Graphical Models 76 (2014) 682–690

Contents lists available at ScienceDirect

Graphical Models

journal homepage: www.elsevier.com/locate/gmod

Hierarchical Laplacian-based compression of triangle meshes ☆

P. Lobaz [a], L. Váša [b,*]

[a] University of West Bohemia, Plzen, Czech Republic
[b] Chemnitz University of Technology, Chemnitz, Germany

ARTICLE INFO

Article history:
Received 22 April 2014
Received in revised form 22 September 2014
Accepted 23 September 2014
Available online 2 October 2014

Keywords:
Compression
Perceptual metric
Laplacian
Discrete shape operator
Encoding
Distortion

ABSTRACT

In this paper, we present an algorithm for efficient encoding of triangle meshes. The algorithm preserves the local relations between vertices by encoding their Laplacian coordinates, while at the same time, it uses a hierarchy of additional vertex constraints that provides global rigidity and low absolute error, even for large meshes. Our scheme outperforms traversal based as well as Laplacian-based compression schemes in terms of both absolute and perceived distortion at a given data rate.

1. Introduction

Triangle meshes are getting increasingly popular as a new medium storing shapes of 3D objects, thanks to recent advances in 3D scanning and even 3D printing. Publishing 3D meshes representing products is an attractive new way of marketing, and sharing 3D meshes may be the next step in computer-aided social interaction. Efficient storage of this kind of data is essential in both allowing displaying of the information even on mobile devices with low bandwidth on the content consumer side, as well as in allowing storage of many highly detailed models on the side of the content distributor.

The task of mesh compression is to store the triangle mesh in a file that is as small as possible. As with other kinds of media, some precision loss is allowable in most applications in order to achieve even smaller file sizes. This problem has been studied for about two decades now, yet only recently scientists started to seriously analyse the perception-related issues arising from the problem. In particular, measuring the amount of distortion due to the precision loss has recently received much attention. New error metrics have been proposed that capture the perceived distortion much better than mean squared error (MSE) and its derivatives. Along with this progress, new compression algorithms have been suggested, which attempt to minimise the perceived distortion.

One of the most efficient algorithms in this regard is the high-pass coding (HPC) proposed by Sorkine et al. [13]. The idea is to express the mesh in terms of local details, using a combinatorial discrete Laplacian. These details are in turn transmitted to the decoder, which solves a linear inverse problem. The approach is very efficient and outperforms all previous methods in terms of perceptual metrics, which usually focus on local similarity of meshes. Despite these advantages, users seem reluctant in adopting this technique, mainly because it lacks a mechanism that would avoid error accumulation. As a result, the performance of HPC in terms of MSE is rather poor, and the algorithm is unable to provide a guarantee

☆ This paper has been recommended for acceptance by Peter Lindstrom.
* Corresponding author.

http://dx.doi.org/10.1016/j.gmod.2014.09.003
1524-0703/

of maximum absolute dislocation of vertex positions. This in turn leads to problems when several meshes interact in 3D space by touching each other – even though the HPC compressed meshes show little perceptual distortion, they may intersect or not touch each other correctly, which is a disturbing artifact.

In this work, we propose an extension of the HPC algorithm which avoids this problem. We include additional data into the data stream, which describe the higher-level relations between vertices of the mesh, thus limiting the error accumulation and providing a considerable improvement as measured by MSE as well as perceptual metrics. Our extension is expressed in the terms of the original HPC encoding algorithm, and thus it does not require significant changes in the encoding/decoding implementation.

2. Related work

Compression of triangle meshes has been intensively studied in the past. Separate approaches have been proposed for compression of mesh connectivity and mesh geometry, and it has been shown that using information from connectivity improves the performance of geometry compression and vice versa. Most commonly, compression schemes build on encoding/decoding connectivity first, followed by a connectivity guided encoding of the geometry.

Several methods of connectivity compression have been proposed in the past. Connectivity is usually encoded without any data loss, while reindexing of indices is usually used in order to reduce the amount of data required. The Topological Surgery approach [15] encodes a vertex spanning tree and a triangle spanning tree that uniquely identify the connectivity. The Edgebreaker scheme [11] provides a guarantee of 4 bits per vertex by encoding the so-called op-codes for each triangle. Valence based coding approaches [1,7] build on the fact that the main part of connectivity information is contained in the vertex degrees. A theoretical bound of 3.245 bits per vertex has been derived [5] under the assumption that every possible connectivity is equally probable. The valence-based encoder provide performance even below this limit, exploiting the higher probability of highly regular (vertex degrees close to 6) nature of most practical meshes.

Encoding of mesh geometry is a task with much more freedom regarding the loss of precision. Most algorithms perform quantization of the floating point values at some stage, while some advanced algorithms have other sources of precision loss, such as neglecting high frequencies in the mesh.

One large class of algorithms works during a mesh traversal, which attaches vertices to the decoded part of the mesh one at the time. The new vertex is predicted in some way, such as using the parallelogram predictor [14] or some of its extensions [3,16]. Finally, a correction vector is encoded that is added to the prediction, yielding the decoded position of the vertex. If both the encoder and the decoder work with the same prediction, then this scheme effectively eliminates error accumulation. An alternative of this approach working with angles rather than position vectors has been also proposed [8].

There have also been more complex algorithms proposed, which do not work in the traversal-based fashion. An algorithm based on eigenvalue decomposition of the mesh connectivity matrix [6] uses transformation of the coordinate functions into a basis of the discrete Laplace operator. A basis reduction is applied to reduce the dimensionality of the data, and the remaining amplitudes are quantized and encoded.

The high-pass encoding [13,2] also builds on the discrete Laplace operator, only this time using it directly to transform the coordinate functions into the Laplacian (delta) coordinates. Together with the anchor points, this data allows the decoder to reconstruct the original vertex positions by solving a sparse system of linear equations. Since we build on this method, it will be described in more detail in Section 3.

Although mean squared error and Hausdorff distance have been used extensively for evaluation of the amount of distortion caused by mesh compression, it has been recently conclusively shown that these metrics provide only limited correlation with distortion perception [4]. User studies have been performed and metrics such as MSDM2 [9], FMPD [19] or DAME [18] have been proposed to provide better correlation with the results. Currently researchers are designing new compression algorithms that minimise the newly proposed metrics, while the traditional metrics keep their relevance in situations where multiple objects interact with each other. While the perception based metrics ensure that each object is visually indistinguishable from the original, the absolute metrics ensure that the interacting objects, such as touching hands or shoe touching a floor, stay in correct position with respect to one another.

Currently, high-pass coding provides the best results in terms of perceived distortion, while the traversal based methods work best in terms of mean squared error. In this paper, we propose an algorithm that outperforms both of these approaches both in terms of perceptual metrics and in terms of mean squared error.

3. Algorithm overview

A triangle mesh $\mathcal{M}$ is defined as a set of vertex positions $v_1, v_2, \ldots, v_V$, representing points in 3D space and referred to as geometry, and a set of index triplets $t_1 = (t_1^1, t_1^2, t_1^3), t_2 = (t_2^1, t_2^2, t_2^3), \ldots, t_T = (t_T^1, t_T^2, t_T^3)$, referred to as connectivity. We assume that the connectivity has been transmitted to the decoder, and thus it is available at both sides of the transmission. The task is to encode the geometry as efficiently as possible.

In the high-pass coding of triangle meshes, the positions of vertices are not encoded as absolute coordinates. Instead, for each vertex, the so-called *Laplacian coordinates* are computed: