

Rate-Distortion Optimized Wavelet-based Irregular Mesh Coding

Jonas El Sayeh Khalil¹, Adrian Munteanu² and Peter Lambert¹

¹ELIS Department, IDLab, Ghent University-iMinds, Sint-Pietersnieuwstraat 41, B-9000, Ghent, Belgium

²Dept. of Electronics and Informatics, Vrije Universiteit Brussel, Pleinlaan 2, B-1050, Brussels, Belgium

Keywords: Quality-scalable Mesh Representation, Wavelet-based Mesh Coding, Rate-Distortion Optimization.

Abstract: This work investigates the optimization of mesh quality at lossy rates for a lossless scalable wavelet-based irregular mesh codec. Whereas previously proposed wavelet-based irregular mesh codecs offer coarse-grain resolution scalability, in this paper we propose a coding scheme which enables fine-grain quality scalability. This is done by avoiding the use of geometric data in the encoding process, which reduces dependencies within the data stream and allows for an unrestricted storage and transmission order of wavelet subband bitplanes and connectivity information. This in turn allows us to perform rate-distortion optimization, whereby the subband bitplanes to be encoded are determined by minimizing distortion subject to an overall target bitrate. Experimental results show that the proposed coding approach offers fine-grain quality scalability, achieves optimality in rate-distortion sense and improves compression performance over the state of the art.

1 INTRODUCTION

As the amount and the quality of obtained media increases, so does the need for efficient compression in multimedia domains. This need ignited advancements in subsequently audio, image, and video compression systems. Today, advanced compression solutions are also required for three-dimensional data. While the gaming industry tries to cope with increasing quality requirements by making use of shader tricks, and the movie industry makes use of huge rendering farms for their rendering needs, we observe that, in general, a scalable coding solution is required. With the rise of virtual and augmented reality applications, the advent of 3D printing and the increasing detail of 3D scanners comes the need for lossless storage and transmission, while offering interactivity in a scalable way.

Wavelet-based techniques have been successful in multimedia compression, including several 3D codecs. A wavelet-based solution uses a set of high-pass and low-pass filters to obtain a low-resolution base mesh where all high-frequency data have been removed, and a set of wavelet subbands containing increasingly higher frequency information. By exploiting fine-grain quality scalability, which allows for scaling the quality of reconstructed data by decoding per wavelet subband bitplane, data can be transmitted such that the distortion in the reconstructed mesh decreases optimally.

Contributions. In this work, we investigate our previous wavelet-based irregular mesh codec offering resolution scalability and propose a coding scheme based on it which offers quality scalability. This improves over the state of the art in two ways: (1) the *coding performance* at low bitrates is improved by our proposed algorithm by performing rate-distortion optimization, and (2) *functionally*, an additional form of scalability is offered without negatively impacting the lossless coding rate.

The remainder of this work is structured as follows: related work is given in Section 2, our resolution-scalable codec we improve upon is described in Section 3, our contributions are detailed in Sections 4 and 5 and we evaluate the proposed codec in Section 6. Section 7 concludes this work.

2 RELATED WORK

Until today, triangle meshes have been the main representation of 3D models for real-time rendering. A mesh is often seen as the combination of both *geometry* information, i.e., the positions of the vertices, and *connectivity* or *topology*, i.e., the edges (and consequently the triangles) between these vertices. Mesh coding has been an active topic of research for over two decades. State-of-the-art single-rate coders, e.g., (Touma and Gotsman, 1998), proved to be insufficient

despite good performance, as single-rate codecs only allow for decoding an *entire* mesh. An ad-hoc solution is the use of *Levels of Detail* (LODs), obtained by simplifying a mesh several times using a decreasing amounts of vertices to approximate the mesh. This allows for improved interactivity but increases the storing cost as each LOD is stored separately.

Pioneering work in *progressive mesh representations* was done by Hoppe (Hoppe, 1996). In this work, Hoppe describes how a mesh can be simplified vertex per vertex, minimizing an energy function at every step. This results in a so-called *continuous LOD* chain: this progressive mesh representation generates a nearly continuous spectrum of LODs where each new level is obtained by splitting one vertex of the previous level. Each obtained level is optimal given a specific number of vertices. Further improvements in compression performance was obtained by grouping the individual increments in batches, each batch splitting half of the decoded vertices (Pajarola and Rossignac, 2000). (Peng and Kuo, 2005) and (Valette et al., 2009) present state-of-the-art codecs which allow for fine-grain vertex-by-vertex resolution control.

