The Visual Computer manuscript No.
(will be inserted by the editor)

Jianhua Wu · Leif Kobbelt

Efficient Spectral Watermarking of Large Meshes with Orthogonal Basis Functions

Abstract Allowing for copyright protection and ownership assertion, digital watermarking techniques, which have been successfully applied for classical media types like audio, images and videos, have recently been adapted for the newly emerged multimedia data type of 3D geometry models. In particular, the widely used spread-spectrum methods can be generalized for 3D datasets by transforming the original model to a frequency domain and perturbing the coefficients of the most dominant basis functions. Previous approaches employing this kind of spectral watermarking are mainly based on multiresolution mesh analysis, wavelet domain transformation or spectral mesh analysis. Though they already exhibit good resistance to many types of real-world attacks, they are often far too slow to cope with very large meshes due to their complicated numerical computations. In this paper, we present a novel spectral watermarking scheme using new orthogonal basis functions based on radial basis functions. With our proposed fast basis function orthogonalization, while observing similar persistence with respect to various attacks as other related approaches, our scheme runs faster by two orders of magnitude and thus can efficiently watermark very large models.

Keywords Digital Watermarking · Large Meshes Watermarking · Radial Basis Functions · Spectral Decomposition

1 Introduction

Watermarking is an established way to provide copyright protection and ownership assertion in the area of steganography. Most availible digital watermarking techniques have been focusing mainly on classical media data types like audio, images and videos because of the dominancy of these data distributed on the Internet [15,8].

Computer Graphics Group
RWTH Aachen University, Germany
E-mail: {wu, kobbelt}@cs.rwth-aachen.de

Fig. 1 The Iphigenie model (left, 1.01M vertices, 2.02M triangles) is watermarked in 86 seconds (middle) using 100 orthogonal basis functions. Although without perceptible visual differences, these two meshes have a maximum Hausdorff distance of 1.64% of the major bounding box diagonal length (right, distances with color coding - blue minimum and red maximum).

Due to their regularly parametrized functional representations, most watermarking schemes are based on *spread-spectrum* methods with signal processing, i.e., the media data have to be transformed into a spectral domain, then the coefficients corresponding to the most perceptually salient basis functions will be modulated with watermarks to achieve *robustness* against possible attacks.

Extending the above spectral watermarking methods to the newly emerged multimedia data type of 3D geometry models is difficult mainly because of the lack of basic 3D signal processing tools like filtering, regular parametrization and frequency analysis. On the other hand, with a drastically increasing availability of 3D datasets and practical geometric applications in recent years, the need for efficient watermarking schemes for 3D models becomes more eminent.

Previous robust 3D spectral watermarking approaches have been adapted and based mainly on multiresolution mesh analysis, wavelet domain transformation or spectral mesh analysis. Although they already exhibit good resistance and robustness to many types of real-world attacks, they are often far too slow to cope with nowadays large meshes due to the involved complicated numerical computations.

In this paper, we present a novel imperceptible spectral watermarking scheme to support ownership claims on triangle meshes of given 3D shapes (cf. Fig. 1). To span the spectral domain for watermarking, it uses a new set of orthogonal basis functions derived from *radial basis functions* (RBFs), which can lead to optimal concentration of the shape information to just a few (low-frequency) modes (cf. Fig. 4). Concluding from extensive tests on different models, our watermarking scheme exhibits almost the same watermarking quality and robustness against various real world attacks as other related spectral approaches. On the other side, by utilizing a fast basis function orthogonalization algorithm, our watermarking scheme runs much faster by two orders of magnitude, hence can process and watermark very large models more efficiently.

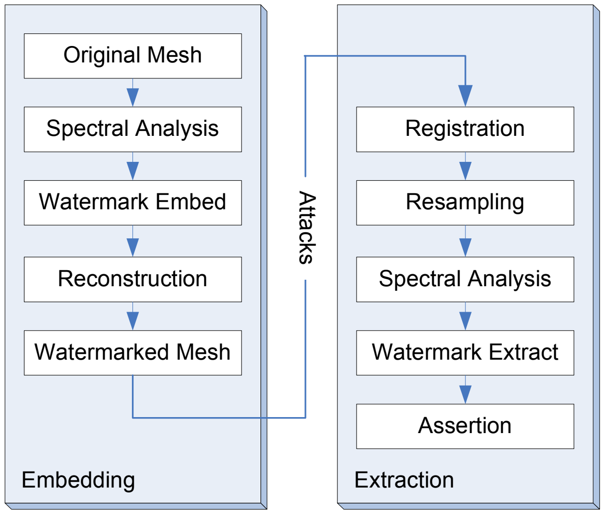

Fig. 2 A typical watermarking scenario.

1.1 Related Work

Most previous works to watermark 3D models have been trying to mimic the common spectral approaches in some alternative ways, though early works on 3D watermarking even did not utilize the spectral idea. Watermarks were embedded into 3D meshes by directly modifying either the geometry, the vertex positions, the topology, the vertex connectivity [18,19,12], or the surface normals [3]. Simple enough, these kind of methods usually can not provide enough robustness with respect to many different types of ordinary attacks.

Kanai et al. [13] decomposed the target mesh into a spectral domain by applying the lazy *wavelet transform* proposed by Lounsbery et al. [16]. Wavelet coefficients were then modified to embed watermarks. Extending this work, a blind watermarking algorithm was presented more recently [31] that can ignore the original mesh information on the detector side. One constraint to these methods is that the input mesh is limited to have a prerequisite semi-regular subdivision connectivity.

Multiresolution analysis is another way to construct the spectral-like domains. While Praun et al. [23] used standard mesh simplification [10] to construct multiresolution hierarchies, Yin et al. [33] adopted the scheme in [11] to perform the multiresolution decomposition. Watermark information can be embedded into some spatial kernels of the low-frequency component of the shape corresponding to the low-resolution representation in the geometry hierarchy.

