

Multiresolution and fast decompression for optimal web-based rendering

Federico Ponchio^a, Matteo Dellepiane^a

^aVisual Computing Lab, ISTI CNR, Pisa, Italy

Abstract

Limited bandwidth is a strong constrain when efficient transmission of 3D data to Web clients and mobile applications is needed. Hence, compression methods can help in reducing the disk occupation and the bandwidth load.

However, while compression ratio is important, the use of Javascript on the Web and low power CPUS in mobile applications led to critical computational costs. Progressive decoding improves the user experience by providing a simplified version of the model that refines with time, and it's able to mask latency. Current approaches obtain this with very poor compression rates or with additional computational costs.

In this paper we present a novel multi-resolution WebGL based rendering algorithm which combines progressive loading, view-dependent resolution and mesh compression, providing high frame rates and a decoding speed of million of triangles per second in Javascript. The method is parallelizable and scalable to very large models.

Keywords: multiresolution, WebGL, 3D Web, web based 3D rendering, online 3D content deployment, mesh compression

1. Introduction

Limited bandwidth and increasing model sizes pose a challenge in the transmission of 3D data to Web clients and mobile applications. Mesh compression is a viable approach to minimize transmission time, and most research focus in this field has been on optimizing compression ratio.

Unfortunately, limited bandwidth often pairs with limited computational power, either because of Javascript environment or low CPU power mobile devices, to the point that for most algorithms decoding time becomes the bottleneck even at moderately low bandwidth. Acceptable rates can be regained reducing compression ratio (for example forfeiting connectivity compression) or using less sophisticate entropy compression algorithms.

A different approach makes use of progressive reconstruction algorithms, which improve the user experience by providing a simplified version of the model that refines while the remaining part of the model is being downloaded. The model converges very quickly at the beginning of the download, and only the details require the full model. However this class of algorithms performs even worse in terms of decoding time (as shown in Limper [1]) or in terms of compression ratio.

Another desirable feature, especially for very large models, is view-dependent resolution: this allows to prioritize the download, decode a specific part of the model and vary resolution of the rendered geometry to maintain a constant screen resolution. This is obtained by maximizing quality at a given frame rate.

Moreover, the 3D models that are now available on the web cover a much broader range of possibilities w.r.t. the past, including point clouds, *triangle soups*, topologically complex geometries, partially textured models. This leads to the necessity to propose a framework which could be robust enough to deal

with different cases.

In this paper we present a novel multi-resolution WebGL based rendering algorithm (Figure 1) which combines progressive loading, view-dependent resolution and a mesh compression providing good rates and a decoding speed of million of triangles per second in Javascript. This method works with textured models, but it can handle non-manifold meshes, and it is also scalable to deal with very large models.

The method is based on a class of multiresolution structures [2, 3] where the “primitive” of the multiresolution is a patch made of thousands of triangles. The processing required to traverse this structure becomes a fraction of triangle based multiresolution algorithms, and allows “batch” operation on the patches: moving data from disk or network to GPU RAM, rendering, and decompression.

In section 3 we describe the improvement made on the multiresolution structure and how the compression algorithm was designed to optimize decoding time while maintaining a good compression ratio. In section 4 we compare it with existing web solutions for mesh compression and progressive visualization, and we analyze the performances when dealing with different classes of 3D models. It represents a solid alternative to current methods, providing a practical mean to handle 3D models on the web.

2. Related Work

This paper is related to several topics in the field of Computer Graphics. Among them, the main are: web-based 3D rendering, progressive and multiresolution rendering approaches,




Figure 1: Progressive refinement of the Happy Buddha: on the upper left corner the size downloaded, on the upper right corner the number of triangles in the refined model. The header and index amount to 8KB

63 and fast decompression methods for 3D models.

64 While a complete overview of all these subjects goes well be-
 65 yond the scope of the paper, in the next subsections we provide
 66 a short description of the state of the art, trying to focus on the
 67 aspects which are more related to the proposed approach.

68

69 2.1. Web-based 3D rendering

70 Three-dimensional content has always been considered as
 71 part of the multimedia family. Nevertheless, especially when
 72 talking about web visualization, its role with respect to images
 73 and videos has always been a minor one. Visualization of 3D
 74 components was initially devoted to external components, such
 75 as Java applets or ActiveX controls [4].

76 After some initial efforts for standardization [5, 6], the pro-
 77 posal of WebGL standard [7], which is a mapping of OpenGL|ES
 78 2.0 [8] specifications in JavaScript, brought a major change.
 79 Several actions related to the use of advanced 3D graphics has
 80 been proposed since then. For a general survey, please refer to
 81 the work by Evans [9]. Since the use of OpenGL commands
 82 needs advanced programming skills, there have been several
 83 actions to provide an "interface" between them and the cre-
 84 ation of web pages. We could subdivide the proposed sys-
 85 tems between *declarative* approaches [10], like X3DOM [11]
 86 or XML3D [12], and *imperative* approaches, like Three.js [13],
 87 SpiderGL [14] and WebGLU [15]. The main difference be-
 88 tween the groups is that the first ones rely on the concept of
 89 *scenegraph*, hence a scene has to be defined in all its elements,
 90 while the second ones provide a more direct interface with the
 91 basic commands. Other systems provide a sort of hybrid ap-
 92 proach [16], where a very simplified scene has to be defined.

93 Evans [9] points out in his survey that declarative approaches
 94 had a major impact in the research community, while imperative
 95 approaches were mainly used in the programming community.
 96 More in general, given the fact that the amount of data that
 97 needs to be sent to the webpage can be quite big, several efforts
 98 about a better organization of generic streamable formats [17,
 99 18] has been proposed. Nevertheless, when complex 3D data
 100 have to be streamed, these structures are not flexible enough to
 101 handle them.

102 In order to face this problem, in the last three years some pro-
 103 gressive compression methods ad hoc for 3D streaming have
 104 been developed. Gobbetti et al. [19] proposed a quad-based
 105 multi-resolution format. Behr et al. [20] transmit different quan-
 106 tization levels of the geometry using a set of nested GPU-friendly
 107 buffers. Lavouè et al. [21] proposed an adaptation for the Web
 108 (reduced decompression time at the cost of a low compression
 109 ratio) of a previous progressive algorithm [22]. Other research
 110 has been also conducted to handle other types of data, like point
 111 clouds [23], which may present different types of issues to face
 112 with.

113 The rendering of textures or textured 3D models has been taken
 114 into account even before the standardization actions. In these
 115 cases the main issue to be faced is the amount of image data:
 116 standard techniques like mip-mapping can be adapted and im-
 117 proved both on the software and hardware side [24]. The is-
 118 sue of handling geometric data has been usually considered a
 119 minor one, due to the usual low complexity of 3D textured
 120 models [25]. Nevertheless, recently complex 3D models with
 121 texture coordinates are available from acquisition devices and
 122 technologies. Next subsections will provide further details.

