Anantrasirichai, N., Canagarajah, CN., Redmill, DW., & Akbari, A. S. (2007). Volumetric reconstruction with compressed data. In *International Conference on Distributed Smart Cameras, 2007. ICDSC '07* (pp. 328 - 335). Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/ICDSC.2007.4357541>

Peer reviewed version

Link to published version (if available):
[10.1109/ICDSC.2007.4357541](https://doi.org/10.1109/ICDSC.2007.4357541)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

VOLUMETRIC RECONSTRUCTION WITH COMPRESSED DATA

N. Anantrasirichai, C. Nishan Canagarajah, David W. Redmill and Akbar Sheikh Akbari

Department of Electrical & Electronic Engineering,
University of Bristol, Bristol, BS8 1UB, UK

ABSTRACT

Volumetric reconstruction is one of the 3D processing technologies for multi-view systems. The existing methods of volume reconstruction have been proposed by exploiting the original views with many constraints. These algorithms may not be suitable for the distributed camera system where multiple views are transmitted via lossy networks before processing at the computing centre. Therefore, this paper proposes a novel volumetric reconstruction from the compressed multi-view images. The algorithm starts with the depth registration, and then the initial volume is refined. Finally the colour selection scheme provides the realistic colour to the volume thereby achieving subjective quality of rendering views as shown in the subjective results. The performance of the proposed scheme reconstructing with the compressed views may be inferior than reconstructing with the original views up to 3 dB, but it can be superior over 3 dB at low bitrate.

Index Terms— volumetric reconstruction, multiview

1. INTRODUCTION

Increasing demands for realism or more natural representations of a scene has opened a wide variety of multi-view imaging researches and applications. In the beginning, stereo vision was introduced to provide three-dimensional (3D) perception. Recently, realistic vision has been further developed to provide free-viewpoint applications. It is normally reconstructed from multiple images which can be achieved with distributed cameras. All of these 3D visual developments have driven various applications, such as virtual view synthesis, teleconference, surveillance, telemedicine, remote education, gaming, free-viewpoint television (FVT) and 3D television (3DTV).

Evidently, huge amount of data is generated in a multi-view system. It requires an efficient process of analysis and interpretation, as well as the efficient management of storage and transmission. Computer vision and computer graphics greatly aid the 3D analysis and interpretation, whilst compression is a key solution to reduce the size of data. Volumetric representation can achieve both requirements. As the sequences are captured from the same scene, the

volumetric representation removes redundant data among views. For view synthesis, the volumetric representation produces consistent 3D information thereby improving the accuracy of depths and view synthesis performance, whilst synthesising directly from depth maps estimated individually may cause errors due to their inconsistency.

In conventional image and video coding, system complexity in transmitters is much higher than in receivers. However, many applications, e.g. sensor networks and multi-camera systems, require the complex process to be done in the receivers because of the limited capacities of the transmitter [1]. The 3D processing is consequently forced to operate in the receiver in which some details of the reference views are lost and may not be sufficient to extract the correct 3D information.

In this paper, we present a novel volumetric reconstruction which is applicable for the distributed camera system. We propose the reconstruction of a 3D object from a sparse set of 2D compressed images which is originally taken from several angles with wide-baseline geometry. The sparse multi-views are generally preferred in the distributed cameras in order to prevent the difficulty of synchronisation and data communication. At the receiver, the intensity balance compensation is applied to the compressed views to adjust the details that may be lost by using the same idea of stereoscopic colour balancing in [2]. Then, the depth of each view is estimated. These depth maps are subsequently utilised to initialise 3D volume. The depth registration method is more suitable to reconstruct the volume compared to the direct voxel searching in the lossy scenario.

We also propose a novel algorithm to refine the shape of the volume. The constructed volume should accurately represent the real scene or object. Its incorrect shape could adversely affect the virtual view synthesis. We therefore adjust the surface by considering the neighbouring voxels. Finally, colours are defined for each voxel. An important aspect of this work is that it should produce a smooth and realistic image. Therefore, the colour defined under the visibility constraints should satisfy the smoothness constraint. That means the artificial colours sometimes give better subjective quality and better viewing experience than the true colours.

The rest of paper is organized as follows: Section 2 reviews the existing volumetric reconstruction techniques.