Recent spectral domain watermarking algorithms employed the *spectral mesh analysis* first proposed by Karni and Gotsman [14]. Eigenvectors of the Laplace matrix to the input mesh, the Laplace basis functions, can span an ideal spectral space for robust watermarking, i.e., leading coefficients corresponding to smallest eigenvalues can be modulated with watermarks [20]. Later this idea was generalized to watermark the topology-free point-based geometry in [7] where k-nearest neighbors have to be constructed to compose substitute Laplace matrices.

Despite the variety of available mesh watermarking algorithms, to our knowledge, none of them has reported to be able to watermark meshes with more than 10^4 vertices mainly because of their involved complicated numerical computations. We will present a *fast yet robust* spectral watermarking algorithm based on the orthogonalization of a small set of radial basis functions that can efficiently handle large meshes even with more than 10^6 vertices (cf. Fig. 1).

Except for our geometry dependant basis functions derived from RBFs, we note that other bases can have similar functionality as ours such as the harmonic functions computed with mesh Laplacian which were recently used for surface deformation and shape approximation [25,34]. Other than the necessary post orthogonalization step, they still have to solve sparse linear systems of Laplace equations, which makes them less efficient as ours. In addition, other than in digital watermarking, the fact that altering the low frequency components of a shape remains nearly invisible to the human eyes has also been observed in mesh compression [26].

2 Overview

Following the most successful spread-spectrum watermarking idea, we also watermark 3D meshes in the spectral domain. The whole watermarking scenario which is typically composed of the watermark embedding and the watermark extraction steps, is illustrated in Figure 2 similar to [20,7].

The major difference of our scheme compared to previous work is that we use a new set of geometry dependant orthogonal basis functions derived from radial basis functions (cf. Section 3) to span the spectral space rather than using Laplace basis functions which emerge from the Laplace matrix that depends only on the mesh connectivity. We will present a fast algorithm to generate these orthogonal basis functions. Compared with the time-consuming eigensolvers for Laplace matrices, our method runs faster by two orders of magnitude and thus can efficiently watermark very large meshes as they are common today. Having the new orthogonal basis functions in hand, the remaining watermarking procedures are quite similar to other spectral watermarking approaches.

The watermark embedding phase (cf. Section 4) first computes a small set of our new orthogonal basis functions. Then the geometry of the original mesh is projected to these basis functions spanning the spectral domain to acquire a set of corresponding spectral coefficients. Watermarks will be encoded into the leading coefficients which can be used later to reconstruct the watermarked mesh together with the unmodified coefficients. Finally this watermarked mesh will be publized with licenses if necessary and possibly receive some attacks.

To assert ownership, the watermark extraction step (cf. Section 5) has to be performed. The possibly attacked test mesh is first aligned by a registration routine and resampled by projections according to the original watermarked mesh. Then this modified attacked mesh will be transformed to the same spectral domain as in the embedding phase. Watermarks can be extracted by comparing the current and original coefficients. The correlation between the extracted and original watermarks will also be found to draw the final ownership assertion.

3 New Orthogonal Basis Functions

The basic idea of spectral watermarking is to represent the geometric information of a mesh with respect to a special basis such that most of the information is captured in just a few coefficients. These coefficients are then used to embed the watermark.

More precisely, let $\mathcal{M}$ be a given mesh with vertex positions $\mathbf{p}_1 \ldots \mathbf{p}_n \in I\!R^3$. Then we want to find a basis $\mathbf{B}_1 \ldots \mathbf{B}_n \in I\!R^n$ such that

$$(\mathbf{p}_1 \ldots \mathbf{p}_n)^T \;=\; \sum_{j=1}^{n} \mathbf{B}_j \, \mathbf{c}_j^T$$

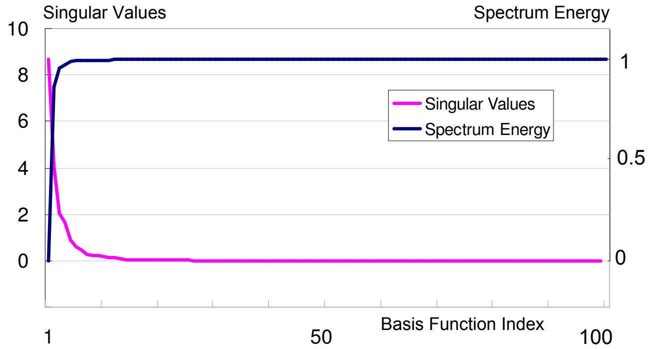

Fig. 3 Singular values $\{w_i\}$ corresponding to the 100 orthogonal basis functions of the Rocker arm mesh (Fig. 4, middle right). The respective spectrum energy $E_i \in [0,1]$ is computed as $E_i = \sqrt{(\sum_{j=1}^{i} w_j^2)/(\sum_{j=1}^{100} w_j^2)}$. Note that most of the energy is concentrated in the leading part.

with coefficients $\mathbf{c}_j \in I\!R^3$. The particular basis should have the property that the approximation error

$$E_k \;=\; \|(\mathbf{p}_1 \ldots \mathbf{p}_n)^T \;-\; \sum_{j=1}^{k} \mathbf{B}_j \, \mathbf{c}_j^T\|$$

decreases as quickly as possible. A very bad example is the canonical basis $\mathbf{B}_j = (0 \ldots 1 \ldots 0)^T$ for which the error E_k decreases only linearly in k. A much better example is the set of basis functions which emerge as eigenvectors from the topological Laplace matrix defined through the connectivity of the given mesh. This basis has been used in [14] where it has been shown empirically that the approximation error E_k decreases so fast that only a few hundreds to thousands of basis coefficients are sufficient to encode even fairly complex shapes [2].

Another important property of the basis is orthogonality since in this case the coefficients $\mathbf{c}_j$ can simply be computed by a dot product

$$\mathbf{c}_j \;=\; (\mathbf{p}_1 \ldots \mathbf{p}_n) \, \mathbf{B}_j.$$

Both, the canonical as well as the Laplace basis are orthogonal. However, both bases have the important drawback that their basis function definition completely ignores the geometry of the input shape. Although the Laplace basis does depend on the mesh connectivity, the geometric shape of the input mesh only plays a minor role under the assumption that the mesh quality (i.e., the shape of the triangles) establishes a loose correlation between geometric and topological distance.

