REGION ADAPTIVE GRAPH FOURIER TRANSFORM FOR 3D POINT CLOUDS

Eduardo Pavez*, Benjamin Girault*, Antonio Ortega*, Philip A. Chou†

*University of Southern California, Los Angeles, California, USA

†Google, Inc., Mountain View, California, USA

ABSTRACT

We introduce the Region Adaptive Graph Fourier Transform (RA-GFT) for compression of 3D point cloud attributes. The RA-GFT is a multiresolution transform, formed by combining spatially localized block transforms. We assume the points are organized by a family of nested partitions represented by a rooted tree. At each resolution level, attributes are processed in clusters using block transforms. Each block transform produces a single approximation (DC) coefficient, and various detail (AC) coefficients. The DC coefficients are promoted up the tree to the next (lower resolution) level, where the process can be repeated until reaching the root. Since clusters may have a different numbers of points, each block transform must incorporate the relative importance of each coefficient. For this, we introduce the $\mathbf{Q}$ -normalized graph Laplacian, and propose using its eigenvectors as the block transform. The RA-GFT achieves better complexity-performance trade-offs than previous approaches. In particular, it outperforms the Region Adaptive Haar Transform (RAHT) by up to 2.5 dB, with a small complexity overhead.

Index Terms— graph fourier transform, 3D point cloud compression, graph signal, multiresolution transform, block transform

1. INTRODUCTION

Driven by applications in virtual and augmented reality, remote sensing, and autonomous vehicles, it is now possible to capture, in real time and at low cost, time varying 3D scenes, public spaces with moving objects, and people. The preferred representation for such data are 3D point clouds, which consist of 1) a list of *3D point coordinates*, and 2) *attributes* associated with those coordinates, such as color. In many applications, large point clouds need to be compressed for storage and transmission, leading to the recent development of a standard by the moving pictures expert group (MPEG) [1].

We present a new transform for attribute compression, which often takes up more than half of the overall bit budget for typical point clouds. Motivated by transforms used in image, video and point cloud compression [2, 3, 4], we construct our transform with the goal of achieving three fundamental properties: 1) orthogonality, 2) a constant basis function, and 3) low complexity. Orthogonality ensures that errors in the transform and point cloud attribute domains are equal. A constant basis function guarantees that an attribute with the same value at all points (the smoothest signal) has the most compact representation (1-sparse signal) in the transform domain. Finally, for the transform to scale to point clouds with a large number of points, N , we require a complexity of $\mathcal{O}(N \log N)$, instead of a naive implementation that uses matrix vector products, which would require $\mathcal{O}(N^2)$ complexity, if the transform is available explicitly, or $\mathcal{O}(N^3)$, if it has to be obtained via eigendecomposition.

We propose the *Region Adaptive Graph Fourier Transform (RA-GFT)*, a multiresolution transform formed by combining spatially localized block transforms, where each block corresponds to a cube in 3D space. The points are organized as a set of nested partitions represented by a tree. Leaf nodes represent points in the original point cloud, while each internal node represents all points within the corresponding subtree (see Fig. 1). Resolution levels are determined by the levels of the tree, with higher resolution corresponding to the deepest level (representing single points), and the coarsest resolution corresponding to the root (representing all points). At each resolution level, attributes belonging to the same group (points that have the same parent in the tree) are passed through an orthogonal transform for decorrelation. Each block transform produces a single approximation (DC) coefficient, and several detail (AC) coefficients. The DC coefficients are promoted to a lower resolution level where the same process can be repeated until reaching the root¹.

Since each internal node in the tree represents a set, possibly containing a different number of points, the block transforms should incorporate the relative importance of the nodes, based on their respective number of descendants. To address this issue, we propose a new graph Fourier transform (GFT) given by the eigenvectors of the $\mathbf{Q}$ -normalized graph Laplacian $\mathbf{L}_\mathbf{Q} = \mathbf{Q}^{-1/2} \mathbf{L} \mathbf{Q}^{-1/2}$, where $\mathbf{Q}$ is a diagonal matrix with diagonal entries containing the number of descendants of each node, and $\mathbf{L}$ is the combinatorial graph Laplacian. In contrast to the normalized or combinatorial Laplacian matrices [5], our new variation operator encodes the local geometry (distances between points) in $\mathbf{L}$ as well as the relative importance of a given set of points. The proposed transform is closely related to the Irregularity Aware Graph Fourier Transform (IAGFT) [6]. Our code is available at: <https://github.com/STAC-USC/RA-GFT>.

