computer Graphics Proc., Ann. Conf. Series (SIGGRAPH)*, Los Angeles, CA, Aug. 1999, pp. 299–306.
- [2] D. Taubman and A. Zakhor, "Multirate 3-D subband coding of video," *IEEE Trans. Image Processing*, vol. 3, pp. 572–588, Sept. 1994.
- [3] A. Wang, Z. Xiong, P. A. Chou, and S. Mehrotra, "3-D wavelet coding of video with global motion compensation," presented at the DCC, Snowbird, UT, Mar. 1999.
- [4] J. R. Ohm, "Three-dimensional subband coding with motion compensation," *IEEE Trans. Image Processing*, vol. 3, pp. 572–588, Sept. 1994.
- [5] J. Y. Tham, S. Ranganath, and A. A. Kassim, "Highly scalable wavelet-based video codec for very low bit-rate environment," *IEEE J. Select. Areas Commun.*, vol. 16, Jan. 1998.
- [6] D. Taubman and A. Zakhor, "A common framework for rate and distortion-based scaling of highly scalable compressed video," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, pp. 329–354, Aug. 1996.
- [7] L. Luo, Y. Wu, J. Li, and Y.-Q. Zhang, "Compression of concentric mosaic scenery with alignment and 3-D wavelet transform," *Proc. SPIE, Image Video Commun. Process.*, vol. 3974-10, Jan. 2000.
- [8] S. Peleg and J. Herman, "Panoramic mosaics by manifold projection," in *IEEE Conf. Computer Vision and Pattern Recognition*, San Juan, PR, June 1997, pp. 338–343.
- [9] P. Rademacher and G. Bishop, "Multiple-center-of-projection images," in *Computer Graphics Proc., Ann. Conf. Series (SIGGRAPH)*, Orlando, FL, July 1998, pp. 199–206.

- [10] S. Peleg and M. Ben-Ezra, "Stereo panorama with a single camera," in *IEEE Conf. Computer Vision and Pattern Recognition*, Fort Collins, CO, June 1999, pp. 395–401.
- [11] *MPEG-4 Video Verification Model 14.2.*, Dec. 1999.
- [12] J. Li and S. Lei, "Arbitrary shape wavelet transform with phase alignment," presented at the Int. Conf. Image Processing, Chicago, IL, Oct. 1998.
- [13] J. Xu, S. Li, and Y. Zhang, "Three-dimensional shape-adaptive discrete wavelet transforms for efficient object-based video coding," presented at the IEEE/SPIE VCIP, Perth, Australia, June 2000.
- [14] C. Zhang and J. Li, "Compression and rendering of concentric mosaic scenery with reference block codec (RBC)," in IEEE/SPIE VCIP, Perth, Australia, June 2000. Accepted.
- [15] Y. Wu, L. Luo, J. Li, and Y.-Q. Zhang, "Rendering of 3-D wavelet compressed concentric mosaic scenery with progressive inverse wavelet synthesis (PIWS)," presented at the IEEE/SPIE VCIP, Perth, Australia, June 2000.
- [16] M. Magnor, P. Eisert, and B. Girod, "Model-aided coding of multi-viewpoint image data," presented at the ICIP, Vancouver, BC, Canada, Sept. 2000.
- [17] M. Magnor and B. Girod, "Model-based coding of multi-viewpoint imagery," presented at the IEEE/SPIE VCIP, Perth, Australia, June 2000.



**Yunnan Wu** received the B.S. degree in computer science from the University of Science and Technology of China, Hefei, in 2000. He is currently pursuing the Ph.D. degree in the Department of Electrical Engineering, Princeton University, Princeton, NJ.

His research interests are in communications, signal processing, and multimedia compression.

Mr. Wu received the 2000 SPIE and IS&T Visual Communication and Image Processing Best Student Paper Award.



**Cha Zhang** received the B.S. and M.S. degrees from the Department of Electronic Engineering, Tsinghua University, Beijing, China, in 1998 and 2000, respectively. He is currently pursuing the Ph.D. degree in the Department of Electrical and Computer Engineering, Carnegie-Mellon University, Pittsburgh, PA.

His research interests include information retrieval, pattern recognition, machine learning, image-based rendering, image-video compression and streaming, and signal processing.



**Jin Li** (S'94–A'95–M'96–SM'99) received the M.S. and Ph.D. degrees in electrical engineering from Tsinghua University, Beijing, China, in 1991 and 1994, respectively.

He was a Research Associate at the University of Southern California, Los Angeles, from 1994 to 1996, and a Member of the technical staff at the Sharp Laboratories of America, Camas, WA, from 1996 to 1999. In 1999, he joined Microsoft Research, and was a Researcher/Project Leader at Microsoft Research China, Beijing, from 1999 to 2000, then moved to Microsoft Research Redmond, Redmond, WA, in 2001.

Since 2000, he has also served as an Adjunct Professor in the Department of Electrical Engineering, Tsinghua University. He is an active contributor to the ISO JPEG 2000/MPEG4 Project, and has published over 50 technical papers in various journals and conferences relating to multimedia compression and communication. He is an Area Editor for the *Journal of Visual Communication and Image Representation*.

Dr. Li was the recipient of the 1994 Ph.D. Thesis Award from Tsinghua University and the 1998 Young Investigator Award from SPIE *Visual Communication and Image Processing*.