In this paper we are constructing a particular orthogonal basis which, on the one hand is optimized for the shape of the input geometry and on the other hand concentrates as much geometric information as possible in as few coefficients as possible. We start by defining a pre-basis $\{\mathbf{B}_j\}$ and then augment it later.

We are using radial basis functions (RBFs) [29,6,17,30,21,22,28] to capture the mesh geometry information. Let $\mathbf{q}_1 \ldots \mathbf{q}_k$ be a set of 3D positions scattered in the

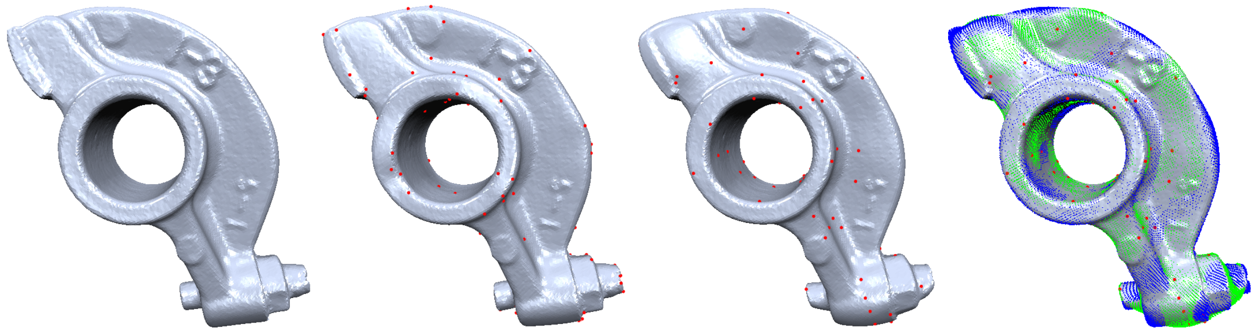

Fig. 4 The Rocker arm mesh (left-most, $40K$ vertices) approximated with 100 orthogonal basis functions based on RBFs using decimated centers (red dots, middle left) and the random uniform centers (middle right). The right-most image shows the vertex-to-vertex deviation vectors from the random center RBF approximation to the original model where vectors pointing inside the object are flipped and rendered as green lines. Notice the good approximation quality of only 100 basis functions and the small differences between two different center placement strategies as well.

vicinity of the mesh surface. There are many possibilities to define these positions, e.g., by selecting the remaining points in $\mathbf{p}_i$ after the mesh decimation with halfedge collapses, or even randomly selecting a uniform subset of the vertices $\mathbf{p}_i$. (See Figure 4 for a comparison, if not otherwise stated, we will always use random uniform selection for efficiency reasons). The $\mathbf{q}_j$ are used as the centers of a set of radial basis functions

$$\phi_j(\mathbf{p}) \;=\; \phi(\|\mathbf{p}-\mathbf{q}_j\|)$$

where we choose $\phi(\cdot)$ to be a monotonically decreasing function with compact support, e.g. $\phi(r) = (1-r)^4_+(4r+1)$ like in [32,22] with $r = \|\mathbf{p}-\mathbf{q}_i\|/\sigma$ and σ its support size (usually half the length of main bounding box diagonal to prevent singularity of the later composed matrix $\mathbf{B}$).

If we evaluate these radial basis functions at all vertex positions $\mathbf{p}_i$, we obtain our set of discrete pre-basis functions

$$\mathbf{B}_j \;=\; [\phi_j(\mathbf{p}_1)\ldots\phi_j(\mathbf{p}_n)]^T.$$

Due to the compact support of $\phi(\cdot)$, the discrete pre-basis functions are very likely to be linearly independent in $I\!R^n$. Moreover if the radial basis functions $\phi_j(\cdot)$ have a sufficient overlap then it turns out that a relatively small number k of pre-basis functions are already sufficient to represent the geometric information of the input mesh fairly accurately (cf. Fig. 4), i.e.,

$$(\mathbf{p}_1\ldots\mathbf{p}_n)^T \;\approx\; \sum_{j=1}^{k}\mathbf{B}_j\,\mathbf{c}_j^T.$$

Moreover, if we choose the distribution of the centers $\mathbf{q}_j$ fairly uniform (and we assume that the vertices $\mathbf{p}_i$ are also distributed uniformly over the mesh surface) then it turns out that the coefficients $\mathbf{c}_j$'s magnitudes do not differ too much.

Notice that the choice of the actual number k of pre-basis functions has influence only on the computation time of the subsequent orthogonalization step and not on the applicability of the approach in general.

Remember that our goal is to construct an orthogonal basis which concentrates most of the geometric information in just a few coefficients. In order to obtain this basis we compute a singular value decomposition (SVD [24]) of the matrix $\mathbf{B}\in R^{n\times k}$ which has the pre-basis functions $\mathbf{B}_j$ as its columns. With this decomposition we find

$$\begin{aligned}(\mathbf{p}_1\ldots\mathbf{p}_n)^T &\approx \mathbf{B}\;(\mathbf{c}_1\ldots\mathbf{c}_k)^T\\ &= \mathbf{U}\,\mathbf{W}\,\mathbf{V}^T\,(\mathbf{c}_1\ldots\mathbf{c}_k)^T\\ &= \mathbf{U}\,(\mathbf{c}'_1\ldots\mathbf{c}'_k)^T.\end{aligned}$$

Hence the orthogonal columns of $\mathbf{U}$ form a new set of discrete basis functions. The corresponding coefficients $\mathbf{c}'_j$ are obtained by first multiplying the initial coefficients $\mathbf{c}_j$ by the orthogonal matrix $\mathbf{V}^T$ and then multiplying the j-th coefficient with the j-th singular value of $\mathbf{B}$.