Multiresolution decompositions for point cloud coding have been proposed based on graph filter bank theory [7, 8, 9]. However, they often lack orthogonality, and the multiresolution representations are built through complex graph partitioning and reduction algorithms, which make them impractical for large point clouds. A Haar-like basis was proposed in [10] for any data that can be represented by a hierarchical tree. This construction is orthogonal and has a constant basis function. Although the other basis functions are spatially localized, they do not exploit local geometry information (distances between points). In addition, there is no efficient algorithm for computing transform coefficients, and matrix vector products need to be used. Also inspired by the Haar transform, [11] proposed sub-graph based filter banks, where a graph is partitioned into connected sub-graphs. For each sub-graph, a local Laplacian based GFT is used. Although [11] and the RA-GFT follow a similar strategy based on nested partitions, the graph construction, transforms, coefficient arrangement, and design goals are different.

¹This process can also stop before reaching the root, so that multiple subtrees, each with its own DC coefficient, are stored. This may be more efficient for large point clouds where there is limited correlation across subtrees.

The RA-GFT can be applied to any type of dataset as long as a nested partition is available. For 3D point clouds, a natural choice is the octree decomposition [12, 13], which can be used to implement RA-GFT for point clouds with $\mathcal{O}(N \log N)$ complexity. This data structure has already been used to design transforms for point cloud attributes [4, 14, 15, 16, 17]. In the block based graph Fourier transform (block-GFT) [14], the voxel space is partitioned into small blocks, a graph is constructed for the points in each block, and the corresponding graph Fourier (GFT) transform is used to represent attributes in the block. Another popular approach is RAHT [4], which is a multiresolution transform formed by the composition of 2×2 orthogonal transforms. The block-GFT can achieve excellent performance if the block size is large enough (more points per block), but this has a significant computational cost, since it requires computing GFTs of graphs with possibly hundreds of points. On the other hand, RAHT has an extremely fast implementation, with a competitive coding performance. Our proposed RA-GFT combines ideas from block-GFT and RAHT: it generalizes the block-GFT approach to multiple levels, while RAHT can be viewed as a special case of RA-GFT that is separable and uses only $2 \times 2 \times 2$ blocks. We demonstrate through point cloud attribute compression experiments that when RA-GFT is implemented with small block transforms, it can outperform RAHT by up to 2.5dB, with a small computational overhead. When the RA-GFT is implemented with larger blocks, we outperform the block-GFT with a negligible complexity overhead.

The rest of this paper is organized as follows. Section 2 introduces the RA-GFT for arbitrary datasets, while Section 3 explains how to implement it for 3D point clouds. Compression experiments and conclusions are presented in Sections 4 and 5 respectively.

Fig. 1: Nested partition represented by a hierarchical tree. Each level of the tree partitions the nodes in the next level. Thus at level $\ell = 1$, $\mathbf{v}_1^1$, $\mathbf{v}_2^1$, and $\mathbf{v}_3^1$ respectively represent the sets $\mathcal{V}_1^1 = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, $\mathcal{V}_2^1 = \{\mathbf{v}_4, \mathbf{v}_5, \mathbf{v}_6, \mathbf{v}_7\}$, and $\mathcal{V}_3^1 = \{\mathbf{v}_8, \mathbf{v}_9\}$, while at level $\ell = 0$, $\mathbf{v}_1^0$ represents the set $\mathcal{V}_1^0 = \{\mathbf{v}_1^1, \mathbf{v}_2^1, \mathbf{v}_3^1\}$.