The idea of constructing scalable representations and compression systems is well-known in signal processing. In this context, wavelets play a major role, being used to generate multiresolution representations of an input signal and to build scalable codecs based on them. Wavelet-based scalable codecs include well-known examples for images (Said and Pearlman, 1996; Munteanu et al., 1999; Taubman and Marcellin, 2001), for video (Taubman and Zakhor, 1994; Ohm, 1994; Andreopoulos et al., 2004), and were introduced for surfaces by Lounsbery et al. (Lounsbery et al., 1997). Essentially, Lounsbery et al. established the link between subdivision schemes and multiresolution analysis for meshes. Subdivision schemes result in *semi-regular* meshes, however, whereas models are most efficiently represented using *irregular* meshes, allowing for adaptive sampling of a surface. This avoids oversampling near areas lacking high frequencies, and undersampling where high frequencies are present. One solution is to convert irregular meshes to semi-regular ones in a preprocessing step (Payan et al., 2015). Semi-regular mesh coding allows for superior compression performance due to the implicit connectivity knowledge (Khodakovsky et al., 2000; Denis et al., 2010). However, not every application allows for a lossy remeshing step, and the inherent oversampling associated with semi-regularity results in a larger number of vertices and, consequently, a higher rendering cost.

Few wavelet transforms for irregular meshes have been proposed in the literature. Bonneau describes a

generalization of Haar-wavelets in (Bonneau, 1998). Valette et al. (Valette et al., 1999; Valette and Prost, 2004a) propose an extension of the work of Lounsbery et al., describing a new subdivision approach which is not restricted to 1-to-4 subdivision, but allows creating one, two, three or four triangles from a single triangle. In (Valette and Prost, 2004b), Valette et al. describe how this wavelet transform can be efficiently encoded, resulting in the state-of-the-art *Wavemesh* codec. Recently we proposed a wavelet-based irregular mesh codec (El Sayeh Khalil et al., 2016) which we will refer to as *EMD+16* in the remainder of this paper. This codec produced improved coding performance over *Wavemesh* in low and midrange bitrates, and additionally tackled a new untouched issue: whereas scalable mesh coding mainly focuses on minimizing the distortion for a given *amount of bits*, efficient rendering requires minimizing the distortion for a given *triangle limit*, which is directly related to a given memory limit.

In this paper we build upon *EMD+16*. Observing that the decoding of higher-range bits only yields marginal improvements in visual quality, we focus on the low and midrange rates and examine the possibility and the effect of performing a rate-distortion optimized storage and transmission of wavelet subband bitplanes. Rate-distortion performance is optimized by identifying at every encoding step the wavelet subband bitplane which maximally reduces distortion in the reconstructed mesh at minimum cost in bits.

3 FEATURE-PRESERVING IRREGULAR MESH CODING

In (El Sayeh Khalil et al., 2016), we describe *EMD+16*, a wavelet-based irregular mesh coding system with implicit preservation of important visual features. We reiterate the most important aspects of this codec, of which an overview is given in Figure 1.

3.1 Wavelet Transform

The wavelet transform is based on the lifting scheme proposed by Sweldens (Sweldens, 1998). This scheme is composed of three steps to perform the forward wavelet transform.

Downsampling and Retriangulation. A low-pass filtering step splits vertices into even and odd vertices. Each odd vertex is surrounded by even neighbouring vertices, forming a patch. After downsampling, the remaining polygonal patches are retriangulated while implicitly preserving geometric features,




Figure 1: Overview of the original mesh coder. The transform step generates a base mesh and a sequence of wavelet subbands. In this work, we modify the coding components to allow for an unrestricted storage and transmission order.

resulting in a lower-resolution approximation of the higher-resolution mesh. Inversely, the upsampling step adds a vertex per patch, and trivially retriangulates each patch by connecting each border vertex with the newly added odd vertex.

Prediction Step. The high-frequency data is predicted using the local patch information, determining a prediction error per patch. To allow the same predictions during reconstruction, additional information is required to determine the low-resolution patches. Additional connectivity data is common to all irregular mesh codecs, but the higher cost *per triangle* is compensated by requiring fewer triangles overall.

Update Step. The low-frequency data can finally be post-processed to avoid aliasing artefacts. In the given codec however, this step is chosen to be a null-operation to avoid smoothing artefacts around geometric features which should remain sharp.