123 2.2. Progressive and Multi-resolution methods

124 An important feature for user experience when rendering
 125 over slow connections or compressed models is progressive-
 126 ness: the possibility to temporarily display an approximated
 127 version of the model and to refine it while downloading or pro-
 128 cessing the rest of the data.

129 The simplest (and widely used) strategy is to use a discrete
 130 set of increasing resolution models (usually known as Level Of
 131 Detail, LOD). The main drawback with this approach is the
 132 abrupt change in detail each time a model is replaced.

133 A change of paradigm was brought by progressive meshes,
 134 introduced by Hoppe [26]. These meshes encode the sequence
 135 of operations of an edge collapse simplification algorithm. This
 136 sequence is traversed in reverse, so that each collapse becomes
 137 a split, and the mesh is refined until the original resolution. An
 138 advantage of progressive techniques is the much more smooth
 139 transition resolution changes, and the possibility to combine it
 140 with selective refining or view-dependent multiresolution, but

141 this high granularity was achieved at the cost of low compression rates: about 37 bpv with 10 bit vertex quantization.

142
143 A large number of progressive techniques were later developed, but as noted in [1], Table 1, the research focus, however, 144 was on rate-distortion performances and speed was mostly neglected. Latest algorithms still run below 200KTs in CPU. 145
146

147 Mobile and web application would be really too slow using these methods. As a compromise, pop buffers [20] propose 148 a method to progressively transmit geometry and connectivity, while completely avoiding compression. 149
150

151 Another desirable feature, especially for large models, is view-dependent loading and visualization. Most multiresolution 152 algorithms were made obsolete by the increased relative performances of GPU over CPU around the first years of 2000. 153
154 It simply became inefficient to operate on the mesh at the level of the single triangle. Several works [27, 28, 2, 29] achieved 155 much better performances by increasing the granularity of the multiresolution to a few thousand triangles. 156
157

158 The main problem when increasing the granularity is ensuring boundary consistency between patches at different resolution: Yoon [27] and Sander [28] both employ a hierarchical spatial 159 subdivision, but while the first simply disables simplification of most boundary edges, which results in scalability problems, 160 the second relies on global, spatial GPU geomorphing to ensure that progressive meshes patch simplification is consistent 161 between adjacent blocks. The works by Cignoni [29, 2] 162 rely instead on a non hierarchical volumetric subdivision and a boundary preserving patch simplification strategies that guarantee 163 coherence between different resolutions while at the same time ensuring that no boundary persists for more than one level. 164
165 While not progressive in a strict sense, given current rendering speed, the density of triangles on screen is so high that popping 166 effects are not noticeable. 167
168

169 An additional issue when dealing with view-dependent multiresolution techniques is the handling of textured models. While 170 the encoding of texture coordinates can be easily taken into account when creating the patches of different resolution, the 171 boundary consistency among them needs to take into account the texture images. Previous multi-resolution methods [30, 31] 172 proposed solutions for this, but they could fail when dealing with complex geometries. 173
174

175 Compression comes as a natural extension to this family of multiresolution algorithms: each patch can be compressed independently 176 from the others as long as the boundary still matches with neighboring patches. A wavelet based compression was developed 177 in [32] for terrains, a 1D Haar wavelet version in [33] for generic meshes on a mobile application. A comprehensive 178 account of compression algorithms and the convergence with view-dependent rendering of large datasets can be found on a 179 recent survey from Maglo et al.[34]. 180
181

182 2.3. Fast Decompression of 3D models

183 Given that decompression speed is a key factor in order to be able to use compressed mesh, there's been some effort by the 184 community to provide solutions. 185
186

196 Gumhold and Straßer [35] developed a connectivity only compression algorithm that was able to decompress at 800KTs in 197
198 1998. Pajarola and Rossignac [36], in 2000, reported 26KTs for a progressive compression algorithm, and developed a high- 199 performance Huffman decoding identifying entropy compression as a possible bottleneck. 200
201

202 Finally, Isenburg and Gumhold in 2003 [37] developed a streaming approach to compression of gigantic meshes reaching an 203 impressive decompression speed of 2MTs. The method accounts also for texture coordinates. A further work on this was 204
205 proposed in 2005 [38]. 206

207 3. Method

208 Our multiresolution algorithm builds upon the methods described on [2, 3], which is recapped in section 3.1 for completeness. 209 In our solution we adopt a improved partition strategy (see section 3.2), and, more importantly, a novel compression 210 scheme (section 3.3) tailored around the need for decompression speed. 211
212

213 3.1. Batched Multiresolution




Figure 2: First column: before refinement. Second column: after refinement. From top to bottom: a visual representation of the geometric patches representing the model, the model with pure geometry, the model with color information.

215 The model is split into a set of small meshes at different res-
216 olutions that can be assembled to create a seamless mesh simply
217 traversing a tree which encodes the dependencies between each
218 patch, using the estimated screen error to select the resolution
219 needed in each part of the model. To build this collection of
220 patches we need a sequence of non-hierarchical volume par-
221 titions (V-partition) of the the model; non hierarchical means
222 essentially that no boundary is preserved between partitions at
223 different levels of the hierarchy.

224 The data structure is composed of a fixed size header de-
225 scribing the attributes of the models, an small index which con-
226 tains the tree structure of the patches and the position of each
227 patch in the file, and the patches themselves. We use HTTP
228 Range requests to download header and index, ArrayBuffers to
229 parse this structures into Javascript; the patches are then down-
230 load prioritizing highest screen error. Figure 2 shows an exam-
231 ple of a model before and after view-dependent refinement.

232 The rendering requires the traversal of the patch tree, which
233 is usually quite small since each patch is in the range of 16-
234 32K vertices, computing the approximated screen space error
235 in pixel from the bounding sphere and the quadric error (or
236 any other error metric) during simplification. The traversal is
237 stopped whenever our triangle budget is reached, the error tar-
238 get is met or the required patches are still not available.

239 Since the rendering can start when the first patch is down-
240 loaded and the model is refined as soon as some patch is avail-
241 able, this is effectively a progressive visualization albeit with
242 higher granularity. On the other hand, this structure is view de-
243 pendent and thus able to cope with very large models, on the
244 order of hundreds of millions of triangles.

245 3.2. Partition

246 Cignoni et al [3] showed that any non-hierarchical sequence
247 of volume partitions can be the base of a patch based multires-
248 olution structure. Good partition strategies minimize bound-
249 aries, thus generating compact cells. In addition, they allow
250 streaming construction and generate well balanced trees even
251 when the distribution of the model triangles is very irregular.
252 The Voronoi structure, while optimal for boundary minimiza-
253 tion and balance, is not suitable for streaming, leading to long
254 processing times. On the other hand the regular spatial sub-
255 division used in [2] might generate unbalanced trees for very
256 irregular models. This may impact on adaptivity.