2. REGION ADAPTIVE GRAPH FOURIER TRANSFORM

2.1. Notation and preliminaries

We use lowercase normal (e.g. α, a), lowercase bold (e.g. $\mathbf{v}$) and uppercase bold (e.g. $\mathbf{A}$) for scalars, vectors and matrices respectively. Vectors and matrices may also be denoted using their entries as $\mathbf{x} = [x_i]$, or $\mathbf{A} = [a_{ij}]$. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$ denote a weighted undirected graph with vertex set $\mathcal{V}$, edge set $\mathcal{E}$ and edge weight matrix $\mathbf{W} \geq 0$. An edge weight is positive, that is $w_{ij} = w_{ji} > 0$ if and only if $(i, j) \in \mathcal{E}$. The graph has $|\mathcal{V}| = n$ nodes. Let $\mathbf{D} = \text{diag}(d_i)$, with $d_i = \sum_j w_{ij}$, be the $n \times n$ degree matrix and let $\mathbf{L} = \mathbf{D} - \mathbf{W}$ be the $n \times n$ combinatorial graph Laplacian matrix. $\mathbf{L}$ is symmetric positive semidefinite, and has eigendecomposition $\mathbf{L} = \mathbf{\Phi} \mathbf{\Lambda} \mathbf{\Phi}^\top$, where the eigenvalues matrix is $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$, and the eigenvalues are $\lambda_1 = 0 \leq \lambda_2, \dots, \leq \lambda_n$. For connected graphs, $\lambda_2 > 0$. The eigenvector associated to λ_1 is $(1/\sqrt{n})\mathbf{1}$.

2.2. RA-GFT block transform

The Region-Adaptive Graph Fourier Transform (RA-GFT) is an orthonormal transform formed from the composition of smaller block transforms. We start by describing the latter. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W}, \mathbf{Q})$ denote a graph as defined in Section 2.1. In addition, define the $n \times n$ node weight matrix $\mathbf{Q} = \text{diag}(q_i)$ ($q_i > 0$), the $\mathbf{Q}$ -normalized Laplacian $\mathbf{L}_\mathbf{Q}$, and its eigendecomposition are given by

$$\mathbf{L}_\mathbf{Q} = \mathbf{Q}^{-1/2} \mathbf{L} \mathbf{Q}^{-1/2} = \mathbf{\Phi} \mathbf{\Lambda} \mathbf{\Phi}^\top, \quad (1)$$

where $\mathbf{\Phi}$ is the matrix of eigenvectors of $\mathbf{L}_\mathbf{Q}$ and $\mathbf{\Lambda} = \text{diag}(\lambda_i)$ is the matrix of eigenvalues. Since $\mathbf{L}_\mathbf{Q}$ is symmetric and positive semi-definite, $\lambda_i \geq 0$ and $\mathbf{\Phi}$ is orthonormal. Moreover, if we order the eigenvectors by their eigenvalues, and assume a connected graph, we have $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$ and the first eigenvector is proportional to $\phi_0 = [\sqrt{q_1}, \dots, \sqrt{q_n}]^\top$. Hence $\mathbf{\Phi}^\top$ maps the vector ϕ_0 to $[\sqrt{q_1 + \dots + q_n}, 0, \dots, 0]$. We define $\mathbf{\Phi}^\top$ to be the elementary block transform of the RA-GFT with inverse $\mathbf{\Phi}$.

2.3. Relation to other transforms

Relation to RAHT. As a special case, consider the two-node graph with $\mathcal{V} = \{\mathbf{v}_1, \mathbf{v}_2\}$, $\mathcal{E} = \{(1, 2)\}$, edge weights $w_{12} = w_{21} = 1$, and node weights $q_1 > 0$ and $q_2 > 0$. Then $\mathbf{L} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$, and $\mathbf{Q} = \text{diag}(q_1, q_2)$, hence the $\mathbf{Q}$ -normalized Laplacian is

$$\mathbf{L}_\mathbf{Q} = \begin{bmatrix} \frac{1}{q_1} & \frac{-1}{\sqrt{q_1 q_2}} \\ \frac{-1}{\sqrt{q_1 q_2}} & \frac{1}{q_2} \end{bmatrix} \quad (2)$$

$$= \begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 0 & \lambda_{\max} \end{bmatrix} \begin{bmatrix} a & b \\ -b & a \end{bmatrix}, \quad (3)$$

where $a = \sqrt{q_1}/\sqrt{q_1 + q_2}$, $b = \sqrt{q_2}/\sqrt{q_1 + q_2}$, and $\lambda_{\max} = q_1^{-1} + q_2^{-1}$. The 2×2 matrix $\mathbf{\Phi}^\top = \begin{bmatrix} a & b \\ -b & a \end{bmatrix}$ is the RAHT butterfly [4]. Hence RAHT is the 2×2 case of the RA-GFT.