Section 3 explains the 3D processing appearing in multi-view communications. Then, the proposed volumetric reconstruction is described in Section 4. The experimental results and discussions are presented in Section 5. Finally, the conclusions and future work are presented in Section 6.

2. EXISTING VOLUMETRIC RECONSTRUCTION SCHEMES

The volumetric reconstruction uses a voxel-based representation of which each element, a *voxel*, represents a value on a regular grid in the 3D space. Techniques for the scene reconstruction from the multi-view sequences can be classified into two categories. The first is direct voxel searching of which the occupancy of the voxels is designed by the silhouette and/or the corresponding pixels on reference views. The *shape-from-silhouette* algorithm is based on volume intersection [3]. The binary data taken from an object's boundary are registered into a volume. Subsequently, the algorithm has been developed by exploiting grey scale, *shape from photo-consistency*. With each voxel it must be determined whether it is transparent or whether it is one of the volume by following photometric constraints [4]. Some algorithms check the voxel visibility by using *plane-sweep*, in which the distance between a plane and cameras is increased, so the existing voxels in the front plane may occlude some voxels in the back plane when projecting to a particular view [5]. A multi-plane-sweep method has been developed by iteratively removing the border voxels which do not satisfy the photo consistency, so-called *Space Carving* [6]. The algorithms in the first category work well with a very narrow-baseline geometry and therefore requires a high number of views.

The second is using previously estimated depth maps to construct a volume. Each depth map is warped to the

geometry of the volume and the opaque or the transparent voxels are easily marked [7]. Several algorithms have been proposed to manage the error and noise in the depth maps, such as using a probability distribution of input data [8], and using a Bayesian method [9]. All of these enhanced methods have been proposed under the assumption of Lambertian reflection and required a great number of reference views.

In the lossy scenario, the depth registration method is more appropriate to reconstruct the volume compared to the direct voxel searching. The results of depth estimation always show one value of depths belonging to its optical ray. Although it may be incorrect, it can be a guide which will be used for adjusting the voxels to the more proper positions in the refinement process.

3. 3D PROCESSING IN MULTI-VIEW COMMUNICATIONS

In the last decade, multi-view communication has been developed under a framework that the 3D processing has been done before the compression process. The 3D information, i.e. disparity/depth and model of the volume, is extracted from the original multi-views. We will call this system as a pre-processing geometric reconstruction system. As illustrated in Fig. 1, the multi-views are applied to the 3D processing for estimating the disparity/depth. Then, the disparity/depth and the original multi-views are employed together so as to reconstruct the 3D volumes if the volume-based approach is selected, whilst the volumetric representation can be skipped if the image-based approach is operating. Note that the disparity/depth estimation can also be omitted if the direct searching voxel algorithm is applied for volumetric representation. Finally, the 3D information and textures, including the residuals from the compensation

Fig. 1. Pre-processing geometry reconstruction. (a) Transmitter. (b) Receiver.

Fig. 2. Post-processing geometry reconstruction. (a) Transmitter. (b) Receiver.

process, are compressed. At the receiver, the 3D information and textures are first decompressed. The reconstructed views are then generated by reverse compensation process. Subsequently, these data will be used for image rendering and view synthesis depending on applications. Obviously, the complexity of this system issues mostly at the transmitter. Besides, the bitstreams comprising 3D information and textures possibly cause the inefficient compression, since the removal of redundancies among all views cannot compensate the great number of bits required for coding 3D information.

Distributed video coding (DVC), in contrast, offers an alternative solution to the complexity balance issued between the transmitter and the receiver [10]. Basically, this system can be achieved based on Slepian-Wolf [11] and Wyner-Ziv [12] coding techniques. With this shared-complexity idea, the multi-view coding system may consequently be changed to the system as shown in Fig. 2. The transmitter operates only for data compression which offers two options; one encodes each view independently leading to the very low complex transmitter, while another includes the prediction process which is more appropriate for the system with the great number of views. The receiver of this system will decompress the multi-views, and then the reconstructed views are used for extracting the 3D information. As a result, the image rendering and view synthesis can operate with the similar approach of the pre-processing geometric reconstruction system. Generally, the volume reconstruction process is used for 3D processing, since it achieves a good image rendering and view synthesis. Some constraints, e.g. intensity constraint, are required to extract and reconstruct the 3D modelling. Obviously, the defect of the compressed multi-view images can inevitably cause the difficulty in 3D processing.