257 In our solution each volume partition is defined by the leaves
258 of a KD-tree built on the triangles of the model; to ensure the
259 non hierarchical condition, the split ratio in the KD-tree al-
260 ternates between 0.4 and 0.6 instead of the usual 0.5. This choice
261 allows for streaming processing of the model and good adaptiv-
262 ity. As a bonus, the very regular shape of the patches (see figure
263 2) may be useful when adding texture support.

264 3.3. Mesh Compression

265 Our multiresolution algorithm imposes a set of constrains
266 to mesh compression:

- 267 • each patch needs to be encoded independently from the
268 others, so the method must be efficient and fast even on
269 small meshes

- 270 • boundary vertices, replicated on neighboring patches, need
271 to remain consistent through compression

- 272 • non manifold models must be supported

273 It would be possible to exploit the redundancy of the data
274 due to the fact that the same surface is present in patches at
275 different levels of resolution. We choose not to do so in order
276 to keep the compression stage independent of the simplification
277 algorithm used and to simplify parallel decompression of the
278 patches. Otherwise, we would have to keep track of and enforce
279 dependencies.

280 3.3.1. Connectivity compression

281 We modified the algorithm presented in [39], to support non
282 manifold meshes and surfaces with handles or holes.

283 We need face-face topology for compression and this is com-
284 puted as follows: we create an array containing three edges for
285 each triangle, and sort it so that edges sharing the same vertices
286 will be consecutive (independently of the order of the edges).
287 The edges are then paired taking orientation into account, and
288 all non paired edges are marked as boundary. Non manifold
289 meshes will simply force the creation of some artificial bound-
290 aries.

291 The encoding process starts with a triangle and expands iter-
292 atively adding triangles. The processed region is always home-
293 omorphic to a disk and if the region meets already considered
294 triangles, we consider the common vertices as duplicated. The
295 boundary of the already processed (encoded or decoded) region
296 is stored as a doubly linked list of oriented edges (*active edges*),
297 The list is actually implemented as an array for performances
298 reasons. A queue keeps track and prioritize the *active edges*.

299 The first triangle adds three active edges to the list; itera-
300 tively an edge is extracted from the queue and, if not marked as
301 processed, the following codes are emitted (see Figure 3):

302 **SKIP** if the edge is a boundary edge, or the adjacent triangle
303 has already been encoded; the edge is marked as processed.

304 **LEFT** or **RIGHT** if the adjacent triangle shares two edges
305 with the boundary; The two edges are marked as processed,
306 a new edge added to the queue and its boundary adjacencies
307 adjusted.

308 **VERTEX** if the adjacent triangle shares only one edge with
309 the boundary, in this case the edges is marked as processed and
310 two new edges added to the queue. If vertex of the new triangle
311 opposing the edge was never encountered before its position
312 is estimated using parallelogram prediction and the difference
313 encoded, otherwise its index is encoded (in literature this case
314 is often referred as a “split”). This is a key difference with [39],
315 where in the second case a SKIP code would be emitted, to keep
316 the encoded region simple.

317 If the mesh is composed of several connected components,
318 the process is restarted for each component.

319 The order in which the active edges are processed is im-
320 portant as we would like to minimize the number of VERTEX
321 split operations, and generate a vertex-cache-friendly triangle
322 order. To do so, we simply prioritize the right edges in the
323 VERTEX operation, so that the encoding proceeds in ‘spirals’.




Figure 3: The four decomposition codes: black arrows represent the front, the red arrow the current edge, in green the new edges added to the front.

324 If the mesh is not homeomorphic to a disk, some split opera-
 325 tions are required. This strategy reduces the number of splits
 326 to less than 1% in our examples, incurring in an average of 0.2
 327 bpv cost.

328 This algorithm is certainly not optimal in term of bitrate,
 329 but it is extremely simple, linear in the number of triangles and
 330 robust to non-manifold meshes; as we will see in the results,
 331 speed is more important than bitrate.

3.3.2. Geometry and vertex attribute compression

333 To ensure consistency between boundary vertices of adja-
 334 cent patches, we adopt a global quantization grid for coordi-
 335 nates, normals and colors. The global grid step for vertex posi-
 336 tion quantization is chosen automatically based on the quadric
 337 errors during the simplification step in construction.

338 Geometry and vertex attributes are encoded as differences
 339 to a predicted value. The distribution of these values exhibit a
 340 bias which we can exploit to minimize the number of bits neces-
 341 sary to encode them. Our strategy is based on the assumption
 342 that most of the bias is concentrated on the position of highest
 343 bit (the \log_2 of the value) of these value while the subsequent
 344 bits are mostly random. We simply store in an array, which is
 345 later entropy coded, the number of bits necessary to encode the
 346 value; the subsequent bits are stored in an uncompressed bit-
 347 stream. In this way we need to decode a single symbol, from a
 348 limited alphabet, and read a few bits from a bitstream to decode
 349 a difference.

350 Each new vertex position, result of a VERTEX code, is es-
 351 timated using a simple parallelogram predictor, and the differ-
 352 ences with the actual position encoded as above. Color infor-
 353 mation is first converted into YCbCr color space and quantized,
 354 we encode the difference with one of the corner of the edge pro-
 355 cessed when emitting the VERTEX code. Normals vector are
 356 estimated using the decode mesh position and connectivity, and
 357 differences encoded as usual.

3.3.3. Texture coordinates and textures

358 The main challenge with multiresolution textured models
 359 lies in the simplification algorithm: it needs to take into ac-
 360 count texture seams, and minimize deformation of the texture
 361 on the surface. We used the algorithm employed in Meshlab
 362 [40], based on quadrics in 5 dimensions which include tex-
 363 ture coordinates in selection of the collapse and the vertex op-
 364 timization computation. An optional step could be added which
 365 would modify the texture to take into account the introduced
 366 distortion as described in the work by Chen [41].

368 Textures coordinates in the dataset are stored per vertex,
 369 replicating vertices on texture seams, while texture images are
 370 stored into the dataset as JPEG binaries and loaded on demand
 371 just like the mesh patches. Finally texture coordinates are com-
 372 pressed using the same parallelogram prediction algorithm em-
 373 ployed for the vertex coordinates.

374 Employing more sophisticated methods would drastically
 375 increase decoding time (see timings in table 2 in [42]) or re-
 376 quire additional linear algebra Javascript libraries, for a limited
 377 decrease in bitrate.