With increasing overlap, i.e., with increasing radius of the radial basis function $\phi_j(\cdot)$, we observe a stronger and stronger decay of the singular values of $\mathbf{B}$ and hence a more and more pronounced concentration of the geometric information to just a few leading coefficients (cf. Fig. 3). In fact, suppressing the later coefficients $\mathbf{c}'_l\ldots\mathbf{c}'_k$ (which have been multiplied by the smaller singular values) causes only an additional squared approximation error of

$$E'^2_l \;=\; \sum_{j=l}^{k}\|\mathbf{c}'_j\|^2$$

due to the orthogonality of the columns of $\mathbf{U}$.

While the concentration effect observed with our orthogonalized basis is similar to the effect observed with the eigenbasis of the Laplace operator, the advantage of our approach is that its computation is significantly faster. In fact, computing the eigenbasis of a very large sparse matrix is a numerically challenging task and usually large meshes are split into smaller patches and processed separately in order to cope with this problem [14, 20].

In our case, however, we do not have to analyze the large $(n\times k)$ matrix $\mathbf{B}$ directly. Instead it is sufficient to decompose the much smaller symmetric $(k \times k)$ matrix $\mathbf{B}^T\mathbf{B}$ into

$$\mathbf{B}^T\mathbf{B} \;=\; \mathbf{V}\,\mathbf{W}\,\mathbf{U}^T\,\mathbf{U}\,\mathbf{W}\,\mathbf{V}^T \;=\; \mathbf{V}\,\mathbf{W}^2\,\mathbf{V}^T.$$

With this decomposition we can easily find the orthogonalized basis by

$$\mathbf{U} \;=\; \mathbf{B}\,\mathbf{V}\,\mathbf{W}^{-1}.$$

Since k is usually much smaller than n, the time spent to compute the SVD is negligible and the computation time is dominated by the multiplication of $\mathbf{B}$ and $\mathbf{V}$.

4 Watermark Embedding

Similar to other spectral watermarking approaches, our watermarking scheme embeds the digital watermarks by modifying the low-frequency components of a given shape in the spectral domain. As we have discussed in Section 3, the new orthogonal basis functions $\{\mathbf{B}_i\}$ derived via SVD will be used to decompose the input mesh into a spectral representation and the coefficients of the leading part of the spectrum which are more robust against attacks, then can be modulated. The watermarked mesh is later produced with an inverse transform using the same basis functions and is ready to be distributed.

More specifically, given the original (large) input mesh $\mathcal{M}_o$ with n vertices, a set of k orthonormal basis functions $\{\mathbf{B}_i\}$ is first computed to span a spectral domain. Then *spectral analysis* is performed by projecting the mesh geometry (all three x, y, z components) onto each basis function $\mathbf{B}_i$ to produce $3k$ mesh spectral coefficients $\{\alpha_i^{(x)}\}, \{\alpha_i^{(y)}\}$ and $\{\alpha_i^{(z)}\}$, i.e.,

$$\alpha_i^{(d)} \;=\; \sum_{j=1}^{n} \mathbf{p}_j^{(d)}\mathbf{B}_{i,j}, \quad i = 1...k, \;\; d \in \{x, y, z\}, \qquad (1)$$

where $\mathbf{B}_{i,j}$ denotes the j-th entry of the i-th basis function. Based on these, a spectral mesh representation $\mathcal{M}'_o$ (cf. Fig. 4 middle) can in turn be assembled to approximate the original mesh $\mathcal{M}_o$ with an inverse transform:

$$\mathcal{M}'_o \;=\; \sum_{i=1}^{k} \alpha_i \mathbf{B}_i.$$

The approximation quality of $\mathcal{M}'_o$ is already quite good. But as we only use a small number of basis functions $(k \ll n)$, we still observe small vertex-to-vertex differences between the approximation $\mathcal{M}'_o$ and original $\mathcal{M}_o$ (cf. Fig. 4 right). These differences Δ will be recorded separately to help the later reconstruction, i.e.,

$$\Delta \;=\; \mathcal{M}_o - \mathcal{M}'_o \;=\; \mathcal{M}_o - \sum_{i=1}^{k} \alpha_i \mathbf{B}_i.$$

The watermarks are a binary bit string $\{b_i\}$ with length $3m$ and $m < k$. To *embed watermarks* into the spectral representation, we first convert $\{b_i\}$ to a sign string $\{b'_i\}$ with $b'_i = -1$ for $b_i = 0$ and $b'_i = 1$ for $b_i = 1$. Then the first $3m$ spectral coefficients are modulated as

$$\beta_i^{(d)} \;=\; \alpha_i^{(d)} \cdot (1 + b'_{3i+d} \cdot \rho), \;\; i = 1...m, \;\; d \in \{x, y, z\},$$

where ρ is the *watermarking amplitude*.

Finally, these modulated coefficients $\{\beta_i\}$, together with the unchanged ones $\{\alpha_i\}$ and the differences Δ, are composed to *reconstruct* the output watermarked mesh $\mathcal{M}_w$,

$$\mathcal{M}_w \;=\; \sum_{i=1}^{m} \beta_i \mathbf{B}_i \;+\; \sum_{i=m+1}^{k} \alpha_i \mathbf{B}_i \;+\; \Delta.$$

5 Watermark Extraction

The watermarked mesh $\mathcal{M}_w$ will be distributed with licenses if necessary and will possibly receive some types of real world attacks. To assert ownership of this attacked test mesh $\mathcal{M}_t$, previous embedded watermarks have to be extracted as we will discuss in the following.

To undo a possible similarity transform or translation attacks, the attacked mesh $\mathcal{M}_t$ first has to be aligned with the watermarked mesh $\mathcal{M}_w$. We have used a typical *registration* toolbox developed in [1] to compute an affine map for the final alignment. In principle it needs user intervention to define three matching point pairs to compute initial absolute orientations followed by an automatic iterative closest point (ICP) algorithm [4] to improve the initial registration till a local error minimum is reached.