The next section presents the proposed volumetric reconstruction. Any depth estimation technique can be used and may generate the depth maps with any uncertainty and noise because of the lossy reference views.

4. PROPOSED VOLUME RECONSTRUCTION

In this section, the reconstruction of a 3D object from a set of 2D images taken from several angles with wide-baseline geometry is proposed. It is assumed that the camera

parameters are known. The association between a 3D volume and a set of 2D-image geometries can be stated as follows: A pixel $\mathbf{x}_m = [i_m, j_m]^T$ in image view m , which has the depth w_{ij}^m and the geometry parameters $\mathbf{P}_m$ (consisting of the intrinsic and extrinsic camera matrix), is mapped to a voxel position $\tilde{\mathbf{x}} = [\tilde{i}, \tilde{j}, \tilde{k}]^T$ with the invertible relation:

$$[\tilde{i}, \tilde{j}, \tilde{k}]^T = \tilde{\mathbf{P}} \mathbf{P}_m^{-1} [i_m w_{ij}^m, j_m w_{ij}^m, w_{ij}^m]^T \quad (1)$$

where $\tilde{\mathbf{P}}$ is the geometry parameters of the volume.

At the receiver of the lossy network, the details lost in each view are inconsistent, so the searching of their corresponding is hard to achieve. This problem happens similar to the problems from non-Lambertian surfaces, as well as luminance and chrominance imbalances. That is, the intensity of a point in 3D space may not be the same as that of its corresponding pixel in 2D image. A voxel that does not satisfy the photometric constraints is transparent or is removed from the volume. Hence, the reconstructed volume may contain holes or too many transparent voxels. Our algorithm therefore includes colour constraints and a tolerance for corresponding matching.

There are three steps to construct a volume which are Volume initialization, Volume refinement and Colour selection:

4.1. Volume Initialization

Firstly, depth maps are registered into a volume. Each registered voxel has a value equal to the number of views that are mapped to this voxel; otherwise, if unmapped, the voxels have a value of 0. The voxel value δ_{ijk} can be expressed as follows:

$$\delta_{ijk} = \begin{cases} n, & \text{if registering depth,} \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where the number of views that are mapped to a voxel $\tilde{\mathbf{x}}$ is n for $n \in \{1, 2, \dots, N\}$, with total N reference views. Then some isolated voxels which do not have any neighbours are detected. These have a high possibility of being errors; therefore, they should be removed from the volume. Next, the volume is projected back to the reference view geometry to build up new depth maps that now include the depth information from other views of which some errors

were removed in the previous step. If there is more than one depth value for a pixel, the front cluster of depths is considered. Spurious noise causes dispersion; hence, the proper depth value $\tilde{d}$ of a particular pixel is

$$\tilde{d} = \begin{cases} d_n, & \text{if only one depth has } \delta_n = \max(\delta), \\ \left(\sum_{k=1}^K \delta_k d_k \right) / \left(\sum_{k=1}^K \delta_k \right), & \text{otherwise} \end{cases} \quad (3)$$

where K is the number of depth values in the front cluster. The depth value projected from the voxel that contains the maximum δ is selected and other depths are presumed to be a dispersion noise which is eliminated. If the maximum δ appears in two or more depth values, the fit depth value is calculated from the average of all depth values of the front cluster with weighted coefficient δ . In this case, no peak corresponding to the right depth is clearly present, so the centre value is selected. Subsequently, the volume is reconstructed again. In this step, some voxels are interpolated in order to connect the surface of the volume among the near voxels. This is necessary because the inconsistent intensity among the corresponding pixels of the compressed views is present.

4.2. Volume Refinement

Volume refinement is to refine the volume by iterating several times. It checks from the outer to the inner surfaces. The iteration is needed because once adjusting colour or position of a voxel possibly changes the results of the projection to 2D images in which the invisible pixel may become the visible pixel. Consequently, rechecking the quality of the reconstructed views improves the accuracy of the volume. In the recursive manner, the reconstructed views are rendered from the volume with colours and compare it with the original reference images. The colours are temporarily defined by taking the median of colour values from all the visible views. Note that the median is less sensitive to extreme values than the mean and this makes it a better measure for highly skewed distributions.