3.3.4. Point clouds

379 Enabling support for point clouds requires a few changes,
 380 merging into the same data structure, the functionality of Batched
 381 Multitriangulation [3] and those of Layered Point Clouds [43]:
 382 the simplification algorithm needs to be replaced with a point
 383 filtering approach, where half of the points are removed at each
 384 level and the compression strategy for the vertex coordinates
 385 cannot rely on parallelogram prediction, in this case, after co-
 386 ordinate quantization we sort the points in z-order and store
 387 the differences between consecutive points using the same ap-
 388 proach used for the meshes.

3.4. Entropy coding

390 We have shown how to convert connectivity, geometry and
 391 attributes into a stream of symbols and bits. It is worth com-
 392 pressing the symbol stream due to the biased probability distri-
 393 bution of the symbols.

394 Entropy decoding is the speed bottleneck in many mesh de-
 395 compression methods, often due to the main goal of minimizing
 396 bit per vertex. Pajarola and Rossignac [36] developed a high-
 397 performance Huffman decoding algorithm in order to overcome
 398 this problem. The main advantage of this method is that it re-
 399 duces the decoding phase to a couple of table lookups. Arith-
 400 metic coding, for example, outperforms Huffman in term of
 401 compression rate, but exhibits lower speed. A problem with this
 402 approach is the initialization time required to create the, possi-
 403 bly very large, decoding tables. It is then not suitable for decod-
 404 ing small meshes where the construction time would dominate
 405 over the decoding time.

406 Unlike Huffman and other variable-length codes, Tunstall
 407 code [44] maps a variable number of source symbols to a fixed
 408 number of bits. Since in decompression the input blocks con-
 409 sists of a fixed number of bits and the output is a variable num-
 410 ber of symbols, Tunstall is slightly less efficient than Huffman,
 411 especially where the bit size of the input block is small. The

412 decoding step is very similar to the high-performance Huffman
 413 algorithm, as it consists in a lookup table and a sequence of
 414 symbols for each entry, but the table size is only determined by
 415 the word size, and a fast method to generate it described in [45].

416 Given an entropic source of M symbols, to generate an opti-
 417 mal encoding table for a word size of N bits, we need to gener-
 418 ate 2^N symbol sequences that have a frequency as close as
 419 possible to 2^{-N} , allows to encode every possible input (it is
 420 complete) and no sequence is a prefix of any other sequence
 421 (it is proper).

422 Tunstall optimal strategy starts with the M symbols as initial
 423 sequences, removes the most frequent sequence A and replaces
 424 it with M sequences concatenating A with every symbol until
 425 we reach 2^N sequences. The most time consuming step of the
 426 algorithm is to find the most probable sequence.

427 If we use a matrix where the first column contains the sorted
 428 symbol in order of probability, and at each step we replace the
 429 sequence with highest probability with M sequences adding a
 430 new column, we can observe that this table is sorted both in
 431 columns and rows (see Figure 4). This allows to select the next
 432 sequence by keeping each row in a queue and using a priority
 433 queue to keep track of which queue has the highest front element.
 434

A 0.50	AA 0.25	BA 0.15	AAA 0.125	BAA 0.075
B 0.30	AB 0.15	BB 0.09	AAB 0.075	BAB 0.045
C 0.10	AC 0.05	BC 0.03	AAC 0.025	BAC 0.015
D 0.10	AD 0.05	BD 0.03	AAD 0.025	BAD 0.015

Figure 4: First four steps in construction of a Tunstall code with four symbols, the sequences A, B, AA, BA are replaced with a new column, beside each sequence, its probability is shown. In green the candidates for the next expansion.

435 To initialize the decoding table the symbol frequencies needs
 436 to be transmitted in advance.

437 Finally, an important advantage of variable-to-fixed coding
 438 is that the compressed stream is random accessible: decoding
 439 can start at any block. This makes it especially suited for paral-
 440 lel decompression in particular GPU decompression. Unfortu-
 441 nately, current limitations in the capabilities of WebGL do not
 442 allow for such an implementation.

443 4. Results

444 The testing of the proposed methods has been divided in
 445 two subsections. The first one will be related to the comparison
 446 with existing Compression methods; the second one will show
 447 the performances in rendering on several types of models, in
 448 order to test the robustness and flexibility of the proposed com-
 449 pression paradigm.

450 A demo page, that shows the comparison and a few ex-
 451 amples, is available at <http://fastdec.duckdns.org> (for reviewers
 452 only)

453 Our implementation has been successfully tested on major
 454 browsers on a variety of platform, from desktop machines to
 455 low end cell phones. The results we report here were measured
 456 on an iCore5 3.1Gh, using Chrome 41. Timings taken other
 457 browsers (e.g.Firefox) where comparable. Regarding the mul-
 458 tiresolution model construction, this is a preprocessing opera-
 459 tion. Compression time is negligible, since it can be performed
 460 at about 1M triangles per second. he most cumbersome part is
 461 the quadric simplification algorithm, that runs at about 60K tri-
 462 angles per second per core. Nevertheless, the model construc-
 463 tion must be performed only once.

464 4.0.1. Comparisons with existing systems

465 As already stated, decompression speed is a key factor when
 466 dealing with web streaming of 3D data. Hence, a tradeoff be-
 467 tween compression and speed has to be found to ensure optimal
 468 performances.

469 For this reason, three comparisons have been performed: the
 470 first one is devoted to the measurement of the decompression
 471 speed of some of the reference methods. The second one fo-
 472 cuses on the mesh compression capabilities.
 473 Finally, a third comparison deals with the issue of textured mod-
 474 els, which have been only partially taken into account by the
 475 state-of-the-art method. For this reason, a test against a refer-
 476 ence commercial solution is shown.

477 4.0.2. Entropy Compression: Comparison

478 We tested, both in C++ and Javascript, compression rates
 479 and decompression speed of:

- 480 • our implementation of Tunstall coding (T)
- 481 • Huffman coding (H), in the high-performance version of
 482 Pajarola [36] (our implementation, C++ only)
- 483 • available implementations of LZMA
 484 in C++: <http://www.7-zip.org/sdk.html>
 485 and Javascript: <https://code.google.com/p/js-lzma/>
- 486 • lz-string, a LZW based Javascript implementation
 487 <http://pieroxy.net/blog/pages/lz-string/index.html>

symbols	C++			Javascript		
	T	H	LZMA	T	LZMA	LZW
4	1058	520	1066	201	19	55
9	369	212	170	145	10	23
13	423	168	95	150	6	20
17	359	136	77	163	6	19
22	332	98	67	180	6	17

Table 1: Decompression speed in million of output symbols per second for Poisson distribution of 32K sequences

488 The results are presented in Table 1, the lenght of 32K has
 489 been chosen since it is typical in our application.

490 Huffman and Tunstall are very similar in term of decom-
 491 pression speed, the difference is mainly in the time required to

492 generate the decoding tables which are much larger for Huff-
 493 man, especially when increasing the number of symbols. We
 494 tested also other probability distributions and found little dif-
 495 ference in terms of speed. LZMA and LZW avoid this startup
 496 cost, however their more complex and adaptive dictionary man-
 497 agement allows them to outperform Huffman and Tunstall in
 498 term of decompression speed only for very small runs (and very
 499 small dictionaries). In terms of compression ratio, Huffman and
 500 LZMA performed quite close to the theoretical minimum, while
 501 Tunstall was about 10% worse.

