

Real-time Tessellation of Terrain on Graphics Hardware

Oscar Ripolles^{*},

Universitat Politecnica de Valencia, Valencia, Spain

Francisco Ramos, Anna Puig-Centelles and Miguel Chover

Universitat Jaume I, Castellón, Spain

Abstract

Synthetic terrain is a key element in many applications that can lessen the sense of realism if it is not handled correctly. We propose a new technique for visualizing terrain surfaces by tessellating them on the GPU. The presented algorithm introduces a new adaptive tessellation scheme for managing the level of detail of the terrain mesh, avoiding the appearance of t-vertices that can produce visually disturbing artifacts. Previous solutions exploited the Geometry Shader capabilities to tessellate meshes from scratch. In contrast, we reuse the already calculated data to minimise the operations performed in the shader units. These features allow us to increase performance through smart refining and coarsening.

Key words: Terrain simulation, Tessellation, Level of Detail, Real-time Rendering, GPU

^{*} Corresponding author. Tel.: +34 963877000 ext. 88442
Email addresses: oripolles@ai2.upv.es (Oscar Ripolles),

1 Introduction

In recent years the research area of procedural modeling has been the focus of a lot of effort. The latest works try to take advantage of the new graphics hardware technology, making it possible for the geometry to be generated at rendering time in the graphics card itself. Thus, instead of specifying the details of a 3D object, we provide some parameters for a procedure that will create the object.

In the field of computer graphics, tessellation techniques are often used to divide a surface into a set of polygons. Thus, we can tessellate a polygon and convert it into a set of triangles or we can tessellate a curved surface. These approaches are typically used to amplify coarse geometry. Programmable graphics hardware has enabled many surface tessellation approaches to be migrated to the GPU, including isosurface extraction (Buatois et al., 2006), subdivision surfaces (Shiue et al., 2005), NURBS patches (Guthe et al., 2005), and procedural detail (Bokeloh and Wand, 2006; Boubekour and Schlick, 2005). In this paper we analyse the possibilities offered by GPU-based tessellation techniques for terrain visualisation.

For many decades terrain simulation has been the subject of research, and there are many solutions in the literature to its realistic and interactive rendering. Most of these solutions simulate terrain as an unbounded surface that is represented in the synthetic environment as a heightmap, which is a regularly-spaced two-dimensional grid of height coordinates. These grids can be later processed by a modeling software or a rendering engine to obtain the 3D sur-

francisco.ramos@uji.es (Francisco Ramos), apuig@uji.es (Anna Puig-Centelles), chover@uji.es (Miguel Chover).

24 face of the desired terrain.

25 Some authors have criticised the use of heightfields, as these data structures
26 store only one height value for any given (x,y) pair. In this sense, in specific
27 cases like caves or complex terrain formations it may be possible to have more
28 than one height value for each position. In our case we will not consider these
29 complex formations and, thus, the use of a squared heightmap will still be
30 adequate.

31 We introduce a new adaptive tessellation scheme for terrain that works com-
32 pletely on the GPU. The main feature of the framework that we are presenting
33 is the possibility of refining or coarsening the mesh while maintaining *coher-*
34 *ence*. By coherence we refer to the reuse of information between changes in
35 the level of detail. In this way, the latest approximation extracted is used in
36 the next step, optimising the tessellation process and improving performance.

37 The rest of the paper is structured as follows. Section 2 presents the state of
38 the art in terrain simulation. Section 3 thoroughly describes our tessellation
39 technique. Section 4 offers some results on the technique presented and, lastly,
40 Section 5 presents some conclusions on the techniques developed and outlines
41 future work.

42 **2 Related work**

43 Digital Terrain Models (DTMs) are usually represented and managed by
44 means of regular or irregular grids. The reader is referred to recent surveys
45 for a more in-depth review of these methods ([Pajarola and Gobbetti, 2007](#);
46 [Rebollo et al., 2004](#)).