The iterative process runs from the first reference view, probably the furthest left view, to the last reference view, probably the furthest right view. For each view, the MSE between the reconstructed view and the original view are calculated. If the error is more than a threshold, the old depth is removed and it also causes the associated voxel to be removed. Then the new depth, which is derived from the neighbouring pixels (a smooth surface assumption) is investigated. If this new depth value meets colour constraints for all visible views, the new associated voxel is generated. It is noticeable that the imperfect camera calibration and unequal resolutions of each of the reference views possibly make a mapped pixel to be displaced from its exact position. A window-based computation, e.g. 3x3 pixels, therefore gives a better result for error calculation than exploiting only one pixel.

If some voxels are removed or newly generated, the images are iteratively reconstructed. This is because the change of the position of a voxel produces a new depth, which may change an invisible point to a visible point and vice versa. The refinement process develops the better disparity/depth maps of every view. The volume constructed by exploiting a single view data is extremely sensitive to the viewpoint and the additional data from other views can significantly improve this volume. That means the depth map of arbitrary viewpoint is also improved with global constraint.

4.3. Colour Selection

From the previous step, a fine volume is constructed. This section illustrates the post-processing colour selection for each voxel. Due to the loss of details, the temporary colours defined in the iterative volume refinement are probably not suitable. Smooth colours and brightness are needed, hence the weighted average of the colours from one or two of the nearest visible views is utilised to define voxel colour $\tilde{C}$. There are two camera configurations to be considered, namely linear camera configuration and planar camera configuration. In the first case, the cameras are located in a single direction, e.g. horizontal or vertical direction, whilst the cameras are located in both directions in the latter case.

4.3.1. Linear camera configuration

First of all, the line L_c of the volume's centre is approximately marked. Then the line L_v along the volume surface that is the intersection of the surface and the line created from L_c to the camera centre c_v of view v are drawn as illustrated in Fig. 3 (a). By scanning along L_c , the voxels of each layer obtain the colours as Equation 6.3 with a condition that a voxel $\tilde{x}$ is visible in at least two views and it is located between L_1 and L_2 . The first two nearest visible views are view1 and view2 respectively.

$$\tilde{C} = \begin{cases} (\alpha_1 C_2 + \alpha_2 C_1) / (\alpha_1 + \alpha_2) & \text{if the condition is true,} \\ C_1 & \text{otherwise} \end{cases} \quad (4)$$

where C_1 and C_2 represent the intensity component (Y) or the colour components (U,V) of view 1 and view 2. α_1 and α_2 are the angles between the line from the voxel $\tilde{x}$ to the nearest point on L_c and the point in the same layer as such voxel on L_1 and L_2 respectively.

4.3.2. Planar camera configuration

In this case, the upper and/or lower cameras (the right and/or left cameras) are also available. Therefore, instead of using the marked centre line L_c , the marked centre point P_c of the object is indicated. Consequently, the α_1 and α_2 are the angles between the line from P_c to the camera centre c_v and the line from P_c to $\tilde{x}$ as illustrated in Fig. 3 (b).

Fig. 3. Colour selection process for (a) linear camera configuration and (b) planar camera configuration.

Fig. 4. The magnified texture ((a) of a doll, (b) of a video tape) of synthesised view by applying different colour definition (1) median, (2) proposed scheme.

The average method may cause object to blur, if C_1 and C_2 are too different. Therefore, the defined colours are set into a *blurred* group, if the difference of C_1 and C_2 is more than a particular threshold, or a *clear* group, if the difference of C_1 and C_2 is less than or equal the threshold. To preserve sharpness, colour $\tilde{C}_{old}$ of the voxel in the blur group is replaced by the colour of neighbouring voxel that has the closest value to $\tilde{C}_{old}$ and has to be in the clear group.

The proposed colour selection noticeably outperforms the traditional median method as shown in the magnified version in Fig. 4. Some noisy colours shown in Fig. 4 (a-1) and (b-1) result from the loss of some details of the compressed reference views. The different colours from the same or different reference views can be mapped to a single voxel. After applying colour selection, the colours of the reconstructed views are smoother as shown in figure (a-2) and (b-2).