502 We did not implement Huffman in Javascript, as we are con-
 503 fident the result would be very similar. On the other hand the
 504 numbers for LZMA change dramatically. Lz-string serves as a
 505 comparison, as a better library, optimized for Javascript. The
 506 poor LZMA performances in Javascript help explain the rela-
 507 tively slow performances of CTM in Limper [1].

508 4.0.3. Mesh Compression: Comparison

509 We used the Happy Buddha model (in Figure 1), to compare
 510 compression ratio and decompression speed with OpenCTM
 511 (CTM) [46] Pop buffers (POP)[1], P3DW [21], WebGL-loader
 512 (CHUN) [47]. We compare our multiresolution (OUR) and,
 513 to test single resolution performances of our compression ap-
 514 proach, a version (FLAT) which loads only the highest reso-
 515 lution level of the model. In each case the model has been
 516 quantized at 11 bit for coordinates and 8 bit for normals, and
 517 includes colors.

	FLAT	OUR	CTM	CHUN	POP	P3DW
MB	1.9	3.9	3.5	2.8	15	4.5
bpv	28	57	51	41	220	66
full	0.4	0.9	5.3	0.06	0.5	10

Table 2: Statistics for the Happy Buddha: model size in megabytes, bit per vertex and time in seconds required to fully decompress the model.

518 Our decompression Javascript implementation can decode
 519 about 1-3 million triangles per second with normals and colors
 520 in a single thread, on a desktop machine and 0.5 MT/s on a
 521 iPhone Five. Performances are somewhat degraded when the
 522 code is run during streaming visualization.

523 An important comparison is with [33], which employs the
 524 same multiresolution batched strategy. For their mobile mul-
 525 tiresolution application they reports compression rates of 45-50
 526 bpv on large colored meshes (which should be compared to our
 527 28bpv). The difference is probably mostly due to the differ-
 528 ent connectivity encoding which, in their case, requires 20bpv
 529 against our 4 or 5bpv. It is difficult to compare the speed of
 530 the two decompression approaches since they run natively in
 531 C# on an iPhone4 while we run in Javascript on the same plat-
 532 form. Our implementation speed is still, if a bit faster than their
 533 50KTS¹, at about 60KTs. The difference is probably due their
 534 more sophisticate (and slow) arithmetic encoding.

535 C++ decompression speed is of course faster, reaching 9MTs,
 536 including colors and normals, and 16MTs for just position and

537 connectivity. The speed reported in [39] of 35KTs for just the
 538 connectivity, as they mention, is due to the dynamic memory
 539 allocation in their implementation.

540 4.0.4. Textured models

541 In the case of textured models, the compression and quan-
 542 tization has to be applied not only on the geometric attributes,
 543 but also on the texture image.

544 Regarding the latest, using 13 bits quantization for a 4096x4096
 545 pixel texture (which amount to half pixel precision) results in
 546 an hardly noticeable distortion,. (see Figure 5). With parallelo-
 547 gram prediction, texture coordinates are encoded in about 8 bpv
 548 (with a reduction of 48 to 1 respect to the standard 6 floats per
 549 face). We have to include the replicated vertices coordinates,
 550 that in our samples amounted to at most 15% of the total amount
 551 in models with many seams. Overall adding texture coordinates
 552 increases the data size of about 25%, and decompression speed
 553 decreases accordingly.

554 We uploaded a few textured model to Sketchfab, for a com-
 555 parison with a state of the art industrial solution: our multireso-
 556 lution structure results on average 10% smaller, although it in-
 557 cludes all the resolutions. Regarding the performances, please
 558 refer to the examples in the demo page (<http://fastdec.duckdns.org>).

559 4.1. Streaming and Rendering

560 Loading the geometry through the Range HTTP request re-
 561 quires an increased number of HTTP calls: one for each patch,
 562 or 30-60 calls every million of triangles. This does not really
 563 impact over performances: the overhead is quite small (about
 564 400 bytes per call) and pipelining (the process of enqueueing
 565 requests and responses between browser and server) ensures
 566 full utilization of the available bandwidth. Random access is
 567 really necessary only to fully exploit the view-dependent char-
 568 acteristics of the multiresolution structure: the code could be
 569 easily modified to load the model with a single call if a higher
 570 number of HTTP calls was problematic on certain web hosting
 571 architectures.

572 In the demo page (<http://fastdec.duckdns.org>) it is possible
 573 to compare the performances of our method w.r.t. existing solu-
 574 tions in the case of a slow connection. Moreover, very complex
 575 geometries are also available for further testing. In the follow-
 576 ing, we show some example of the performances of the com-
 577 pression method in several types of models, in order to test its
 578 flexibility.

579 4.1.1. Point clouds

580 Point clouds are a quite common type of models, especially
 581 when large environments are taken into account. Terrestrial
 582 laser canners, but also UAV may provide dense point clouds.
 583 Ad-hoc solutions for encoding and rendering have been de-
 584 vised, but their implementation is usually very hard to be ex-
 585 tended to triangulated surfaces.

586 On the contrary, the method proposed in this paper can be seam-
 587 lessly applied also in point clouds: the data that are compressed
 588 and streamed are only the vertices attributes.

¹The number is extrapolated from the decoding time of a large mesh given in their paper




Figure 5: Example of textured model from left to right: texture coordinates quantized at 12 bits, at 13 bits and uncompressed. Notice the slight distortion when precision is lower than half a pixel.




Figure 6: Pompeii point cloud: original PLY file, 95M points, 2.26 Gb; uncompressed multires cloud, 1.68 Gb; compressed cloud, 326 Mb. Top left: the full compressed model, top right: a detail. Bottom left: a detail of the uncompressed point cloud . Bottom right: the same detail of the compressed point cloud



Figure 7: Bonifacio rendered in a browser: original PLY file, 84M triangles, 1.6GB; uncompressed model, 2.54 Gb; compressed model, 158 Mb. Top left: the full compressed model, top right: a detail. Bottom left: a detail of the uncompressed model. Bottom right: the same detail of the compressed model

Figure 6 shows an example of a point cloud of Insula V of Pompei, obtained with terrestrial laser scanning. The top part shows the compressed model and a point of view where all the details are visible. The bottom part of the Figure shows the difference between the compressed and uncompressed point clouds: the quantization of the original data brings to a reduction of points, since the quantization tends to "regularize" the points grid (in this case the quantization step was 0.5 mm). This effect is not noticeable in triangulated surfaces, but it may be an issue for point clouds. In this case, it's possible to change the quantization step in order to find the best tradeoff between compression and data quality.