3.2 Wavelet Subband Coder

The incremental data consists of two key components. Firstly, the connectivity information has to be encoded to allow for locating the patches. For every edge of the intermediate mesh, a single bit describes if the edge was a *new* edge which was constructed by retriangulation, or if the edge was present in the higher-resolution mesh. Secondly, the wavelet coefficients are represented as vectors in global orthonormal space and indicate the prediction errors of the odd vertices. Due to *Successive Approximation Quantization* (SAQ), the wavelet coefficients can be coded bitplane by bitplane. The quantizers are defined as:

$$Q_p(w) = \begin{cases} s(w) \left\lfloor \frac{|w|}{2^p \Delta} + \frac{\xi}{2^p} \right\rfloor, & \text{if } \frac{|w|}{2^p \Delta} + \frac{\xi}{2^p} > 0 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

with $2^p \Delta$ the quantization cell width, ξ the width of the deadzone, and $s(\cdot)$ the sign function.

The connectivity information and wavelet coefficients are both encoded using an octree structure.

This octree structure decouples the encoding process from any mesh traversal order and allows for exploiting spatial locality. Within the octrees, the connectivity samples are encoded at the edge middles, while the geometry samples are encoded at the predicted positions. For the encoder and decoder to operate in a synchronized fashion, a template mesh is used to map the samples. (El Sayeh Khalil et al., 2016) does not give requirements for this, but indicates that the *previous resolution* is used for the template mesh.

4 GEOMETRY-INDEPENDENT TEMPLATE MESHES

The original design of EMD+16 allows for resolution scalability and only for a limited form of quality scalability: at each resolution, the decoder can determine the amount of quality bits for the newly added vertices; however, a new resolution can only be decoded after fully decoding the previous resolution. This is a disadvantage of using this previous resolution as a template mesh for coding the connectivity information and wavelet coefficients of the next resolution. Dependencies between data blocks are shown in Figure 2(a), clearly showing that this approach allows only one order in which data can be transmitted, depicted in Figure 2(b).

Instead, we explore quality-scalable encoding of irregular meshes where the wavelet subbands are encoded at possibly different quality levels. Moreover, the subband bitplanes are encoded in a rate-distortion optimal manner, ensuring minimal distortion in the reconstructed mesh for any target bitrate.

We propose to use the base mesh as a first template mesh again. However, instead of reconstructing each resolution using its previous resolution, effectively doing an identical inverse transform on the template mesh as performed on the actual mesh, we propose to perform a modified inverse transform which only uses connectivity information. This results in a template mesh which increases in resolution by *only predicting* the added vertices without requiring the decoded

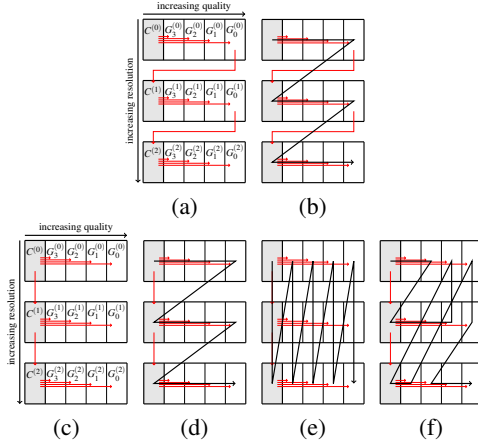


Figure 2: Data dependencies and coding orders. Fig. 2(a) shows dependencies in EMD+16, and Fig. 2(b) the only possible coding order; Fig. 2(c) shows dependencies in the proposed framework, Figs. 2(d), 2(e) and 2(f) possible coding orders. For resolution j , $C^{(j)}$ represents the connectivity data and $G_i^{(j)}$ the geometry data at bitplane i .

wavelet coefficients to reposition these vertices. The result is that the connectivity and geometry of a resolution can be decoded as soon as the connectivity of the previous resolution is known. The resulting dependencies are depicted in Figure 2(c): the encoder is able to store the data blocks in an unrestricted order as long as (a) the order of connectivity information blocks is maintained, (b) geometry information within each resolution is stored in the correct order, and (c) connectivity information for a specific resolution is encoded before geometry information. This allows for storing and transmitting the blocks in any order: resolution per resolution as before (Fig. 2(d)), bitplane by bitplane or purely quality-scalable (Fig. 2(e)), or in any arbitrary order that preserves the dependencies as indicated above (e.g., Fig. 2(f)).

5 RATE-DISTORTION OPTIMIZATION

Rate-distortion optimization requires encoding data blocks such that the distortion is minimal at all bitrates. Such distortion optimization depends on the distortion measure used and as such does not yet have an unambiguous solution. Our aim is to show that such optimizations are indeed *possible* by proposing a rate-distortion optimization algorithm, proving that an optimized subband bitplane storage and transmission order is enabled by our codec architecture.