Relation to IAGFT. Consider $\mathbf{Z} = \mathbf{Q}^{-1} \mathbf{L}$, the fundamental matrix of the $(\mathbf{L}, \mathbf{Q})$ -IAGFT [6]. Clearly $\mathbf{Z}$ is related to $\mathbf{L}_\mathbf{Q}$ by a similarity transform [6, Remark 1]: $\mathbf{Z} = \mathbf{Q}^{-1/2} \mathbf{L}_\mathbf{Q} \mathbf{Q}^{1/2}$. Thus $\mathbf{Z}$ and $\mathbf{L}_\mathbf{Q}$ have the same set of eigenvalues, $\mathbf{\Lambda}$, and the matrix $\mathbf{U}$ of eigenvectors of $\mathbf{Z}$ is related to $\mathbf{\Phi}$ by $\mathbf{U} = \mathbf{Q}^{-1/2} \mathbf{\Phi}$, that is, $\mathbf{Z} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^{-1} = \mathbf{Q}^{-1/2} (\mathbf{\Phi} \mathbf{\Lambda} \mathbf{\Phi}^\top) \mathbf{Q}^{1/2}$. It can be shown [6, Thm. 1] that the columns of $\mathbf{U}$ are orthonormal under the $\mathbf{Q}$ -norm, i.e., $\mathbf{U}^\top \mathbf{Q} \mathbf{U} = \mathbf{I}$. The IAGFT is defined as $\mathbf{F} = \mathbf{U}^{-1}$. So $\mathbf{F} = \mathbf{U}^\top \mathbf{Q} = \mathbf{\Phi}^\top \mathbf{Q}^{1/2}$, or $\mathbf{\Phi}^\top = \mathbf{F} \mathbf{Q}^{-1/2}$. Hence the IAGFT is a block transform of the RA-GFT applied to a $\mathbf{Q}^{1/2}$ -weighted signal.

2.4. Definition of full RA-GFT

We now show how the RA-GFT block transforms are composed to form the full RA-GFT. Consider a list of N points $\mathbf{V}$, a real-valued attribute signal $\mathbf{a}$ and node weights $\mathbf{q}$ on those points:

$$\mathbf{V} = [\mathbf{v}_i], \quad \mathbf{a} = [a_i], \quad \mathbf{q} = [q_i].$$

Here, $\mathbf{v}_i$ is an abstract point, and can be considered a node or vertex of a discrete structure. Now suppose we are given a hierarchical partition or nested refinement of the points, as illustrated in Fig. 1. Let L be the depth of the hierarchy, and for each level ℓ from the root ($\ell = 0$) to the leaves ($\ell = L$), let $\mathbf{v}_i^\ell$ be the i th of M_ℓ nodes at level ℓ . At level L , we have $M_L = N$ and $\mathbf{v}_i^L = \mathbf{v}_i$. For level $\ell < L$, let $\mathcal{C}_i^\ell \subseteq \{1, \dots, M_{\ell+1}\}$ denote the indices of the children of node $\mathbf{v}_i^\ell$.

Fig. 2: RA-GFT for point cloud represented by tree structured partition of Figure 1. Circles represent DC coefficients that are further processed by block transforms, while squares contain RA-GFT coefficients.

and let $\mathcal{D}_i^\ell \subseteq \{1, \dots, N\}$ denote the indices of the descendants of $\mathbf{v}_i^\ell$. We are also given for all $\ell < L$ and $i \in \{1, \dots, M_\ell\}$, a graph $\mathcal{G}_i^\ell = (\mathcal{V}_i^\ell, \mathcal{E}_i^\ell, \mathbf{W}_i^\ell, \mathbf{Q}_i^\ell)$, where $\mathcal{V}_i^\ell = \{\mathbf{v}_k^{\ell+1} : k \in \mathcal{C}_i^\ell\}$ is the set of children of node $\mathbf{v}_i^\ell$, $\mathcal{E}_i^\ell$ is a set of edges between the children, $\mathbf{W}_i^\ell$ is a matrix of edge weights, and $\mathbf{Q}_i^\ell = \text{diag}(q_k^{\ell+1} : k \in \mathcal{C}_i^\ell)$ is the diagonal matrix of child node weights, where $q_k^{\ell+1} = \sum_{j \in \mathcal{D}_k^{\ell+1}} q_j$ is the sum of the weights of all $n_k^{\ell+1} = |\mathcal{D}_k^{\ell+1}|$ descendants of child $\mathbf{v}_k^{\ell+1}$. Let $\Phi_{\ell,i}^\top$ be the RA-GFT block transform of the graph $\mathcal{G}_i^\ell$.