48 The most common regular structures are quadtrees and binary trees (bintrees).
49 These structures with regular connectivity are suitable for terrain, as the input
50 data usually come as a grid of values. In this sense, regular approaches have
51 produced some of the most efficient systems to date ([Pajarola and Gobbetti,](#)
52 [2007](#)).

53 Quadtrees are found in the literature in many papers with CPU-based solu-
54 tions ([Lindstrom et al., 1996](#); [Pajarola, 1998](#)) as well as GPU-based approxima-
55 tions ([Schmiade, 2008](#)). This latter approach proposed a tessellation algorithm
56 on the GPU, although the pattern selection process was very complex.

57 The ROAM method (Real-time Optimally Adapting Meshes) ([Duchaineau](#)
58 [et al., 1997](#)) is a widely known method based on the use of bintrees. They
59 use two priority queues to manage split and merge operations, obtaining high
60 accuracy and performance. As an attempt to improve this solution, in ([Apu](#)
61 [and Gavrilova, 2004](#)) the authors eliminated the priority queue for merges to
62 exploit frame-to-frame coherence.

63 Some authors proposed using bintrees where each node contains, instead of a
64 single triangle, a patch of triangles ([Levenberg, 2002](#); [Pomeranz, 2000](#)). Algo-
65 rithms like ([Cignoni et al., 2003](#); [Schneider and Westermann, 2006](#)) batched
66 the triangular patches to the graphics hardware. The Batched Dynamic Adap-
67 tive Mesh (BDAM) proposed in ([Cignoni et al., 2003](#)) used triangle strips to
68 increase performance, although it was based on complex data structures and
69 costly processes which still produced popping artifacts. From a different per-
70 spective, ([Larsen and Christensen, 2003](#)) used patches of quads to manage

71 terrain rendering on the GPU. Later, (Schneider and Westermann, 2006) re-
72 duced bandwidth requirements by simplifying the mesh and using progressive
73 transmission of geometry. Recently, in (Bosch et al., 2009) the authors pro-
74 posed the use of precalculated triangle patches to develop a GPU-intensive
75 solution.

76 The projected grid concept offered an alternative way to render displaced sur-
77 faces with high efficiency (Johanson, 2004). The idea was to create a grid with
78 vertices that were evenly-spaced in post-perspective camera space. This repre-
79 sentation provided spatial scalability and a fully GPU-based implementation
80 was described. In (Schneider et al., 2006), Schneider et al. used the projective
81 grid method to render infinite terrain in high detail. They generated the ter-
82 rain in real time on the GPU by means of fractals. The work in (Livny et al.,
83 2008) proposed using ray tracing to guide the sampling of the terrain, being
84 the technique able to produce both regular and irregular meshes.

85 As an improvement over binary trees, (Losasso and Hoppe, 2004) introduced
86 *geometry clipmaps*, caching the terrain in a set of nested regular grids. These
87 grids were stored as vertex buffers, and mipmapping was applied to prevent
88 aliasing. As vertex buffers cannot be modified on the GPU, this approach was
89 later improved by using vertex textures to avoid having to use the CPU to
90 modify the grids (Asirvatham and Hoppe, 2005). This work was also extended
91 to handle spherical terrains (Clasen and Hege, 2006).

92 Lastly, it is worth mentioning that bintree hierarchies are also useful for decom-
93 pressing terrain surfaces on the GPU. In this sense, (Lindstrom and Cohen,
94 2010) presented a fast, lossless compression codec for terrains on the GPU,
95 and demonstrated its use for interactive visualisation.

97 Irregular grids are commonly known as TINs (Triangulated Irregular Net-
98 works) and represent surfaces through a polyhedron with triangular faces.
99 These solutions are less constrained triangulations of the terrain and, in gen-
100 eral, need fewer triangles although their algorithms and data structures tend
101 to be more complex.