We construct our optimization algorithm by considering the distortions introduced by the wavelet transform. With $N^{(j)}$ the number of wavelet coeffi-

cients for resolution j , the remaining distortion $D_p^{(j)}$ related to this j^{th} resolution decoded up to bitplane p is given by:

$$D_p^{(j)} = \sum_{i=1}^{N^{(j)}} \alpha_i^{(j)} d_{i,p}^{(j)}, \quad (2)$$

with $d_{i,p}^{(j)}$ the distortion on the odd vertex $o_i^{(j)}$ when the most significant bits of wavelet coefficient i of resolution j are decoded up to the p^{th} bitplane, i.e.,

$$\begin{aligned} d_{i,p}^{(j)} &= |\mathbf{x}_i^{(j)} - \tilde{\mathbf{x}}_{i,p}^{(j)}| \\ &= |\mathbf{x}_i^{(j)} - (\tilde{\mathbf{x}}_i^{(j)} + \mathbf{w}_{i,p}^{(j)})| \end{aligned} \quad (3)$$

In Equation 3, $\tilde{\mathbf{x}}_i^{(j)}$ is the predicted position of odd vertex $o_i^{(j)}$ and $\tilde{\mathbf{x}}_{i,p}^{(j)}$ is its reconstructed position when decoding the most significant bitplanes of the accompanying wavelet coefficient $\mathbf{w}_i^{(j)}$ up to bitplane p . To simplify notations, we drop the superscript j and subscript i as the following paragraphs always handle a specific odd sample i of a specific resolution j .

If we denote the total number of bitplanes by p_{\max} , then $\mathbf{w}_{p_{\max}} = 0$, i.e., when no bitplanes are encoded, the wavelet coefficient is zero and the prediction is not corrected: $d_{p_{\max}} = |\mathbf{x} - \tilde{\mathbf{x}}|$. The distortion becomes exactly 0 as soon as the least-significant bitplane is decoded: $\mathbf{w}_0 = \mathbf{w} = \mathbf{x} - \tilde{\mathbf{x}}$ and $d_0 = |\mathbf{x} - (\tilde{\mathbf{x}} + \mathbf{w}_0)| = |\mathbf{x} - (\tilde{\mathbf{x}} + \mathbf{x} - \tilde{\mathbf{x}})| = 0$.

The wavelet coefficients are encoded using SAQ, so let $q_p = Q_p(w)$; the wavelet coefficient up to bitplane p can be dequantized as $w_p = Q_p^{-1}(q_p)$, with:

$$Q_p^{-1}(q_p) = \begin{cases} 0, & q_p = 0 \\ s(q_p) \left(|q_p| - \frac{\xi}{2^p} + \delta \right) 2^p \Delta, & q_p \neq 0 \end{cases} \quad (4)$$

where $0 \leq \delta < 1$ determines the placement of w_p within the quantization cell. In the current implementation, it is chosen to be $\delta = 0.5$.

To find weights α for the wavelet coefficients, we first assign weights β_k to *all* vertices. These weights indicate an estimation of the effect on the full-resolution mesh of repositioning vertices. After downsampling, the position of each removed odd vertex $o_i^{(j)}$ is influenced by the positions of its neighbours. Hence, the weights of these neighbour vertices of o each increase by β_o/d with d the amount of neighbours. The weight of a wavelet coefficient is given by its accompanying odd vertex: $\alpha = \beta_o$.

At the highest resolution, each vertex (and consequently each wavelet coefficient) has a weight of $1/N$. If the downsampling terminates at a tetrahedral base mesh counting four vertices, each vertex (and consequently each wavelet coefficient on average) has a weight of $1/4$. This indicates that a single wavelet




Figure 3: Generic rate-distortion curve. With decreasing bitplane p the rate increases and the distortion decreases.

coefficient at a lower resolution is equally valuable as a multitude of wavelet coefficients at a higher resolution. Note that the use of such weights is similar to the scaling of wavelet coefficients prior to encoding – see e.g. (Said and Pearlman, 1996), rendering biorthogonal transforms approximately unitary.

The complete algorithm for assigning weights is presented in Algorithm 1.