The full $N \times N$ RA-GFT $\mathbf{T}$ is a composition of $N \times N$ orthonormal transforms, $\hat{\mathbf{a}} = \mathbf{T}\mathbf{a} = \mathbf{T}_0\mathbf{T}_1 \cdots \mathbf{T}_{L-1}\mathbf{a}$, with inverse $\mathbf{a} = \mathbf{T}^{-1}\hat{\mathbf{a}} = \mathbf{T}_{L-1}^{-1} \cdots \mathbf{T}_1^{-1}\mathbf{T}_0^{-1}\hat{\mathbf{a}}$, where

$$\mathbf{T}_\ell = \mathbf{B}\text{diag}(\mathbf{T}_{\ell,1}, \dots, \mathbf{T}_{\ell,M_\ell}), \quad (4)$$

is an $N \times N$ orthonormal block diagonal matrix and block $\mathbf{T}_{\ell,i}$ is an $n_i^\ell \times n_i^\ell$ orthonormal matrix for transforming the n_i^ℓ descendants of the i th node $\mathbf{v}_i^\ell$ at level ℓ . Specifically,

$$\mathbf{T}_{\ell,i} = \mathbf{P}_{\ell,i}^{-1} \mathbf{B}\text{diag}(\Phi_{\ell,i}^\top, \mathbf{I}) \mathbf{P}_{\ell,i} \quad (5)$$

where $\mathbf{P}_{\ell,i}$ is an $n_i^\ell \times n_i^\ell$ permutation that collects the lowpass coefficients of the child nodes $\mathbf{v}_k^{\ell+1}$, $k \in \mathcal{C}_i^\ell$ for processing by the RA-GFT block transform $\Phi_{\ell,i}^\top$ to produce lowpass and highpass coefficients for parent node $\mathbf{v}_i^\ell$. When the RA-GFT is implemented using all levels of the tree, it produces a single approximation or low pass (DC) coefficient, and $N - 1$ detail or high pass (AC) coefficients. Figure 2 depicts the application of the RA-GFT for the nested partition depicted in Figure 1. It can be seen inductively that $\mathbf{T}$ maps the signal $[\sqrt{q_i}]$ to a single lowpass (DC) coefficient equal to $(\sum_i q_i)^{1/2}$, while all other, highpass (AC), coefficients are equal to 0. Thus the first, lowpass (DC) basis function of $\mathbf{T}$ is proportional to $[\sqrt{q_i}]$. In the usual case when $q_i = 1$ for all i , the first basis function of $\mathbf{T}$ is constant, as desired. This can be verified using Figure 2. Assume the attributes a_i and weights q_i are all equal to 1. Since the block transforms at level $\ell = 1$ have sizes 3, 4 and 2, the only non zero coefficients at level $\ell = 1$ are $\hat{a}_{1,0}^1 = \sqrt{3}$, $\hat{a}_{2,0}^1 = \sqrt{4}$ and $\hat{a}_{3,0}^1 = \sqrt{2}$. Then at level $\ell = 0$, the weight matrix is $\mathbf{Q}_0^0 = \text{diag}(3, 4, 2)$, hence the first column of $\Phi_{0,1}^\top$ is $(1/\sqrt{9})[\sqrt{3}, \sqrt{4}, \sqrt{2}]^\top$. Then the only nonzero transform coefficient produced at level $\ell = 0$ is $\hat{a}_{1,0}^0 = \sqrt{9}$. Therefore the RA-GFT of the $\mathbf{a} = \mathbf{1}$ vector is a 1-sparse vector.

3. APPLICATION TO POINT CLOUDS

In a point cloud, the vertices $\mathbf{V} = [\mathbf{v}_i]$ represent the coordinates (x, y, z) of real points in space; the attributes $\mathbf{a} = [a_i]$ represent colors or other attributes of the points; and the weights $\mathbf{q} = [q_i]$ represent the relative importance of the points. The weights are usually set to be constant ($q_i = 1$), but may be adjusted to reflect different

regions of interest [18]. We assume points are voxelized. A voxel is a volumetric unit of the domain of a 3-dimensional signal, analogous to a pixel in the 2-dimensional case. Let J be a positive integer, and partition the space into $2^J \times 2^J \times 2^J$ voxels. We say a point cloud is voxelized with depth J if all the point coordinates take values in the integer grid $\{0, 1, \dots, 2^J - 1\}$. A voxelized point cloud can be organized into a hierarchical structure. The process is described in [4, 17]. The voxel space $\{0, 1, \dots, 2^J - 1\}^3$ is hierarchically partitioned into sub-blocks of size $b_\ell \times b_\ell \times b_\ell$, where b_ℓ is a power of 2. These block sizes allow for a hierarchy with L levels, where $\prod_{\ell=1}^L b_\ell = 2^J$. Levels are ordered according to resolution. The partition is constructed by generating a point cloud for each resolution level. That is, beginning with $\mathbf{V}^L = \mathbf{V}$, we have for $\ell < L$:

$$\mathbf{V}^\ell = \text{unique} \left(\text{floor} \left(\frac{\mathbf{V}^{\ell+1}}{b_{\ell+1}} \right) \right), \quad (6)$$

where the unique function removes points with equal coordinates. Each point $\mathbf{v}_i^\ell$ in $\mathbf{V}^\ell$ has children

$$\mathcal{V}_i^\ell = \left\{ \mathbf{v}_k^{\ell+1} : \mathbf{v}_i^\ell = \text{floor} \left(\frac{\mathbf{v}_k^{\ell+1}}{b_{\ell+1}} \right) \right\}. \quad (7)$$

With the children we form a graph $\mathcal{G}_i^\ell = (\mathcal{V}_i^\ell, \mathcal{E}_i^\ell, \mathbf{W}_i^\ell, \mathbf{Q}_i^\ell)$, where there is an edge between nodes $\mathbf{v}_j^{\ell+1}, \mathbf{v}_k^{\ell+1} \in \mathcal{V}_i^\ell$ if the distance between $\mathbf{v}_j^{\ell+1}$ and $\mathbf{v}_k^{\ell+1}$ is less than a threshold, in which case the weight w_{jk} is set to a decreasing function of that distance. Using this hierarchy and set of graphs, the RA-GFT is constructed and applied to the attributes. The point hierarchies described in (6) and (7) can be obtained in $\mathcal{O}(N \log N)$ time [4, 17]. At resolution $\ell + 1$ we need to construct $M_\ell = \mathcal{O}(N)$ block transforms. For constant block sizes ($b_\ell = \mathcal{O}(1)$), each block transform can be obtained and applied in $\mathcal{O}(1)$ time. Since there are $L = \mathcal{O}(\log N)$ resolution levels, and each level is $\mathcal{O}(N)$, the RA-GFT has complexity $\mathcal{O}(N \log N)$.

4. EXPERIMENTS

In this section, we evaluate the RA-GFT in compression of color attributes of the “8iVFBv2” point cloud dataset² [19], which consists of four sequences: “longdress”, “redandblack”, “loot” and “soldier”, and compare its performance to block-GFT [14] and RAHT [4]. Colors are transformed from RGB to YUV spaces, and each of the Y, U, and V components are processed independently by the transform. For all transforms, we perform uniform quantization and entropy code the coefficients using the adaptive run-length Golomb-Rice algorithm [20]. Distortion for the Y component is given by

$$PSNR_Y = -10 \log_{10} \left(\frac{1}{T} \sum_{t=1}^T \frac{\|\mathbf{Y}_t - \hat{\mathbf{Y}}_t\|_2^2}{255^2 N_t} \right),$$

²<https://jpeg.org/plenodb/pc/8ilabs/>

RA-GFT $b_L = 2$	RA-GFT $b_L = 4$	RA-GFT $b_L = 8$	RA-GFT $b_L = 16$	block-GFT $b = 8$	block-GFT $b = 16$
0.2098×10^3	0.5243×10^3	6.12×10^3	104.25×10^3	5.97×10^3	104.13×10^3

Table 1: Complexity estimate of the RA-GFT and block-GFT as a function of the block size.