102 ([Hoppe, 1998](#)) proposed specialising his View-dependent Progressive Mesh
103 (VDPM) framework for arbitrary meshes that represent terrain. With a more
104 intense GPU exploitation, ([Dachsbacher and Stamminger, 2004](#)) proposed a
105 costly procedural solution that needs three rendering passes to obtain the
106 geometry.

107 Delaunay triangulation is one of the main techniques used to create the ter-
108 rain mesh. In computational geometry, a Delaunay triangulation for a set of
109 points is a triangulation where no point is inside the circumcircle (circle which
110 passes through all the vertices of the triangle) of any generated triangle ([De-
111 launay, 1934](#)). This triangulation has been widely used in terrain solutions
112 ([De Berg et al., 2008](#); [Rabinovich and Gotsman, 1997](#)). The main problem
113 with Delaunay triangulation is that it relies on smoothing morphing between
114 two triangulations generated in two successive frames, but the triangulations
115 may be very different. As an improvement, ([Cohen-Or and Levanoni, 1996](#))
116 proposed a solution that involved blending between two levels of Delaunay
117 triangulation without adding more triangles. More recently, ([Liu et al., 2010](#))
118 proposed a new technique where points from a DEM are initially given an
119 importance value in order to guide the adaptive triangulation in real time.
120 Their proposal allowed for smooth morphing and tried to eliminate very small

121 triangles which could produce visual disturbances.

122 As a conclusion, we can note that techniques based on irregular grids tend to
123 be more complex and less suitable for GPU computations.

124 **3 Our GPU-based Tessellation Scheme**

125 In this paper we propose a new adaptive tessellation algorithm that works com-
126 pletely on the GPU. Moreover, this algorithm is able to offer view-dependent
127 approximations where more detail is added in areas of interest. Our algorithm
128 will be based on bintrees, creating the hierarchy on the GPU using some spe-
129 cific equations. As mentioned above, there have been different proposals with
130 similar aims, although our scheme is easier to implement while still robust and
131 efficient.

132 A successful tessellation algorithm is based on the selection of the most suit-
133 able tessellation patterns to amplify the triangles. These patterns affect the
134 algorithm chosen to refine and coarsen the geometry. As our aim is to process
135 the mesh in a *Geometry Shader*, each triangle is to be processed separately
136 and in parallel. Thus, it will be necessary to develop a technique to alter the
137 geometry of the different triangles without any communication between them.
138 Moreover, the algorithms must assure that no cracks or holes appear on the
139 surface mesh.

140 In the remainder of this section we will address the selection of the patterns
141 and also the algorithms to manage the terrain surface.

143 It is possible to find in the literature different proposals of tessellation patterns.
144 Among them, we have selected the seven patterns presented in (Schmiade,
145 2008). Figure 1 presents, on the left side, an initial rectangular triangle where
146 its hypotenuse and catheti (or *legs*) are labeled as H , C_1 and C_2 respectively.
147 Next, the seven tessellation patterns are presented, where the edges of the
148 original triangle that need refinement are depicted in red.

149 These patterns assure that no *t-vertices* are produced. A *t-vertex* appears after
150 a tessellation step when two edge junctions make a *t-shape* (McConnell, 2006).
151 To clarify the appearance of *t-vertices*, Figure 2 presents the initial geometry
152 of a terrain composed of three triangles. According to some criterion, it is
153 decided to refine the middle triangle and vertex v_5 is introduced, outputting
154 4 triangles. Then, if we decided to apply a heightmap to this geometry, we
155 would find holes around vertex v_5 , as the triangle on top has no reference to
156 this vertex. In the example offered, vertex v_5 would be a *t-vertex*.

157 The different patterns presented in Figure 1 offer a robust tessellation and
158 avoid cracks and holes. We must note that our tessellation algorithm is based
159 on the use of an edge-based criterion to decide which triangles to refine and
160 which to coarsen. This is an improvement over triangle-based criteria, as
161 they tend to introduce *t-vertices*. In this sense, each pattern shows the tessel-
162 lation that would be necessary depending on the combination of edges that
163 need refinement. We use the center point of each edge to perform the appro-
164 priate calculations. For example, pattern 3 considers a situation in which the
165 hypotenuse needed refinement and a new vertex has been added to create two
166 new triangles.