Algorithm 1: Assigning weights to vertices.

```

1: for all  $v \in M$  do
2:    $v.\beta \leftarrow 1/N$ 
3: end for
4: while downsampling do
5:   for all  $v_o \in M_{i,o}$  do
6:      $d \leftarrow \text{valence}(v_o)$ 
7:     for all  $v \in \text{neighbourhood}(v_o)$  do
8:        $v.\beta \leftarrow v.\beta + v_o.\beta/d$ 
9:     end for
10:  end for
11: end while
    
```

At each resolution, a rate-distortion curve such as shown in Figure 3 can be found. Optimization comes down to considering the rate-distortion curves for every resolution, and coding at every step the information which introduces the largest distortion decrease at the lowest rate. With P_j the last encoded bitplane of resolution j , we would encode the next bitplane of resolution j' with

$$j' = \arg \max_{j:P_j \neq 0} \frac{D_{P_j}^{(j)} - D_{P_{j-1}}^{(j)}}{R_{P_{j-1}}^{(j)} - R_{P_j}^{(j)}} \quad (5)$$

Up to this point, we did not handle connectivity blocks. To start decoding a specific resolution, the connectivity information of all previous resolutions has to be decoded. This decoding comes at a rate but does not introduce a distortion decrease in mean-squared-error sense. Furthermore, as the most-significant bitplanes are mostly zero, decoding these highest bitplanes also requires rate often without decreasing distortion either. Hence, Equation 5 is adapted to consider multiple bitplanes together to find the most optimal slope. We generalize the definition of P_j to encompass *any* data block required to encode

resolution j . Consequently, P_j will start at $p_{\max} + 1$ as it counts for all p_{\max} wavelet subband bitplanes and an additional data block for the connectivity information. With L the first unencoded resolution, we will encode k' bitplanes of resolution j' using:

$$(j', k') = \arg \max_{j \in [0, L], k \in [1, P_j]} \frac{D_{P_j}^{(j)} - D_{P_{j-k}}^{(j)}}{R_{P_{j-k}}^{(j)} - R_{P_j}^{(j)}} \quad (6)$$

For this, we make two conventions:

- $R_{p_{\max}+1}^{(j)} = 0; R_{p_{\max}}^{(j)} = R_{\text{conn}} \neq 0$
- $D_{p_{\max}+1}^{(j)} = D_{p_{\max}}^{(j)}$

This states that encoding the connectivity information, i.e., the first data block of a resolution, introduces a rate of $R_{\text{conn}}^{(j)}$ while not decreasing the distortion; this corresponds to adding vertices to the decoded mesh without refining their locations. The distortions $D_p^{(j)}, \forall p \in [0, p_{\max}]$ are defined in Equation 2.

6 EXPERIMENTAL EVALUATION

To evaluate our codec, we assume a raw storage of the base mesh. In practice, the wavelet transform stops at a resolution which is small enough to benefit from scalability, while remaining qualitative enough to be used as a base resolution. To evaluate, however, the transform is applied until no more downsampling can be performed. The base mesh is included in the results by counting $B = 3Qn_v + 3 \lceil \log_2(n_v) \rceil n_t$ bits, with Q the amount of quantization bits per vertex, n_v the number of vertices in this base mesh and n_t the number of triangles. The base meshes were clearly negligible, hence the produced coding numbers can be entirely ascribed to the wavelet subband codec itself.

We have used the METRO tool (Cignoni et al., 1998) to obtain the distortion values, reported as *root-mean-squared* (RMS) errors, while rates are expressed in amount of *bits per vertex* (bpv). The resulting rate-distortion curves have been made convex by removing non-convex ratepoints.

6.1 Accuracy of the Template Mesh

Using a template mesh which does not take into account geometry information, i.e., a template which develops using only connectivity information, had an effect on the implementation: possibly overlapping vertices in the template mesh have to be handled properly, by repositioning these vertices equally by both the encoder and decoder *without relying on the encounter order*. Note that these geometric modifications do not alter the geometry of the decoded meshes

Table 1: Additional cost when using an accurate template. All models have been encoded using 12 bit precision.

Model (#verts)	Rate increase
teapot (1, 292)	+0.068bpv (+0.21%)
beethoven (2, 521)	+0.187bpv (+0.55%)
triceratops (2, 832)	-0.017bpv (-0.052%)
elk (5, 194)	+0.015bpv (+0.049%)
fandisk (6, 475)	-0.015bpv (-0.057%)
maxplanck (7, 399)	-0.013bpv (-0.044%)
venushead (8, 268)	+0.022bpv (+0.074%)
bimba (8, 857)	+0.094bpv (+0.31%)
horse (19, 851)	+0.022bpv (+0.087%)
screwdriver (65, 538)	+0.182bpv (+0.89%)
rabbit (67, 039)	+0.073bpv (+0.32%)
dino (129, 026)	+0.074bpv (+0.37%)
Average	+0.23%

after inverse wavelet transformation; they can only alter where samples are encoded in the octrees.

Both the use of this updated template mesh as well as the occasional repositioning of overlapping sample positions can have an influence on the exploitability of the spatial correlations within the data. However, the effects of using the new geometry-agnostic template proved to be negligible.