After registration, a *resampling* phase is usually necessary to deal with those attacks that may modify the mesh topology like simplification or remeshing. The goal is to map the original topology of $\mathcal{M}_w$ to $\mathcal{M}_t$ as our basis functions $\{\mathbf{B}_i\}$ are defined in the vertex indices order. Only when the two meshes have the same vertex order, the comparisons between their spectral coefficients to extract the watermarks will be reasonable. Because the registration has already minimized the distances between $\mathcal{M}_t$ and $\mathcal{M}_w$, we can use a simple nearest point search strategy to obtain the resampled test mesh $\mathcal{M}'_t$: the topology of $\mathcal{M}'_t$ is the same as $\mathcal{M}_w$ and each vertex $\mathbf{v}'_i$ of $\mathcal{M}'_t$ is fixed as the nearest point of the corresponding vertex $\mathbf{p}_i$ of $\mathcal{M}_w$ on the mesh surface $\mathcal{M}_t$. Note that like [23] this resampling step also marks vertices as cropped if the nearest distances are larger than a user-defined threshold to account for possible cropping attacks.

Having the registered and resampled test mesh $\mathcal{M}'_t$, we perform a similar *spectral analysis* to it as Equ. (1) to get another set of $3m$ spectral coefficients $\{\gamma_i^{(x)}\}, \{\gamma_i^{(y)}\}$ and $\{\gamma_i^{(z)}\}$. Then for each x, y or z component, watermarks are *extracted* as follows,

$$c_i^{(d)} \;=\; \begin{cases} 1, & \text{if } \gamma_i^{(d)} > \alpha_i^{(d)} + \varepsilon, \\ 0, & \text{if } \gamma_i^{(d)} < \alpha_i^{(d)} - \varepsilon, \\ N/A, & \text{otherwise}, \end{cases}$$

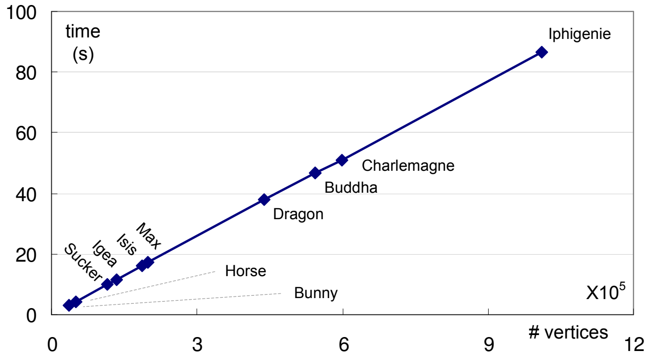

Fig. 5 Computation times for embedding the watermarks into various models.

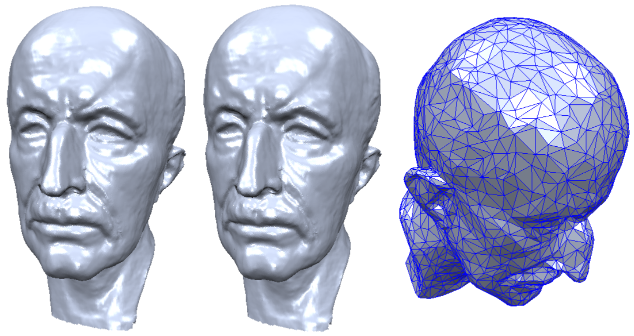

Fig. 6 The Max model (left, $200K$ vertices), watermarked mesh (middle) and the test mesh under both similarity transform and simplification attacks (right, 1% of original size). The extracted correlation is $R = 1$.

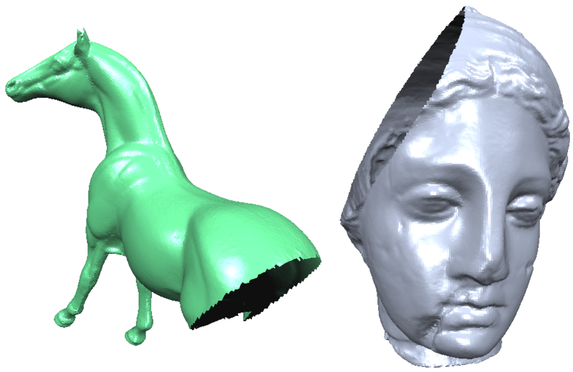

Fig. 7 The cropped Horse model (left, $51K$ vertices) and the cropped Igea model (right, $134K$ vertices). The correlation is $R = 1$ in both cases.

where ε is the detection sensitivity and is set to 0.1ρ.

Since various attacks might disturb some of the embedded watermarks, bitwise comparisons have to be performed to create the correlation between embedded watermarks $\{b_i\}$ and extracted watermarks $\{c_i^{(d)}\}$, i.e., the correlation will be

$$R = \frac{1}{3m} \cdot \sum_{i=1}^{m} \sum_{d=x}^{z} (\, c_i^{(d)} == b_{(3i+d)}\,). \qquad (2)$$

The final *ownership assertion* can then be made as in [7]. If the correlation R is larger than a specified threshold (e.g. 0.75), we affirm that the attacked test mesh $\mathcal{M}_t$ contains the originally embedded watermarks.

6 Results

We have tested our watermarking algorithm on various typical 3D mesh datasets. The overall performance including the timing measurement and memory estimation is first to be discussed. The watermarking robustness of our scheme will be verified with a lot of diverse real world attacks. A comparison between our orthogonal basis functions based on RBFs and the Laplace basis functions will also be conducted. All experiments shown here have been done on a commodity Linux PC with a 3.2GHz P-IV CPU and 2GB main memory. Some important parameters are usually set as following (if not otherwise specified): basis function number $k = 100$, watermark length $3m = 24$ and watermarking amplitude $\rho = 0.01$.