167 In addition to the edge-based criterion, we also must select a metric to decide
168 whether an edge should be refined or coarsened. Although in most cases the
169 distance to the camera is the selected metric, it would be possible to apply
170 more accurate heuristics that balance the perceived visual quality and the
171 extraction time that the tessellation process needs.

172 3.2 Tessellation algorithm

173 At each iteration of the tessellation process, the algorithm checks each trian-
174 gle to see whether it is necessary to refine or coarsen it. More precisely, the
175 algorithm checks the center point of each edge according to the selected met-
176 ric. The resulting combination of edges that need processing indicates which
177 pattern should be applied.

178 For the correct performance of our tessellation scheme, it is necessary for each
179 triangle to store:

- 180 • The spatial, texture and normal coordinates.
- 181 • A number indicating the *id* of the triangle.
- 182 • A number coding the tessellation patterns that have been applied (*patternInfo*).

183 The need of storing the *id* and *patternInfo* values is due to the fact that we
184 must know how a particular triangle was created in order to know how we
185 should modify it when swapping to a lower level of detail. This information is
186 crucial as it allows the algorithm to coarsen the geometry without having to
187 return to the initial mesh. This is one of the main contributions of our algo-
188 rithm, as coherence among extracted approximations considerably increases
189 the rendering performance.

190 On the one hand, the *id* value will uniquely identify each of the generated
191 triangles and will also allow us to calculate the *id* of its parent triangle. This
192 *id* number is calculated following the formula:

$$193 \quad id = id * maxOutput + originalTris + childType \quad (1)$$

194 The *maxOutput* value is understood as the maximum number of triangles
195 that can be output from a parent triangle using the available patterns. The
196 patterns presented in Figure 1 involve outputting a maximum number of four
197 new triangles and, as a consequence, in our case *maxOutput* is a value equal
198 to 4. The *originalTris* constant refers to the number of initially existing tri-
199 angles on the source mesh, which depends on the input mesh the application
200 uses. Finally, *childType* is a value used to differentiate between the triangles
201 output from a parent triangle. As the patterns output a maximum number of
202 four new triangles, in our case the *childType* is a value ranging from 0 to 3.
203 The *childType* value assures that each *id* is different, allowing us to distin-
204 guish between the triangles that belong to the same parent. This feature is
205 compulsory for the correct simplification of the mesh.

206 On the other hand, the *patternInfo* number is used to store all the patterns
207 that have been applied to refine a triangle. Equation 2 has been specifically
208 prepared to code the different patterns applied to a triangle in one single value.
209 In this equation, *latestPattern* refers to the type of the latest pattern, the
210 one that we have used to create this triangle. The value *numberOfPatterns*
211 refers to the number of available patterns that we can apply in our tessellation
212 algorithm. In our case we use the seven patterns presented in Figure 1 and, as a
213 consequence, *numberOfPatterns* should be equal to 7. It is worth mentioning
214 that, initially, all the *patternInfo* values are equal to 0. This *patternInfo*

215 value will be the same for all the triangles belonging to the same parent.
216 This piece of information is important to know how a particular triangle was
217 created and, consequently, how we should modify it when swapping to a lower
218 level of detail.

$$219 \quad patternInfo = patternInfo * numberOfPatterns + latestPattern \quad (2)$$

220 *Refining algorithm*

221 When *refining* the mesh, the algorithm checks the center point of each edge
222 to see whether they need refinement. Depending on the combination of edges
223 that need more detail, the algorithm selects a tessellation pattern (see Figure
224 1) and generates the adequate number of triangles. For each new triangle, the
225 algorithm calculates its spatial coordinates, texture information and any other
226 information needed for rendering. To clarify the process, Figure 3 presents an
227 example of how the tessellation process works. We present the initial mesh
228 composed of three triangles, initially labeled with *ids* 0, 1 and 2. The dotted
229 line represents the plane that we will use to define which area of the mesh needs
230 refinements, being the area on the left the one that requires more detail. Each
231 of the initial triangles goes through the extraction process of the algorithm
232 that we are presenting.