5. EXPERIMENTAL RESULTS

The performances of the proposed volumetric reconstruction were tested with the compressed multi-view images (post-processing) and were compared to the pre-processing geometric reconstruction system. In experiments, the dynamic programming proposed in [13] was utilised to estimate depth maps. The volumetric compression proposed in [14] was employed in the pre-processing system, whilst in-band disparity compensated view filtering scheme proposed in [15] was employed for data compression in the post-processing. Three multi-view test sequences: *Head*¹, *Santa*¹ and *Leo*², were used for testing the proposed scheme. Note that the object of the Leo sequence was processed separately from its background. The results are discussed as follows:

For the Head sequence, the results in Fig. 8 were synthesised from three reference views at bitrate of 35 kbits/view. The average PSNR values of the compressed reference views are 30.15 dB and 36.04 dB for the pre-processing scheme (using the volume-based coding) and the post-processing scheme (using the I-DCVF coding) respectively. The results of the post-processing scheme are significantly better than those of the pre-processing scheme for the reason that the reconstructed views of the I-DCVF achieve the higher quality than those of the volumetric coding at this bitrate thereby attaining the exceptional result of the volume representation at the receiver. It can be seen that, for example, the poster with white background located at the right of the image shows clearer texts in Fig. 8 (b) than in Fig. 8 (a). Fig. 5 shows the view synthesis performance of the post-processing and pre-processing methods. The quality of the synthesised images at low bitrates is improved by using the post-processing because the limitation from 3D information of the pre-processing is not encountered.

For the Santa sequence, at the bitrate of 20 kbits/view, the qualities of the synthesised views of the I-DCVF and the volumetric coding are slightly different as shown in Fig. 6. However, the results of the I-DCVF coding contain unclear details, e.g. the texture of the background near the Santa toy. These blurred compressed images produce the noticeable errors in the estimated depth maps, especially at the background because it is difficult to find the correct matching of correspondences. However, the synthesised views show acceptable perceptions, since the background of this sequence is not complicated. The colour selection method proposed in section 4.3 can put the appropriate colour to each voxel; therefore, the users can perceive the smooth, realistic and natural images as shown in Fig. 9 (b).

¹ Head and Santa sequences are from University of Tsukuba.

² Leo sequence are from University of Bristol.

Fig. 5. Rate-distortion performance of the *Head* sequence with 4 views.

Fig. 6. Rate-distortion performance of the *Santa* sequence with 15 views.

Fig. 7. Rate-distortion performance of the *Leo* sequence with 13 views.

For the *Leo* sequence, Fig. 7 and Fig. 10 show the interesting results at low bitrate. Although, the reconstructed reference views of the I-DCVF gain better quality than those of the volumetric coding at the same bitrate, the 3D processing at the receiver generates the noticeable loss of

some details at the synthesised views, especially at the regions where the colours of the object and background are totally different such as the area of the video tape. Moreover, the object and the background of this sequence are individually operated, so the incorrect volumetric representation is more observable when compared to the *Santa* sequence. However, the reconstructed object of the pre-processing scheme contains noisy texture, e.g. the side of the toy, because more high frequency loss comparing to the post-processing result.

From above experimental results, some observations can be noted. The 3D processing at the receiver can notably improve the quality of the view synthesis as the performance of the post-processing scheme is better than using the information sent with the coded texture in the pre-processing scheme, particularly at very low bitrates. When the bitrate increases, it shows the inferior performance owing to the lack of the true 3D information. However, its performance increases when the bitrate increases until the quality of the compressed views is close to the original views. The quality of the synthesised views is consequently similar to that of the pre-processing scheme. Although, the true 3D information cannot be extracted exactly from the low-quality reference views, the viewers may feel realistic with the synthesised results.

Fig. 11 illustrates the synthesised view of the proposed volumetric reconstruction against that of the plane-sweep method which were done at the receiver side. The result of the plane-sweep method clearly shows the loss of background details. Moreover, it generates noise around the edges of the object. These problems occur due to the lack of intensity consistency among the compressed reference views.