6.1 Overall Performance

Computation times of our watermark embedding process are measured and summarized in Figure 5 as a function of input model sizes. Other recent spectral watermarking schemes like [20, 7] need much longer running times than ours, as they have to solve large eigensystems which will be impractical for meshes with more than 10^4 vertices. Even when taking the differences of computing hardware into account, we still find that (e.g. for the bunny model reported in both papers) our algorithm runs more than 100 times faster, which are two orders of magnitude.

Actually, as our watermarking method runs so fast, experiments show that the performance bottleneck arises from the memory consumption. For example, our non-optimized watermarking proto-system needs more than 1GB Linux process memory to deal with the Iphigenie mesh and the basis functions themselves require about 400MB space already. However, this is just the reality that all spectral-based approaches have to face as basis functions always need to be computed for both watermark embedding and extraction phases.

In addition, as our watermarking scheme only modifies the “low-frequency” component of a given shape, the visual differences between the original mesh and the watermarked mesh are almost imperceptible (cf. Fig. 1, 6 and 10). And this will make the watermarking scheme more robust to malicious attacks.

6.2 Robustness Against Attacks

We test our watermarking method with different mesh models to verify the robustness of embedded watermarks under various types of real world attacks. The watermarks will be generated randomly many times for the same testing object and the extracted correlation R (cf. Equa. 2) will be the mean value. If the correlation is larger than 0.75, the ownership will be claimed.

Similarity Transforms can be handled by the first registration process (cf. Fig. 6). As our registration results are precise enough, the similarity transform attack will have very few influences on the correlation R;
Cropping is dissolved by the resampling process to mark the cropped vertices and neglect them in the basis functions when performing the spectral analysis. The correlation R is always 1 (cf. Fig. 7);
Simplification is simulated with a standard QEM-based greedy decimation scheme [9]. The resampling process will equalize the mesh topology for watermark extraction (cf. Fig. 8). Note that even under extreme simplification, ownership still can be successfully asserted (cf. Fig. 6 and 8);
Additive noise will randomly modify the vertex positions in the normal directions. Figure 9 shows that our watermarking scheme is robust against strong noises;
Remeshing attack tries to improve the shape quality of mesh primitives where the initial topology will be greatly destroyed (cf. Fig. 10). Thus the resampling process is also necessary for watermark extraction.
Smoothing will apply certain iterations of Laplace operators [27] to the watermarked mesh. Figure 11 compares the robustness effects of different iteration numbers to the extracted correlations.

The above testing scenarios show that we can correctly assert the ownership of the watermarked mesh. We also test our watermarking scheme with *false-positive* attacks, e.g., to perform the same watermark extraction step to a mesh without actrually having embedded the designated watermarks or with other watermarks embedded. The extracted correlations are always below 0.5. Regarding the specified threshold, our watermarking algorithm will not incorrectly assert that a model is watermarked when it is not.

In summary, without perceptible differences between the original and watermarked meshes, our watermarking scheme is robust against a lot of different real world attacks, even the combined attacks as well (cf. Fig. 6), thus can lead to accurate ownership assertion and copyright protection.

6.3 Comparing to Laplace Basis Functions

Figure 11 compares our new orthogonal basis functions derived from RBFs to the Laplace basis functions (LBFs), the (leading) eigenvectors of the Laplace matrix of the input mesh [14] adopted by recent spectral watermarking schemes [20, 7]. It is interesting to see that our basis functions can capture more shape details than LBFs with the same number of basis functions. This is because ours are geometry-aware derived functions while LBFs are merely derived from the mesh topology. On the other hand, it is still easy to find that our watermarking scheme can have almost the same robustness against attacks and watermarking quality as previous spectral methods based on LBFs.

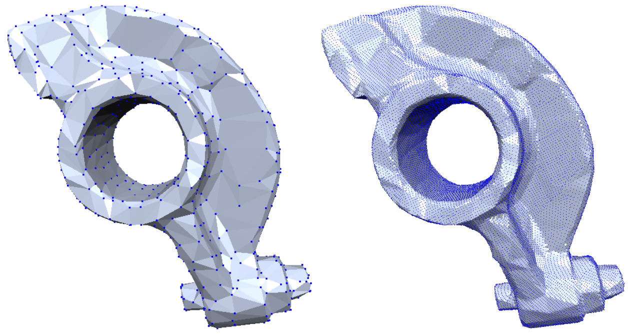

Fig. 8 The Rocker arm mesh (in Fig. 4) under simplification attack (left, 2% of original size) and the resampled mesh (right). The correlation is $R = 0.96$.

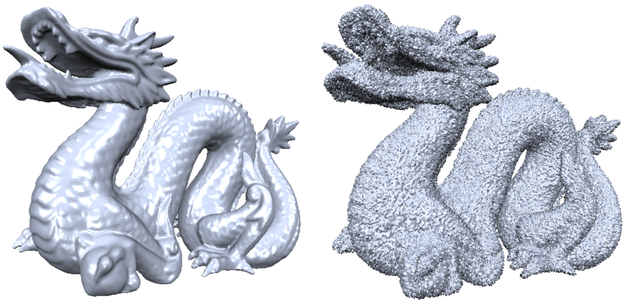

Fig. 9 The watermarked Dragon model (left, $438K$ vertices) and attacked by additive noise(right) with maximum deviation of 1% to main bounding box diagonal length. The correlation is $R = 1$.

7 Conclusions

In this paper, we have presented a novel spectral watermarking scheme using a new set of orthogonal basis functions derived from radial basis functions. While observing similar watermarking quality and robustness against attacks as other related approaches, our method runs much faster by two orders of magnitude thus leading to efficient watermarking of very large models.

The proposed watermarking algorithm is fast, but it is also very flexible. In fact, we also find that our watermarking technique using RBFs can be easily applied to point-sampled geometry by slightly changing the extraction process like adding registration and resampling tools for point sets. Since RBFs are used, the method can be extended for watermarking higher dimensional data, e.g. points with colors, meshes with texture coordinates, and even animated meshes.