233 In the specific case of triangle number 2, the algorithm detects that none of
234 its edges needs refinement and, as a consequence, no change will be made.
235 Nevertheless, the algorithm detects that triangle with *id* 0 needs refinement
236 because the center point of some of its edges is on the left of the dotted line.
237 Then, the algorithm chooses pattern 6 as it reflects the combination of edges
238 to be refined. Using this pattern, the tessellation process algorithm generates

239 the three new triangles shown in the figure. It can be seen how the *id* values
240 of the new triangles are calculated following the formula 1, assuring that no
241 repeated *id* is given. The triangle with *id* 1 is similarly refined using pattern
242 6.

243 Following on with the refinement process, the next tessellation step shows that
244 different patterns have been applied to obtain different types of tessellation.
245 In the figure we depict how triangle 5 is refined with pattern 2 and triangle
246 7 with pattern 7. It is important to mention that this figure also includes the
247 *patternInfo* of the different triangles, which is calculated using Equation 2.

248 *Coarsening algorithm*

249 A different process should be applied when diminishing the detail of the mesh.
250 The *id* and the *patternInfo* values of the triangles have been precisely given
251 in order to simplify the coarsening process. Using the example given in the
252 previous subsection, let us suppose that we want to reduce the detail and
253 return to the state shown in the middle of Figure 3. In this case, each of
254 the triangles located on the left of the dotted line would execute the same
255 coarsening process.

256 When tessellating a triangle, for example with the pattern that is used when
257 the hypotenuse and both legs are refined (see pattern 7 in Figure 1), four
258 triangles are output. Nevertheless, only one of these triangles will be needed
259 when diminishing the detail. As we will see in the remaining of this Section,
260 three of them will be discarded and the other one will be modified to recreate
261 the geometry of the parent triangle.

262 When coarsening the mesh, the first step is to find out whether the triangle
 263 that we are processing can be discarded or if it is the triangle in charge of
 264 retrieving the geometry of the parent triangle. The *childType* used when cal-
 265 culating the *id* of each triangle is necessary for this particular differentiation.
 266 In those cases where this value is equal to 0, the algorithm assumes this trian-
 267 gles is in charge of recovering the geometry of the parent triangle; if the value
 268 is not equal to 0, the triangle is discarded. The *childType* can be retrieved by
 269 using Equation 3.

$$270 \quad \textit{childType} = \textit{mod}((\textit{id} - \textit{originalTris}), \textit{maxOutput}) \quad (3)$$

$$271 \quad \textit{id} = (\textit{id} - \textit{originalTris})/\textit{maxOutput} \quad (4)$$

272 The second step entails knowing which pattern was applied to create the
 273 existing triangle. This is due to the fact that for each pattern we will perform
 274 different calculations for retrieving the three vertices of the parent triangle.
 275 In this situation, the *patternInfo* value helps us to know which pattern was
 276 applied, as the latest pattern can be obtained with the next equation:

$$277 \quad \textit{latestPattern} = \textit{mod}(\textit{patternInfo}, \textit{numberOfPatterns}) \quad (5)$$

278 Once we know which pattern was applied, we calculate the position of the
 279 vertices and we output the new geometry with the new *id* value obtained in
 280 Equation 4 and the new *patternInfo* value obtained with Equation 6. The way
 281 we calculate these values assures that we will be able to continue coarsening
 282 the mesh or refining it without any problem.

$$283 \quad \textit{patternInfo} = \textit{patternInfo}/\textit{numberOfPatterns} \quad (6)$$

284 Following on with the example presented in Figure 3, let us suppose that we

285 are processing the triangle with *id* 34. If we calculate its *childType* we obtain a
286 value greater than 0, indicating that it can be discarded. Nevertheless, triangle
287 31 has a *childType* value equal to 0 and, thus, it is the one used to recover the
288 parent triangle. The *id* of the parent triangle can be obtained with Equation 4.
289 In this case, the *latestPattern* would indicate that pattern 7 was applied and
290 we would calculate the spatial coordinates of the parent triangle accordingly.
291 Once again we would like to remember that all these operations have been
292 coded in the shaders, so that the algorithm knows which operations to perform
293 depending on the type of pattern applied.

294 *Global algorithm*

295 Once we have described the main characteristics of our algorithm, we must
296 consider how the refining and coarsening processes work together. The refine-
297 ment process is executed at each frame while the criterion is met and until
298 we reach a maximum tessellation level, which is defined by the application.
299 Similarly, the coarsening process is performed at each frame until the original
300 geometry is obtained.

301 Nevertheless, in a real application the surface representing the terrain is refined
302 and coarsened at the same time, as the tessellation conditions are modified
303 while the user navigates through the scene. Our algorithm is capable of han-
304 dling multiple levels of detail on the mesh, as the tessellation is applied to
305 each edge of the triangles individually. In this sense, it is possible that in the
306 triangle some edges need refinement and some need simplification. In these
307 cases, and as it happened in the examples above, the algorithm would choose
308 the most suitable pattern that fits this situation.

309 We must note that we will not store precomputed patterns on GPU memory
310 as other solutions do ([Boubekeur and Schlick, 2005](#)). We just code in the
311 *Geometry Shader* the seven cases that we follow (see [Figure 1](#)) so that the
312 coordinates of the new vertices can be calculated from the coordinates of
313 existing vertices when refining and coarsening the triangles.

314 Finally, it is worth mentioning that the last step of our algorithm includes
315 retrieving the height of the newly computed vertices from the heightmap,
316 which is previously stored in the GPU. It can be seen as the use of a displace-
317 ment map to alter the position of each vertex ([Szirmay-Kalos and Umenhoffer,](#)
318 [2008](#)).

319 In [Figure 4](#) we present an overview of the management of the heightmaps.
320 First, the whole heightmap is allocated into main memory. Before the ren-
321 dering stage starts, the area to be initially processed is uploaded to graphics
322 memory. When the area of interest changes, a new texture should be up-
323 loaded to GPU memory. In order to avoid GPU stalls during these texture
324 streamings, our approach uses asynchronous updates by means of Pixel Buffer
325 Objects. A Pixel Buffer Object [Elhassan \(2005\)](#) is simply an array of bytes in
326 GPU memory. However, this type of object can improve performance because
327 it allows the graphics driver to streamline writing to video memory and to
328 schedule asynchronous transfers. Thus, CPU does not need to wait for the
329 texture transfer to be completed. In [Figure 5](#), we present a graphical compar-
330 ison between the conventional manner to load a texture from main memory
331 to graphics memory and the alternative method by using a Pixel Buffer Ob-
332 ject. The conventional method requires the CPU to perform all the transfer
333 processes. On the contrary, with the PBO, the CPU still has to perform the
334 transfer of data, but transferring data from the PBO to the texture object is

335 managed by the GPU. Therefore, OpenGL performs these transfer operations
336 without the CPU intervention and asynchronous operations in memory can
337 be scheduled while rendering.

338 The texture upload could be triggered when the user approaches one of the
339 limits of the terrain. When this situation happens, the systems starts stream-
340 ing the new heightmap to the PBO, replacing those areas which are no longer
341 used. At the same time,

342 4 Results

343 In this section we will study the performance of our tessellation method by
344 analysing the visual quality obtained as well as the calculation time of the
345 extracted approximation. Our scheme was programmed with GLSL and C++
346 on a Windows Vista Operating System. The tests were carried out on a Pen-
347 tium D 2.8 GHz. with 2 GB. RAM and an nVidia GeForce 8800 GT graphics
348 card.