4.1.2. Dense triangulated models

View-dependent progressive method have been especially devised to handle dense, triangulated 3D models. For this reason, the proposed method is able to provide optimal performances even when hundreds million triangles have to be taken into account.

In the following, two examples of complex geometries are shown.

Figure 7 shows the 3D model of a 3-meter tall statue which was acquired with triangulation structured light scanner. The compressed model, which is nearly 10 % of the original PLY file, exhibits a detail that is undistinguishable from the uncompressed version.

In Figure 8 we show our system rendering the Portalada, a 180M triangles model at 30fps. The triangle budget has been

fixed at 1M triangles and the streaming requires 2-3 seconds to reach full resolution on a good connection. The original model is 3.6GB, while the compressed multiresolution model is 838MB. The Figure also shows how the view-dependent paradigm is able to handle different resolutions of different parts of the model when peculiar points of view are shown.

4.1.3. Non-optimal, topologically complicated models

Some of the solutions proposed for progressive view dependent rendering proved to be limited since their basic assumptions on data processing didn't take into account that most of the more complex 3D models come from acquisition devices or techniques. This leads very often to the presence of geometric artifacts or unbalanced density.

Figure 9 shows two examples where the method deals with non-optimal geometries. On the left side, a model exhibiting strong topological artifacts. On the right side, a model with very unbalanced data density. In both cases, the method is able to deal with the issues and provide an accurate and reliable rendering.

5. Conclusion

The method proposed in this paper provides good compression ratio, progressive visualization, fast decoding and view dependent rendering. It proves effective in a wide range of bandwidth availability, computing power and rendering capabilities. Moreover, it is able to handle a wide variety of 3D models types. This means that also very complex geometries can be now explored in real time with average connections speeds.

Many mesh compression algorithms for mobile and web application do not employ topological connectivity compression often because it is believed to be excessively complex or slow and limited to manifold meshes. We prove that, if implemented correctly, this is not the case, and the choice of the entropy compression algorithm can play a much more important role. The proposed method represents the current best tradeoff between data compression and rendering performances.

5.1. Future improvements

The current implementation of the compression algorithm has the advantage of being able to encode most of the attributes of a 3D model.

Nevertheless, improvements in both compression and rendering performances can be obtained by further exploitation of the characteristics of some types of models.

For example, in the case of Point clouds, the rendering paradigm plays a key role to obtain a satisfying visualization. The current rendering method could be improved by implementing and extending existing approaches [48]. The attributes (i.e. radius) that could be used for efficient rendering can be easily inserted in the compression framework.




Figure 8: Portalada rendered in a browser: original PLY file, 180M triangles, GB; compressed model, 621 Mb. Top left: the full model, top right: a detail of the figure above the arch, middle right: the resolution of the model as seen from the middle left view point (without frustum culling)




Figure 9: Left: a model with severe topological issues. Right: a model with very imbalanced vertex distribution

670 The compression and rendering of textured models can be
671 further improved, by working on ways to better compress and
672 handle textures, or moving to other texturing paradigms. An
673 example could be the projective textures (a similar approach on
674 point clouds was recently proposed by Arikan [49]), that could
675 remove the need for parametrization, and open to even more
676 complex datasets.

677 Acknowledgements

678 The research leading to these results was funded by EU
679 FP7 project ICT Harvest4D (<http://www.harvest4d.org/> , GA
680 n. 323567) and EU INFRA Project Ariadne (GA n. 313193,
681 <http://www.ariadne-infrastructure.eu/>).

682 References

683 References

684 [1] Limper, M., Wagner, S., Stein, C., Jung, Y., Stork, A.. Fast delivery of
685 3d web content: A case study. In: Proceedings of the 18th International
686 Conference on 3D Web Technology. Web3D '13; New York, NY, USA:
687 ACM. ISBN 978-1-4503-2133-4; 2013, p. 11–17.
688 [2] Cignoni, P., Ganovelli, F., Gobbetti, E., Marton, F., Ponchio, F.,
689 Scopigno, R.. Adaptive tetrapuzzles: Efficient out-of-core construction
690 and visualization of gigantic multiresolution polygonal models. ACM
691 Trans on Graphics (SIGGRAPH 2004) 2004;23(3):796–803.
692 [3] Cignoni, P., Ganovelli, F., Gobbetti, E., Marton, F., Pon-
693 chio, F., Scopigno, R.. Batched multi triangulation. In: Proceed-
694 ings IEEE Visualization. Conference held in Minneapolis, MI,
695 USA: IEEE Computer Society Press; 2005, p. 207–214. URL:
696 <http://vcg.isti.cnr.it/Publications/2005/CGGMP05>.
697 [4] Microsoft. Microsoft ActiveX Controls.
698 <http://msdn.microsoft.com/>; 2013.
699 [5] Raggett, D.. Extending WWW to support platform independent virtual
700 reality. Technical Report 1995;.
701 [6] Don Brutzmann, L.D.. X3D: Extensible 3D Graphics for Web Authors.
702 Morgan Kaufmann; 2007.
703 [7] Khronos Group. . WebGL - OpenGL ES 2.0 for the Web. 2009.
704 [8] Khronos Group. . OpenGL ES - The Standard for Embedded Accelerated
705 3D Graphics. 2009.
706 [9] Evans, A., Romeo, M., Bahrehmand, A., Agenjo, J., Blat, J.. 3d
707 graphics on the web: A survey. Computers & Graphics 2014;41(0):43 –
708 61. doi:<http://dx.doi.org/10.1016/j.cag.2014.02.002>.
709 [10] Jankowski, J., Ressler, S., Sons, K., Jung, Y., Behr, J., Slusallek,
710 P.. Declarative integration of interactive 3d graphics into the world-wide
711 web: Principles, current approaches, and research agenda. In: Proceed-
712 ings of the 18th International Conference on 3D Web Technology. Web3D
713 '13; New York, NY, USA: ACM. ISBN 978-1-4503-2133-4; 2013,
714 p. 39–45. URL: <http://doi.acm.org/10.1145/2466533.2466547>.
715 doi:10.1145/2466533.2466547.
716 [11] Behr, J., Eschler, P., Jung, Y., Zöllner, M.. X3dom: a
717 dom-based html5/x3d integration model. In: Proceedings of the
718 14th International Conference on 3D Web Technology. Web3D '09;
719 New York, NY, USA: ACM. ISBN 978-1-60558-432-4; 2009, p.
720 127–135. URL: <http://doi.acm.org/10.1145/1559764.1559784>.
721 doi:10.1145/1559764.1559784.
722 [12] Sons, K., Klein, F., Rubinstein, D., Byelozorov, S., Slusallek,
723 P.. Xml3d: Interactive 3d graphics for the web. In: Proceedings
724 of the 15th International Conference on Web 3D Technology. Web3D
725 '10; New York, NY, USA: ACM. ISBN 978-1-4503-0209-8; 2010, p.
726 175–184. URL: <http://doi.acm.org/10.1145/1836049.1836076>.
727 doi:10.1145/1836049.1836076.
728 [13] Dirksen, J., editor. Learning Three.js: The JavaScript 3D Library for
729 WebGL. Packt Publishing; 2013.