This can also be seen in the final bitrates. Table 1 lists the changes in lossless bitrates for several models. We observe that the bitrate increases on average by 0.23%, while in some cases the bitrate even goes down. These results show that lowering the accuracy of the template mesh does not hinder the exploitation of spatial correlations, on the contrary it is clear that topological locality information is preserved. Vertices that are located closely together will have mapped vertices in the template mesh which are also located closely together due to topological proximity, albeit possibly at another global position (hence, falling within a different octree cell) due to geometry information not being taken into account.

6.2 Rate-distortion Optimization

Rate-distortion optimization orders the data such that the quality gains come at the lowest rates. The results for the three models considered in (El Sayeh Khalil et al., 2016) are given in Figure 4. For *fandisk* and *horse* the improvements compared with a resolution-scalable transmission order are clear at low bitrates, obtaining similar qualities at a lower cost. At higher bitrates the improvements are minimal. This indicates that the resolution-scalable transmission at high bitrates is in general already nearly optimal in rate-distortion sense.

Visual results of our proposed codec are given in Figure 5. As it was also observed in (El Sayeh Khalil

et al., 2016), the information in the lower resolution wavelet subbands contributes the most to the shape of such densely sampled models, while higher resolutions only serve to increase the quality when rendering from very nearby: Figure 5(c) already resembles very closely the original mesh depicted in Figure 5(d).

In theory, a rate-distortion optimization should *always* be on par or better; in practice this is the case if the optimization algorithm calculates distortions in the spatial domain. Our results show that an optimization based on the proposed weighted wavelet coefficients while ignoring topological information also proves to give superior results. The coding overhead is minimal: the storage increases only a fraction of a bit per vertex, due to the need to identify the resolution of each subsequent data block.

6.3 State of the Art Comparison

For comparison, we mainly investigated the overall improvement over *EMD+16*. As in (El Sayeh Khalil et al., 2016), we also compare with Wavemesh (Valette and Prost, 2004b) and IPR (Valette et al., 2009). *EMD+16* was implemented using ratepoints at every resolution; this is more coarse-grain than our proposed codec which gives a ratepoint after every bitplane. For Wavemesh we made use of the publicly available software with *wavelet geometrical criterion* enabled. We also made a comparison with IPR for the models for which we received their decoded results.

Figure 6 shows a comparison with the state of the art. It shows even more competitive results brought by the proposed codec compared with *EMD+16*. In the case of the feature-rich model *fandisk*, our results after rate-distortion optimization improve over both Wavemesh and IPR. In the case of *horse*, the results are entirely on-par with or better than previous work, even at the lowest rates. Finally, for the feature-poor model *rabbit*, results remain nearly unchanged compared with *EMD+16*; IPR remains the better solution.

Results over a small set of models are summarized in Table 2. This table makes use of a measure similar to the Bjøntegaard delta rate (Bjøntegaard, 2001), as was also described in (El Sayeh Khalil et al., 2016). It interpolates the ratepoints within a limited rate range, samples the distortion values and measures the differences in rate at these samples. This way we can find a maximal, minimal and average value of these differences, with a positive difference indicating that a state-of-the-art codec requires more bits for the same quality, and a negative difference indicating that the state-of-the-art codec outperforms our proposed codec. In this case, the limited rate range over which

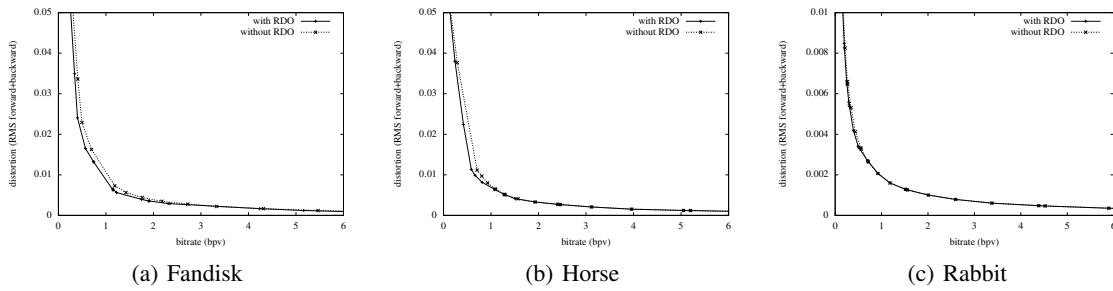


Figure 4: Rate-distortion optimization.