349 4.1 Visual Results

350 First, we offer some visual results of the tessellation algorithm that we have
351 described. Figures 6 and 7 present a mesh in wireframe where different tessel-
352 lations have been applied. These figures show how the tessellation process is
353 capable of increasing the detail of an input mesh without introducing cracks
354 or other artifacts.

355 Figure 6 presents a tessellation case where an initial mesh (on top) is refined

356 according to the distance to the camera. In this Figure the height values are
357 recovered from a heightmap stored as a texture on the GPU. We have also
358 included an image of the texturing process that can also be applied in our pro-
359 cess, as the algorithm can calculate the texture coordinates when tessellating
360 the surface mesh.

361 We can find another tessellation example in Figure 7 where five tessellation
362 steps are presented. In this case, we have considered that a fictitious frustum
363 has been located on the mesh to guide the tessellation process which considers
364 the distance to the camera. It is important to mention that some areas of
365 the mesh that are outside the frustum are also tessellated in order to avoid
366 T-vertices, as we explained when describing our proposal. From a different
367 perspective, in this case we have tested our method with a heightmap in geotiff
368 format [Sazid and Ramakrishnan \(2003\)](#). Figure 4 shows the area that covers
369 the map, which has a size greater than 1.5 GBytes and an error of around
370 25 meters. This terrain is located in Spain and has been extracted from a
371 public web service. The whole map was initially allocated in main memory. In
372 case of requiring more space than that available in main memory, it would be
373 compulsory to resort to out of core techniques [Silva et al. \(2002\)](#); [Varadhan
374 and Manocha \(2002\)](#). On the GPU side, graphics memory has a limited size
375 (512 Mbytes in the graphics card that we have used). Thus, we also made use
376 of an OpenGL extension (PBO or Pixel Buffer Object [Elhassan \(2005\)](#)) that
377 enabled us to stream the heightmap from the CPU to the GPU, as commented
378 in Section 3.2.

380 In order to evaluate the performance of our tessellation technique, we have
381 conducted some tests where an initial mesh composed of 4 triangles is tes-
382 sellated. The detail of the input mesh is first increased and later coarsened
383 following a smooth trajectory of the camera.

384 Figure 8 presents the time needed for tessellating and rendering the input
385 mesh at different tessellation levels. In this case the tessellation depends on
386 the distance to the camera. Table 1 presents the results obtained in this test,
387 helping us to show how the calculations for tessellation suppose an average
388 increase of 60%.

389 For offering further tessellation experiments, Figure 9 presents the results of
390 a similar test where all the geometry is tessellated at the same time, without
391 any specific criterion. In this case, the obtained geometry will be composed of
392 2^n triangles, where n is the tessellation step. In this case, we can observe how
393 the cost of the tessellation is exponential, offering very high temporal costs
394 when outputting a large number of triangles. Again, the results are depicted
395 in Table 2 to help us analyse the way this tessellation algorithm works. It
396 is worth mentioning that, in our simulation, we will never include so many
397 triangles as only those areas that need detail will be tessellated. Nevertheless,
398 we considered it to be interesting in order to show how the temporal cost of
399 the algorithm can be affected by the quantity of output triangles.

401 An important contribution of the proposed approach is the possibility of ex-
402 ploiting coherence among the extracted tessellations. Table 3 presents the
403 temporal results of a scenario similar to that presented in Figure 8, where
404 the distance to the camera is used to guide the tessellation. These temporal
405 costs include visualisation and tessellation of the input mesh. The column
406 on the right offers the results without coherence maintenance, which nearly
407 double the cost of our coherence-based algorithm. These results show that we
408 can offer better performance as our tessellation scheme can exploit coherence
409 among extracted tessellation, in contrast to previous solutions which had to
410 start again from the input mesh.