730 [14] Di Benedetto, M., Ponchio, F., Ganovelli, F., Scopigno, R.. Spidergl:
731 a javascript 3d graphics library for next-generation www. In: Proceed-
732 ings of the 15th International Conference on Web 3D Technology. Web3D
733 '10; New York, NY, USA: ACM. ISBN 978-1-4503-0209-8; 2010, p.
734 165–174. URL: <http://doi.acm.org/10.1145/1836049.1836075>.
735 doi:10.1145/1836049.1836075.
736 [15] DeLillo, B.. WebGLU: A utility library for working with WebGL .
737 <http://webglu.sourceforge.org/>; 2009.
738 [16] Potenziani, M., Callieri, M., Dellepiane, M., Corsini, M.,
739 Ponchio, F., Scopigno, R.. 3dhop: 3d heritage on-
740 line presenter. Computer & Graphics 2015;52:129–141. URL:
741 <http://vcg.isti.cnr.it/Publications/2015/PCDCPS15>.
742 [17] Limper, M., Thöner, M., Behr, J., Fellner, D.W.. SRC
743 - a streamable format for generalized web-based 3d data transmis-
744 sion. In: The 19th International Conference on Web3D Technology,
745 Web3D '14, Vancouver, BC, Canada, August 8-10, 2014. 2014, p.
746 35–43. URL: <http://doi.acm.org/10.1145/2628588.2628589>.
747 doi:10.1145/2628588.2628589.
748 [18] Sutter, J., Sons, K., Slusallek, P.. Blast: A binary large struc-
749 tured transmission format for the web. In: Proceedings of the Nine-
750 teenth International ACM Conference on 3D Web Technologies. Web3D
751 '14; New York, NY, USA: ACM. ISBN 978-1-4503-3015-2; 2014,
752 p. 45–52. URL: <http://doi.acm.org/10.1145/2628588.2628599>.
753 doi:10.1145/2628588.2628599.
754 [19] Gobbetti, E., Marton, F., Rodriguez, M.B., Ganovelli, F., Di Benedetto,
755 M.. Adaptive quad patches: An adaptive regular structure for web distri-
756 bution and adaptive rendering of 3d models. In: Proceedings of the 17th
757 International Conference on 3D Web Technology. Web3D '12; New York,
758 NY, USA: ACM. ISBN 978-1-4503-1432-9; 2012, p. 9–16.
759 [20] Limper, M., Jung, Y., Behr, J., Alexa, M.. The pop buffer: Rapid pro-
760 gressive clustering by geometry quantization. Computer Graphics Forum
761 2013;32(7):197–206.
762 [21] Lavoué, G., Chevalier, L., Dupont, F.. Streaming compressed 3d data
763 on the web using javascript and webgl. In: Proceedings of the 18th In-
764 ternational Conference on 3D Web Technology. Web3D '13; New York,
765 NY, USA: ACM. ISBN 978-1-4503-2133-4; 2013, p. 19–27.
766 [22] Lee, H., Lavou, G., Dupont, F.. Rate-distortion optimization for progres-
767 sive compression of 3d mesh with color attributes. The Visual Computer
768 2012;28(2):137–153.
769 [23] Evans, A., Agenjo, J., Blat, J.. Web-based visuali-
770 sation of on-set point cloud data. In: Proceedings of the
771 11th European Conference on Visual Media Production. CVMP
772 '14; New York, NY, USA: ACM. ISBN 978-1-4503-3185-
773 2; 2014, URL: <http://doi.acm.org/10.1145/2668904.2668937>.
774 doi:10.1145/2668904.2668937.
775 [24] Mitting, M., GmbH, C.. Advanced virtual tex-
776 ture topics. In: ACM SIGGRAPH 2008 Games. SIG-
777 GRAPH '08; New York, NY, USA: ACM; 2008, p. 23–51.
778 URL: <http://doi.acm.org/10.1145/1404435.1404438>.
779 doi:10.1145/1404435.1404438.
780 [25] Cohen-Or, D., Mann, Y., Fleishman, S.. Deep compres-
781 sion for streaming texture intensive animations. In: Proceedings
782 of the 26th Annual Conference on Computer Graphics and Inter-
783 active Techniques. SIGGRAPH '99; New York, NY, USA: ACM
784 Press/Addison-Wesley Publishing Co. ISBN 0-201-48560-5; 1999,
785 p. 261–267. URL: <http://dx.doi.org/10.1145/311535.311564>.
786 doi:10.1145/311535.311564.
787 [26] Hoppe, H.. Progressive meshes. In: Proceedings of the 23rd Annual Con-
788 ference on Computer Graphics and Interactive Techniques. SIGGRAPH
789 '96; New York, NY, USA: ACM. ISBN 0-89791-746-4; 1996, p. 99–108.
790 [27] Yoon, S.E., Salomon, B., Gayle, R., Manocha, D.. Quick-vdr: inter-
791 active view-dependent rendering of massive models. In: Visualization,
792 2004. IEEE. 2004, p. 131–138. doi:10.1109/VISUAL.2004.86.
793 [28] Sander, P.V., Mitchell, J.L.. Progressive buffers: View-
794 dependent geometry and texture lod rendering. In: Proceed-
795 ings of the Third Eurographics Symposium on Geometry Pro-
796 cessing. SGP '05; Aire-la-Ville, Switzerland, Switzerland: Euro-
797 graphics Association. ISBN 3-905673-24-X; 2005, URL:
798 <http://dl.acm.org/citation.cfm?id=1281920.1281941>.
799 [29] Cignoni, P., Ganovelli, F., Gobbetti, E., Marton, F., Pon-
800 chio, F., Scopigno, R.. Batching meshes for high per-