Fig. 3: Distortion rate curves for color compression.

where $T = 300$ is the number of frames in the sequence, N_t is the number of points in the t -th frame, and $\mathbf{Y}_t$ and $\hat{\mathbf{Y}}_t$ represent the original and decoded signals of frame t . Rate is reported in bits per voxel [bpv] $R = \left(\sum_{t=1}^T B_t \right) / \left(\sum_{t=1}^T N_t \right)$, where B_t is the number of bits used to encode the YUV components of the t -th frame.

4.1. Compression of color attributes

Each point cloud in the 8iVFB dataset is voxelized with depth $J = 10$. Therefore $L \leq 10$ and its value will depend on the block sizes. We implement several RA-GFTs, each with a different block size at the highest resolution (level L), but with the same block sizes $b_\ell = 2$ for all other resolutions ($\ell < L$). When b_L is equal to 2, 4, 8, and 16, L is equal to 10, 9, 8 and 7 respectively. For the block-GFT we choose block sizes $b = 8$ and $b = 16$. Graphs are constructed by adding edges if the distance between a pair of point coordinates is

below a fixed threshold, while edge weights are the reciprocal of the distance, as in [14]. Distortion rate curves are shown in Figure 3.

The RA-GFT provides substantial gains over RAHT. When the block size is smallest $b_L = 2$, the corresponding RA-GFT outperforms RAHT up to 0.5db for the “longdress” sequence, and up to 1db for the “loot” sequence. Similar results were obtained for other sequences, not shown due to lack of space. Coding performance improves as the block size b_L increases, up to 2.5 dB over RAHT on both sequences. The block-GFT also has this property, however, it requires a large block size ($b = 16$) to consistently outperform the RAHT for all sequences. This could occur because for smaller blocks, the block-GFT DC coefficients may still be highly correlated. Since the RA-GFT with small blocks can be viewed as an extension of the block-GFT to multiple levels, the transform coefficients of the proposed transform are less correlated.

4.2. Complexity analysis

At each level of RA-GFT, multiple GFTs of different sizes are constructed, so that the overall computational complexity is dominated by the number and size of those transforms. At resolution level ℓ , the i th transform is a $|\mathcal{V}_i^\ell| \times |\mathcal{V}_i^\ell|$ matrix, which is obtained by eigen-decomposition in roughly $|\mathcal{V}_i^\ell|^3$ operations. As a proxy for the total number of operations required by the RA-GFT we use

$$K = \sum_{\ell=0}^{L-1} \sum_{i=1}^{M_\ell} |\mathcal{V}_i^\ell|^3. \quad (8)$$

Eq. (8) is consistent with Section 3, since $|\mathcal{V}_i^\ell|^3 = \mathcal{O}(1)$, $M_\ell = \mathcal{O}(N)$ and $L \leq \log(N)$. We consider a collection of T point clouds. The t -th point cloud has N_t points, and the quantity (8) for the RA-GFT on that point cloud is denoted by K_t . The complexity proxy for the RA-GFT with a given tree structured nested partition is defined as $C = (\sum_{t=1}^T K_t) / (\sum_{t=1}^T N_t)$. We compute this quantity for the RA-GFT and block-GFT depicted in Figure 3. Our results are shown in Table 1, and are computed from the first 10 point clouds of each sequence of the “8iVFB” dataset, for a total of $T = 40$ point clouds. For $b_L = 8$, the increase in complexity from the block-GFT to the RA-GFT is only 2.5%; for $b_L = 16$ the increase is negligible. More importantly, for smaller blocks, the complexity of the RA-GFT is orders of magnitude lower than of block-GFT.

5. CONCLUSION

By allowing multiple block sizes, and multiple levels of resolution, the proposed RA-GFT can be viewed as a generalization of the block-GFT and the RAHT, reaching coding efficiency comparable to block-GFT, with computational complexity slightly higher than RAHT. Therefore, the RA-GFT achieves better performance-complexity trade-offs than these transforms. By using a non-separable transform on larger blocks the RA-GFT can exploit local geometry more efficiently than the RAHT. By applying transforms with small blocks at multiple resolutions, the RA-GFT can approach the performance of the block-GFT with a reduced complexity. In addition, for large transform blocks at resolution L , the RA-GFT outperforms the block-GFT, with a negligible complexity increase.