Figure 5: Visual results at specific rates. Figures 5(a), 5(b) and 5(c) show the model at increasing rates. Observe that 3bpv is still in the low bitrate range considering a lossless rate of 33.3bpv for this 18 bit quantized fertility model shown in 5(d).

Table 2: Rate savings w.r.t. the state of the art. To obtain the same quality at rates up to 3bpv, the numbers indicate in bpv (Δ_{avg}) the average rate savings, (Δ_{max}) the largest rate savings, and (Δ_{min}) the smallest rate savings. A positive value means more rate is required than the proposed codec, a negative value means the proposed codec performs worse.

Model	Coder	Δ_{avg}	Δ_{max}	Δ_{min}
teapot (1, 292)	EMD+16	+0.22	+1.00	0.0
beethoven (2, 521)	EMD+16	+0.32	+0.81	0.0
triceratops (2, 832)	EMD+16	+0.07	+0.59	0.0
	Wavemesh	-0.72	+0.43	-1.30
elk (5, 194)	EMD+16	+0.16	+0.60	0.0
	Wavemesh	+0.10	+1.20	-0.28
fanDisk (6, 475)	EMD+16	+0.08	+0.73	0.0
	Wavemesh	+0.25	+0.93	+0.04
	IPR	+0.43	+2.90	-0.03
maxplanck (7, 399)	EMD+16	+0.09	+0.39	0.0
	Wavemesh	-0.25	+0.04	-0.37
venushead (8, 268)	EMD+16	+0.06	+0.73	-0.04
bimba (8, 857)	EMD+16	+0.10	+0.39	0.0
	Wavemesh	+0.28	+1.1	+0.14
horse (19, 851)	EMD+16	+0.10	+0.39	0.0
	Wavemesh	+0.15	+0.84	-0.17
	IPR	+2.40	+7.60	+0.55
screwdriver (65, 538)	EMD+16	+0.01	+0.78	0.0
	Wavemesh	-0.04	+0.01	-2.1
rabbit (67, 039)	EMD+16	+0.011	+0.29	0.0
	Wavemesh	+0.06	+1.3	0.0
	IPR	-0.15	+0.40	-0.80
dino (129, 026)	EMD+16	+0.02	+0.94	0.0
	Wavemesh	-0.08	0.0	-2.50

we measure is taken up to 3bpv for the proposed codec. Furthermore, our rate-distortion optimization

and quality-scalable decoding improves all results of EMD+16. At such low bitrates, relatively high gains of up to 1bpv are obtained. Observe that the minimal rate difference is zero as both the proposed codec and EMD+16 start at the same base mesh, resulting in the same distortion at the same rate. This confirms that our rate-distortion optimized codec never performs worse than EMD+16, as the minimal differences are never negative. The comparison with Wavemesh and IPR shows that in most cases the proposed codec is also more favourable at these low bitrates, except for *dino* where we are on par at best.

7 CONCLUSIONS

We have shown how to achieve quality scalability and optimized rate-distortion performance for the wavelet-based irregular mesh codec *EMD+16* which employs octree-based encoding of both connectivity and geometry information. Gains up to 1bpv are obtained at low bitrates, while at higher bitrates the original codec design is nearly optimal in rate-distortion sense. Furthermore, the results and working implementation serve as a proof-of-concept that an unrestricted storage and transmission of subband bitplanes can be provided using the proposed framework.

ACKNOWLEDGEMENTS

The research activities as described in this paper were funded by Ghent University, iMinds, Flanders Inno-

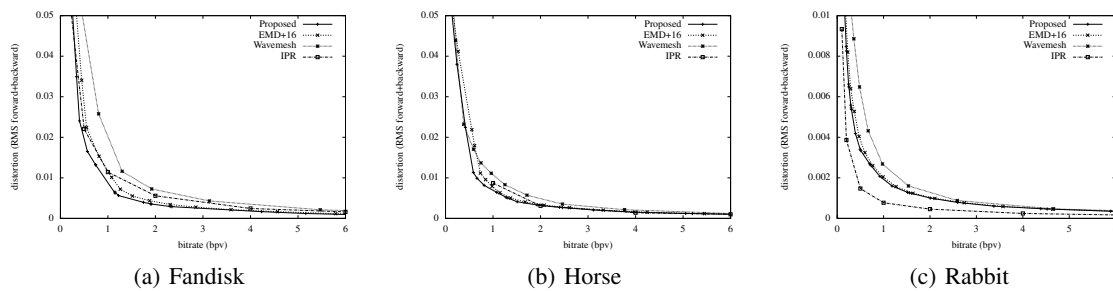


Figure 6: Comparison with the state of the art.

vation & Entrepreneurship (VLAIO), the Fund for Scientific Research-Flanders (FWO-Flanders), and the European Union.