- 801 formance terrain visualization. In: Second Annual Confer-
802 ence of Eurographics Italian Chapter. Conference Series;
803 Milan (ITALY): Eurographics Association; 2003, URL:
804 <http://vcg.isti.cnr.it/Publications/2003/CGMPS03b>.
- 805 [30] Sander, P.V., Snyder, J., Gortler, S.J., Hoppe, H.. Texture map-
806 ping progressive meshes. In: Proceedings of the 28th Annual Confer-
807 ence on Computer Graphics and Interactive Techniques. SIGGRAPH
808 '01; New York, NY, USA: ACM. ISBN 1-58113-374-X; 2001, p.
809 409–416. URL: <http://doi.acm.org/10.1145/383259.383307>.
810 doi:10.1145/383259.383307.
- 811 [31] Borgeat, L., Godin, G., Blais, F., Massicotte, P., Laha-
812 nier, C.. Gold: Interactive display of huge colored and
813 textured models. In: ACM SIGGRAPH 2005 Papers. SIG-
814 GRAPH '05; New York, NY, USA: ACM; 2005, p. 869–877.
815 URL: <http://doi.acm.org/10.1145/1186822.1073276>.
816 doi:10.1145/1186822.1073276.
- 817 [32] Gobbetti, E., Marton, F., Cignoni, P., Benedetto, M.D., Ganovelli,
818 F.. C-BDAM – compressed batched dynamic adaptive meshes for ter-
819 rain rendering. *Computer Graphics Forum* 2006;25(3):333–342. *Proc.*
820 *Eurographics* 2006.
- 821 [33] Rodríguez, M.B., Gobbetti, E., Marton, F., Tinti, A..
822 Compression-domain seamless multiresolution visualization of gigan-
823 tic triangle meshes on mobile devices. In: Proceedings of the
824 18th International Conference on 3D Web Technology. Web3D '13;
825 New York, NY, USA: ACM. ISBN 978-1-4503-2133-4; 2013, p.
826 99–107. URL: <http://doi.acm.org/10.1145/2466533.2466541>.
827 doi:10.1145/2466533.2466541.
- 828 [34] Maglo, A., Lavoué, G., Dupont, F., Hudelot, C.. 3d
829 mesh compression: Survey, comparisons, and emerging
830 trends. *ACM Comput Surv* 2015;47(3):44:1–44:41. URL:
831 <http://doi.acm.org/10.1145/2693443>. doi:10.1145/2693443.
- 832 [35] Gumhold, S., Straßer, W.. Real time compression of triangle
833 mesh connectivity. In: Proceedings of the 25th Annual Confer-
834 ence on Computer Graphics and Interactive Techniques. SIGGRAPH
835 '98; New York, NY, USA: ACM. ISBN 0-89791-999-8; 1998, p.
836 133–140. URL: <http://doi.acm.org/10.1145/280814.280836>.
837 doi:10.1145/280814.280836.
- 838 [36] Pajarola, R., Rossignac, J.. Squeeze: Fast and progressive decompression
839 of triangle meshes. In: *Proc. Computer Graphics Int'l (CGI) 2000*.
840 2000, p. 173–182.
- 841 [37] Isenburg, M., Gumhold, S.. Out-of-core compression for gi-
842 gantic polygon meshes. *ACM Trans Graph* 2003;22(3):935–
843 942. URL: <http://doi.acm.org/10.1145/882262.882366>.
844 doi:10.1145/882262.882366.
- 845 [38] Isenburg, M., Lindstrom, P., Snoeyink, J.. Streaming compression of
846 triangle meshes. In: Proceedings of the Third Eurographics Symposium
847 on Geometry Processing. SGP '05; Aire-la-Ville, Switzerland, Switzer-
848 land: Eurographics Association. ISBN 3-905673-24-X; 2005, URL:
849 <http://dl.acm.org/citation.cfm?id=1281920.1281939>.
- 850 [39] Floriani, L.D., Magillo, P., Puppo, E.. Compressing tins. In: In Proceed-
851 ings of the 6th ACM Symposium on Advances in Geographic Information
852 Systems. 1998, p. 145–150.
- 853 [40] Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F.,
854 Ranzuglia, G.. Meshlab: an open-source mesh processing tool. In:
855 Sixth Eurographics Italian Chapter Conference. 2008, p. 129–136. URL:
856 <http://vcg.isti.cnr.it/Publications/2008/CCCDGR08>.
- 857 [41] Chen, C.C., Chuang, J.H.. Texture adaptation for progressive
858 meshes. *Computer Graphics Forum* 2006;25(3):343–350. URL:
859 <http://dx.doi.org/10.1111/j.1467-8659.2006.00953.x>.
860 doi:10.1111/j.1467-8659.2006.00953.x.
- 861 [42] Va, L., Brunnett, G.. Efficient encoding of texture coordinates guided
862 by mesh geometry. *Computer Graphics Forum* 2014;33(5):25–34. URL:
863 <http://dx.doi.org/10.1111/cgf.12428>. doi:10.1111/cgf.12428.
- 864 [43] Gobbetti, E., Marton, F.. Layered point clouds: A simple
865 and efficient multiresolution structure for distributing and render-
866 ing gigantic point-sampled models. *Comput Graph* 2004;28(6):815–
867 826. URL: <http://dx.doi.org/10.1016/j.cag.2004.08.010>.
868 doi:10.1016/j.cag.2004.08.010.
- 869 [44] Tunstall, B.. Synthesis of Noiseless Compression
870 Codes. Georgia Institute of Technology; 1967. URL:
871 <http://books.google.it/books?id=NMLInQECAAJ>.
- 872 [45] Baer, M.. Efficient implementation of the generalized tun-
873 stall code generation algorithm. In: Proceedings of the 2009
874 IEEE International Conference on Symposium on Informa-
875 tion Theory - Volume 1. ISIT'09; Piscataway, NJ, USA: IEEE
876 Press. ISBN 978-1-4244-4312-3; 2009, p. 199–203. URL:
877 <http://dl.acm.org/citation.cfm?id=1701495.1701536>.
- 878 [46] Geelnard, M.. OpenCTM. <http://openctm.sourceforge.net/>;
879 2009.
- 880 [47] Chun, W.. WebGL models: end-to-end. In: Cozzi, P., Ric-
881 cio, C., editors. *OpenGL Insights*. CRC Press; 2012, p. 431452.
882 <http://www.openglintsights.com/>.
- 883 [48] Botsch, M., Hornung, A., Zwicker, M., Kobbelt, L.. High-quality
884 surface splatting on today's gpus. In: *Point-Based Graphics, 2005.*
885 *Eurographics/IEEE VGTC Symposium Proceedings*. 2005, p. 17–141.
886 doi:10.1109/PBG.2005.194059.
- 887 [49] Arikian, M., Preiner, R., Scheiblauer, C., Jeschke, S., Wimmer, M..
888 Large-scale point-cloud visualization through localized textured surface
889 reconstruction. *Visualization and Computer Graphics, IEEE Transactions*
890 *on* 2014;20(9):1280–1292. doi:10.1109/TVCG.2014.2312011.