6. CONCLUSIONS AND FUTURE WORK

A novel volumetric reconstruction for compressed multi-view images is proposed. The volume is initially reconstructed by exploiting depth registration. Then this volume is refined by removing outlying voxels. By iteration, the voxels in the volume are further removed, if they produce high error compared to the reference images. Finally particular colours are defined for each voxel to make the viewer perceive a smooth and realistic image. The experimental results show that the proposed volumetric reconstruction works well in the lossy scenarios. The view synthesis performance is insignificantly inferior than using the original reference views. Moreover, it achieves the 3D processing at low bitrates. In the future, this proposed scheme will be combined with the distributed video coding techniques.

(a)

(b)

Fig. 8. The disparity maps and the synthesised view of the Head sequence (3 views) reconstructed at 35 kbits/view (a) pre-processing (PSNR=28.45 dB) and (b) post-processing (PSNR=32.86 dB).

(a)

(b)

Fig. 9. The disparity maps and the synthesised view of the Santa sequence (15 views) reconstructed at 20 kbits/view (a) pre-processing (PSNR=28.85 dB) and (b) post-processing (PSNR=29.26 dB).

(a)

(b)

Fig. 10. The disparity maps and the synthesised view of the Leo sequence (13 views) reconstructed at 75 kbits/view (a) pre-processing (30.73 dB), (b) post-processing (28.95 dB).

(a)

(b)

Fig. 11. Synthesised views of Santa: (a) proposed scheme; (b) plane sweep method.

7. REFERENCES

- [1] Z.Q. Luo, M. Gastpar, J. Liu, A. Swami, "Distributed Signal Processing in Sensor Networks," *IEEE Sig. Proc. Mag.*, pp. 14-15, July 2006.
- [2] R.E.H. Franich, "Disparity Estimation in Stereoscopic Digital Images," PhD thesis, Delft University of Technology, Department of Electrical Engineering, Information Theory Group, 1996.
- [3] A. Laurentini, "The visual hull concept for silhouette-based image understanding," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 16, no. 2, pp. 150-162, Feb. 1994.
- [4] P. Eisert, E. Steinbach and B. Girod, "Automatic reconstruction of stationary 3-D objects from multiple uncalibrated camera views," *IEEE Tran. Circuits Syst. Video Technol.*, vol.10, no. 2, pp.261-277, Mar. 2000.
- [5] L.S. Davis, ed., "Foundations of Image Understanding," © Kluwer, Boston, 2001,
- [6] K. N. Kutulakos, S. M. Seitz, "A Theory of Shape by Space Carving," *Int. J. Comp. Vision*, vol. 38, no.3, pp. 199-218, 2000.
- [7] P. Beardsley, P. Torr, and A. Zisserman, "3D model acquisition from extended image sequences," in *Proc. European Conf. Comp. Vision*, UK, pp. 683-695, 1996.
- [8] R. Koch, M. Pollefeys, L. Van Gool, "Robust Calibration and 3D Geometric Modeling from Large Collections of Uncalibrated Images," in *Proc. DAGM Pattern Recognition Symp.*, pp. 413-420, 1999.
- [9] P. Gargallo and P. Sturm, "Bayesian 3D Modeling from Images using Multiple Depth Maps," *Proc. CVPR'05*, vol.2, pp. 885-891, 2005.
- [10] B. Girod, A.M. Aaron, S.Rane, "Distributed video coding," *Proc. IEEE*, vol. 93, no.1, pp. 71-83, January 2005.
- [11] D. Slepian, J. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inform. Theory*, vol. 19, no.4, pp.471-480, July 1973.
- [12] A. D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the receiver," *IEEE Trans. Inform. Theory*, vol. 22, no.1, pp.1-10, January 1976.
- [13] N. Anantrasirichai, C. Nishan Canagarajah, David W. Redmill and David R. Bull, "Dynamic Programming for Multi-view Disparity/Depth Estimation," in *Proc. Int. Conf. Acoust., Speech, Signal Processing*, vol.2, 2006, pp. 269-272.
- [14] N. Anantrasirichai, C. Nishan Canagarajah, David W. Redmill and David R. Bull, "Colour Volumetric Compression for Realistic View Synthesis Applications," in *Proc. Int. Conf. Acoust., Speech, Signal Processing*, pp. I-1061-1064, Apr. 2007
- [15] N. Anantrasirichai, Nishan Canagarajah, and David R. Bull, "Multi-View Image Coding with Wavelet Lifting and In-Band Disparity Compensation," in *Proc. IEEE Int. Image Processing*, pp.33-36, September 2005.