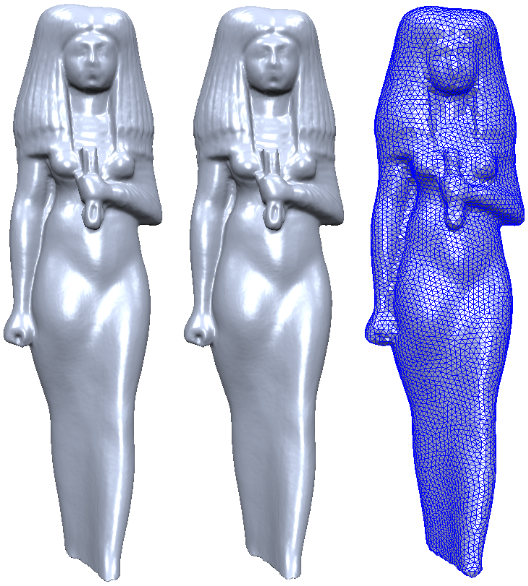

Fig. 10 The Isis mesh (left, $188K$ vert.), watermarked mesh (middle) and remeshed mesh (right, $8K$ vert.). The correlation is $R = 1$.

Future work can be extended in many directions: Relationships between our new orthogonal basis functions and LBFs have to be more clearly understood. Also more types of attacks like free-form constrained modeling [5] need to be handled. Finally, out-of-core watermarking implementations are necessary as massive models are usually more precious and have higher demands on the protection.

Acknowledgements We would like to thank the anonymous reviewers for their insightful comments and Stephan Bischoff for proofreading the paper. The models in the paper are partially courtesy of the computer graphics group at the Stanford University and the Cyberware website.

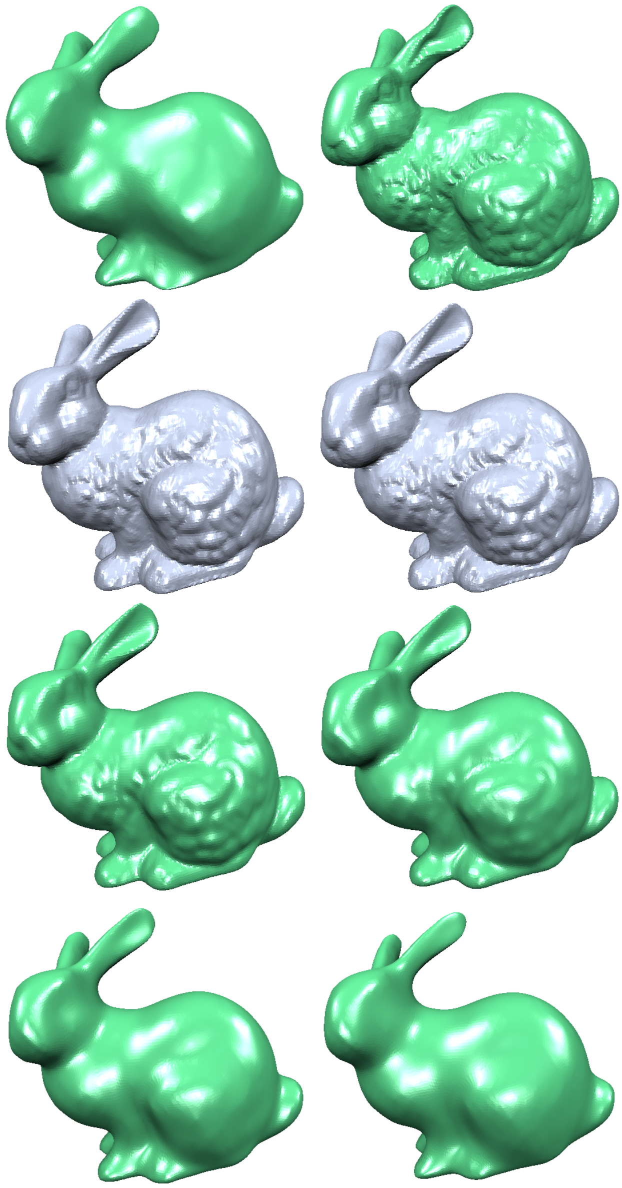

Fig. 11 The Bunny mesh ($35K$ vertices) approximated in first row with 100 LBFs (left) and RBFs (right) and watermarked in second row with corresponding basis functions. The two bottom rows show smoothing attacks on RBF watermarked mesh with $5, 15, 50$ and 100 iterations of Laplace operators respectively from top to bottom and left to right. The corresponding extracted correlations are $1, 0.96, 0.83$ and 0.8 respectively. Under same attacks, the LBF watermarking correlations are $1, 0.96, 0.83$ and 0.83.

References

1. Alekseyev, S.: Volumetric surface reconstruction with level set methods. In: Diploma Thesis. RWTH Aachen University, Germany (2003)
2. Ben-Chen, M., Gotsman, C.: On the optimality of spectral compression of mesh data. ACM Transactions on Graphics **24**(1), 60–80 (2005)
3. Benedens, O.: Geometry-based watermarking of 3d models. IEEE Computer Graphics and Applications **19**(1), 46–55 (1999)
4. Besl, P., McKay, J.: A method for registration of 3-d shapes. IEEE Trans. on Pattern Recognition and Machine Intelligence **14**(2), 239–255 (1992)
5. Botsch, M., Kobbelt, L.: An intuitive framework for real-time freeform modeling. ACM Transactions on Graphics. Special issue for SIGGRAPH conference **23**(3), 630–634 (2004)
6. Carr, J.C., Beatson, R.K., Cherrie, J.B., Mitchell, T.J., Fright, W.R., MacCallum, B.C., Evans, T.R.: Reconstruction and representation of 3d ojbects with radial basis functions. In: Proceedings ACM SIGGRAPH 2001, pp. 67–76 (2001)
7. Cotting, D., Weyrich, T., Pauly, M., Gross, M.: Robust watermarking of point-sampled geometry. In: Proc. of The International Conference on Shape Modeling and Applications 2004, pp. 233–242 (2004)
8. Cox, I.J., Miller, M.L., Bloom, J.A.: Digital Watermarking. Morgan Kaufmann (2001)
9. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: Proceedings of ACM SIG-