6. REFERENCES

- [1] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. Cesar, P.A. Chou, R.A. Cohen, M. Krivokuća, S. Lasserre, and Z. Li, "Emerging MPEG standards for point cloud compression," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 1, pp. 133–148, 2018.
- [2] M. Vetterli and J. Kovacevic, *Wavelet and Subband Coding*, Englewood Cliffs, NY: Prentice-Hall, 1995.
- [3] N. Ahmed, T. Natarajan, and K.R. Rao, "Discrete cosine transform," *IEEE transactions on Computers*, vol. 100, no. 1, pp. 90–93, 1974.
- [4] R.L. De Queiroz and P.A. Chou, "Compression of 3D point clouds using a region-adaptive hierarchical transform," *IEEE Transactions on Image Processing*, vol. 25, no. 8, pp. 3947–3956, 2016.
- [5] D.I. Shuman, S.K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *Signal Processing Magazine, IEEE*, vol. 30, no. 3, pp. 83–98, May 2013.
- [6] B. Girault, A. Ortega, and S.S. Narayanan, "Irregularity-aware graph fourier transforms," *IEEE Transactions on Signal Processing*, vol. 66, no. 21, pp. 5746–5761, 2018.
- [7] A. Anis, P.A. Chou, and A. Ortega, "Compression of dynamic 3D point clouds using subdivisional meshes and graph wavelet transforms," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6360–6364.
- [8] D. Thanou, P.A. Chou, and P. Frossard, "Graph-based compression of dynamic 3D point cloud sequences," *IEEE Transactions on Image Processing*, vol. 25, no. 4, pp. 1765–1778, 2016.
- [9] P.A. Chou, M. Koroteev, and M. Krivokuća, "A volumetric approach to point cloud compression, part I: Attribute compression," *IEEE Transactions on Image Processing*, 2019.
- [10] M. Gavish, B. Nadler, and R.R. Coifman, "Multiscale wavelets on trees, graphs and high dimensional data: theory and applications to semi supervised learning," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010, pp. 367–374.
- [11] N. Tremblay and P. Borgnat, "Subgraph-based filterbanks for graph signals," *IEEE Transactions on Signal Processing*, vol. 64, no. 15, pp. 3827–3840, 2016.
- [12] C.L. Jackins and S. L. Tanimoto, "Oct-trees and their use in representing three-dimensional objects," *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 249–270, 1980.
- [13] D. Meagher, "Geometric modeling using octree encoding," *Computer graphics and image processing*, vol. 19, no. 2, pp. 129–147, 1982.
- [14] C. Zhang, D. Florencio, and C. Loop, "Point cloud attribute compression with graph transform," in *2014 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2014, pp. 2066–2070.
- [15] R.A. Cohen, D. Tian, and A. Vetro, "Attribute compression for sparse point clouds using graph transforms," in *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2016, pp. 1374–1378.
- [16] R.A. Cohen, D. Tian, and A. Vetro, "Point cloud attribute compression using 3-D intra prediction and shape-adaptive transforms," in *2016 Data Compression Conference (DCC)*. IEEE, 2016, pp. 141–150.
- [17] E. Pavez, P.A. Chou, R.L. De Queiroz, and A. Ortega, "Dynamic polygon clouds: representation and compression for VR/AR," *APSIPA Transactions on Signal and Information Processing*, vol. 7, 2018.
- [18] G. P. Sandri, V. Figueiredo, P. A. Chou, and R. L. de Queiroz, "Point cloud coding incorporating region of interest coding," in *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019.
- [19] E. d'Eon, B. Harrison, T. Myers, and P.A. Chou, "8i voxelized full bodies-a voxelized point cloud dataset," *ISO/IEC JTC1/SC29 Joint WG11/WG1 (MPEG/JPEG) input document WG11M40059/WG1M74006*, 2017.
- [20] H.S. Malvar, "Adaptive run-length/Golomb-Rice encoding of quantized generalized gaussian sources with unknown statistics," in *Data Compression Conference (DCC'06)*. IEEE, 2006, pp. 23–32.