REFERENCES

- Andreopoulos, Y., Munteanu, A., Barbarien, J., van der Schaar, M., Cornelis, J., and Schelkens, P. (2004). In-band motion compensated temporal filtering. *Signal Process. Image*, 19(7):653–673. Special Issue on Sub-band/Wavelet Interframe Video Coding.
- Bjøntegaard, G. (2001). Calculation of average PSNR differences between RD-curves. Technical Report VCEG-M33, ITU-T SG16/Q6, Austin, TX, USA.
- Bonneau, G.-P. (1998). Multiresolution analysis on irregular surface meshes. *IEEE Trans. Vis. Comput. Graphics*, 4(4):365–378.
- Cignoni, P., Rocchini, C., and Scopigno, R. (1998). METRO: measuring error on simplified surfaces. *Comput. Graph. Forum*, 17(2):167–174.
- Denis, L., Satti, S. M., Munteanu, A., Cornelis, J., and Schelkens, P. (2010). Scalable intraband and composite wavelet-based coding of semiregular meshes. *IEEE Trans. Multimedia*, 12(8):773–789.
- El Sayeh Khalil, J., Munteanu, A., Denis, L., Lambert, P., and Van de Walle, R. (2016). Scalable feature-preserving irregular mesh coding. *Comput. Graph. Forum*.
- Hoppe, H. (1996). Progressive meshes. In *Proc. 23rd SIGGRAPH Conf. Computer Graphics*, pages 99–108.
- Khodakovsky, A., Schröder, P., and Sweldens, W. (2000). Progressive geometry compression. In *Proc. 27th SIGGRAPH Conf. Computer Graphics and Interactive Techniques*, pages 271–278.
- Lounsbery, M., DeRose, T. D., and Warren, J. D. (1997). Multiresolution analysis for surfaces of arbitrary topological type. *ACM Trans. Graph.*, 16(1):34–73.
- Munteanu, A., Cornelis, J., Van der Auwera, G., and Cristea, P. (1999). Wavelet-based lossless compression scheme with progressive transmission capability. *Int. J. Imag. Syst. Tech.*, 10(1):76–85.
- Ohm, J. R. (1994). Three-dimensional subband coding with motion compensation. *IEEE Trans. Image Process.*, 3(5):559–571.
- Pajarola, R. and Rossignac, J. R. (2000). Compressed progressive meshes. *IEEE Trans. Vis. Comput. Graphics*, 6(1):79–93.
- Payan, F., Roudet, C., and Sauvage, B. (2015). Semi-regular triangle remeshing: A comprehensive study. *Comput. Graph. Forum*, 34(1):86–102.
- Peng, J. and Kuo, C.-C. J. (2005). Geometry-guided progressive lossless 3D mesh coding with octree (OT) decomposition. In *Proc. 32nd SIGGRAPH Internat. Conf. Computer Graphics and Interactive Techniques*, pages 609–616.
- Said, A. and Pearlman, W. A. (1996). A new, fast, and efficient image codec based on set partitioning in hierarchical trees. *IEEE Trans. Circuits Syst. Video Technol.*, 6(3):243–250.
- Sweldens, W. (1998). The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29(2):511–546.
- Taubman, D. and Zakhor, A. (1994). Multirate 3-D subband coding of video. *IEEE Trans. Image Process.*, 3(5):572–588.
- Taubman, D. S. and Marcellin, M. W. (2001). *JPEG 2000: Image Compression Fundamentals, Standards and Practice*. Kluwer Academic Publishers, Norwell, MA, USA.
- Touma, C. and Gotsman, C. (1998). Triangle mesh compression. In *Proc. Graphics Interface Conf.*, pages 26–34.
- Valette, S., Chaine, R., and Prost, R. (2009). Progressive lossless mesh compression via incremental parametric refinement. *Comput. Graph. Forum*, 28(5):1301–1310.
- Valette, S., Kim, Y.-S., Jung, H.-Y., Magnin, I., and Prost, R. (1999). A multiresolution wavelet scheme for irregularly subdivided 3d triangular mesh. In *Proc. Internat. Conf. Image Processing*, pages 171–174.
- Valette, S. and Prost, R. (2004a). Wavelet-based multiresolution analysis of irregular surface meshes. *IEEE Trans. Vis. Comput. Graphics*, 10(2):113–122.
- Valette, S. and Prost, R. (2004b). Wavelet-based progressive compression scheme for triangle meshes: Wavemesh. *IEEE Trans. Vis. Comput. Graphics*, 10(2):123–129.