GRAPH 1997, pp. 209–216 (1997)
10. Gotsman, C., Gumhold, S., Kobbelt, L.: Simplification and compression of 3d-meshes. In: Tutorials on Multiresolution in Geometric Modeling. Springer (2002)
11. Guskov, I., Sweldens, W., Schröder, P.: Multiresolution signal processing for meshes. In: Proceedings ACM SIGGRAPH 1999, pp. 325–334 (1999)
12. Harte, T., Bors, A.: Watermarking 3d models. In: International Conference on Image Processing, pp. 661–664 (2002)
13. Kanai, S., Date, H., Kishinami, T.: Digital watermarking of 3d polygons using multiresolution wavelet decomposition. In: Proc. Sixth IFIP WG 5.2 GEO-6, pp. 296–307 (1998)
14. Karni, Z., Gotsman, C.: Spectral compression of mesh geometry. In: Proceedings ACM SIGGRAPH 2000, pp. 279–286 (2000)
15. Katzenbeisser, S., Petitcolas, F.A.P.: Information Hiding Techniques for Steganography and Digital Watermarking. Artech House Books (2000)
16. Lounsbery, M., DeRose, T.D., Warren, J.: Multiresolution analysis for surfaces of arbitrary topological type. ACM Transactions on Graphics **16**(1), 34–73 (1997)
17. Morse, B.S., Yoo, T.S., Chen, D.T., Rheingans, P., Subramanian, K.R.: Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions. In: Proc. of Shape Modeling and Applications 2001, pp. 89–98 (2001)
18. Ohbuchi, R., Masuda, H., Aono, M.: Watermarking three-dimensional polygonal meshes. In: Proc. ACM Multimedia '97, pp. 261–272 (1997)
19. Ohbuchi, R., Masuda, H., Aono, M.: Watermarking three-dimensional polygonal models through geometric and topological modifications. IEEE Journal on Selected Areas in Communications **16**(4), 551–559 (1998)
20. Ohbuchi, R., Mukaiyama, A., Takahashi, S.: A frequency-domain approach to watermarking 3d shapes. Computer Graphics Forum **21**(3), 373–382 (2002). (Proc. Eurographics 2002)
21. Ohtake, Y., Belyaev, A., Seidel, H.P.: A multiscale approach to 3d scattered data interpolation with compactly supported basis functions. In: Proc. of the International Conference on Shape Modeling and Applications, pp. 153–161 (2003)
22. Ohtake, Y., Belyaev, A., Seidel, H.P.: 3d scattered data approximation with adaptive compactly supported radial basis functions. In: Proc. of the International Conference on Shape Modeling and Applications, pp. 31–39 (2004)
23. Praun, E., Hoppe, H., Finkelstein, A.: Robust mesh watermarking. In: Proceedings ACM SIGGRAPH 1999, pp. 49–56 (1999)
24. Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes in C: The Art of Scientific Computing, Second Edition. Cambridge University Press (1995)
25. Sorkine, O., Cohen-Or, D., Irony, D., Toledo, S.: Geometry-aware bases for shape approximation. IEEE Transactions on Visualization and Computer Graphics **11**(2), 171–180 (2005)
26. Sorkine, O., Cohen-Or, D., Toledo, S.: High-pass quantization for mesh encoding. In: Proceedings of the Eurographics Symposium on Geometry Processing, pp. 42–51. Eurographics Association (2003)
27. Taubin, G.: A signal processing approach to fair surface design. Computer Graphics (SIGGRAPH 1995) **29**(Annual Conference Series), 351–358 (1995)
28. Tobor, I., Reuter, P., Schlick, C.: Reconstruction of implicit surfaces with attributes from large unorganized point sets. In: Proc. of the International Conference on Shape Modeling and Applications, pp. 19–30 (2004)
29. Turk, G., O'Brien, J.F.: Variational implicit surfaces. In: Techical Reports GIT-GVU-99-15 (1999)
30. Turk, G., O'Brien, J.F.: Modeling with implicit surfaces that interpolate. ACM Transactions on Graphics **21**(4), 855–873 (2002)
31. Uccheddu, F., Corsini, M., Barni, M.: Wavelet-based blind watermarking of 3d models. In: Proc. Multimedia and Security Workshop on Multimedia and Security, pp. 143–154 (2004)
32. Wendland, H.: Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. Advances in Computational Mathematics **4**(4), 389–396 (1995)
33. Yin, K., Pan, Z., Shi, J., Zhang, D.: Robust mesh watermarking based on multiresolution processing. Computer and Graphics **25**(3), 409–420 (2001)
34. Zayer, R., Rössl, C., Karni, Z., Seidel, H.P.: Harmonic guidance for surface deformation. In: Proceedings of Eurographics 2005. Blackwell (2005). To appear

Leif Kobbelt is a full professor and the head of the Computer Graphics Group at the RWTH Aachen University, Germany. His research interests include all areas of Computer Graphics and Geometry Processing with a focus on multiresolution and freeform modeling as well as the efficient handling of polygonal mesh data. He was a senior researcher at the Max-Planck-Institute for Computer Sciences in Saarbrücken, Germany from 1999 to 2000 and received his Habilitation degree from the University of Erlangen, Germany where he worked from 1996 to 1999. In 1995/96 he spent a post-doc year at the University of Wisconsin, Madison. He received his Master's (1992) and Ph.D. (1994) degrees from the University of Karlsruhe, Germany. Over the last years he has authored many research papers in top journals and conferences and served on several program committees.

Jianhua Wu graduated in 2002 with a master's degree in computer science from the Tsinghua University in Beijing, China. He then joined the Computer Graphics Group at the RWTH Aachen University, Germany, where he is currently pursuing his Ph.D.'s degree. His research interests mainly focus on efficient surface representations for geometry processing and efficient geometric data structures for distributed and mobile multimedia communications.