411 **5 Conclusions**

412 In this article we have presented a new fully-GPU tessellation technique which
413 offers view-dependent approximations. The scheme proposed avoided the ap-
414 pearance of T-vertices and other artifacts that can produce holes in the surface
415 of a terrain. Another important aspect of this tessellation algorithm was the
416 coherence exploitation, as it is capable of reusing the latest approximations
417 when refining and coarsening the mesh. In this sense, we minimise the oper-
418 ations to perform in both cases, reducing the temporal cost involved in the
419 tessellation process. This coherence maintenance is possible by storing some
420 small pieces of information in each triangle, which is sufficient for altering the
421 level of detail. It is important to underline that previous solutions were not ca-
422 pable of managing coherence, and thus entailed costlier tessellation processes.
423 In addition, we have also considered a simple yet efficient approach to manage

424 the heightmap information on the GPU.

425 A triangle-based criterion.....

426 For future work we would like to use larger terrains and consider out-of-core
427 meshes, where all the geometry of the mesh does not fit within the memory on
428 the GPU. From a different perspective, the appearance of Directx 11 involves
429 further advances in computer graphics. Among the new stages of the rendering
430 pipelines, we could highlight the tessellation unit, which will be able to produce
431 semi-regular tessellations ([Tariq, 2009](#)) by itself. In this sense, for future work
432 we would like to study the possibilities offered by the new tessellation units,
433 in order to adapt our algorithm to this new framework.

434 **Acknowledgements**

435 This work has been supported by the Spanish Ministry of Science and Tech-
436 nology (projects TSI-2004-02940, TIN2007-68066-C04-02 and TIN2007-68066-
437 C04-01) by Bancaja (project P1 1B2007-56) and by ITEA2 (project IP08009).

438 **References**

- 439 Apu, R. A., Gavrilova, M. L., 2004. Gtvis: Fast and efficient rendering system
440 for real-time terrain visualization. In: International Conference on Compu-
441 tational Science and Applications (ICCSA). pp. 592–602.
- 442 Asirvatham, A., Hoppe, H., 2005. Terrain rendering using GPU-based geom-
443 etry clipmaps. In: GPU Gems 2. pp. 27–45.
- 444 Bokeloh, M., Wand, M., 2006. Hardware accelerated multi-resolution geometry

445 synthesis. In: I3D '06: Proceedings of the 2006 symposium on Interactive
446 3D graphics and games. pp. 191–198.

447 Bosch, J., Goswami, P., Pajarola, R., 2009. Raster: Simple and efficient terrain
448 redering on the GPU. In: Proceedings Eurographics Areas Papers. pp. 35–
449 42.

450 Boubekeur, T., Schlick, C., 2005. Generic mesh refinement on GPU. In:
451 HWWS '05: Proceedings of the ACM Siggraph/Eurographics conference
452 on Graphics hardware. pp. 99–104.

453 Buatois, L., Caumon, G., Lvy, B., 2006. GPU accelerated isosurface extraction
454 on tetrahedral grids. In: International Symposium on Visual Computing. pp.
455 383–392.

456 Cignoni, P., Ganovelli, F., Gobbetti, E., Marton, F., Ponchio, F., Scopigno,
457 R., 2003. Planet-sized batched dynamic adaptive meshes (p-bdam). In: VIS
458 '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03). p. 20.

459 Clasen, M., Hege, H.-C., 2006. Terrain rendering using spherical clipmaps. In:
460 EUROVIS - Eurographics /IEEE VGTC Symposium on Visualization. pp.
461 91–98.

462 Cohen-Or, D., Levani, Y., 1996. Temporal continuity of levels of detail in
463 delaunay triangulated terrain. In: VIS '96: Proceedings of the 8th conference
464 on Visualization. pp. 37–42.

465 Dachsbacher, C., Stamminger, M., 2004. Rendering procedural terrain by ge-
466 ometry image warping. In: Proceedings of Eurographics Symposium on Ren-
467 dering. pp. 103–110.

468 De Berg, M., Cheong, O., Van Kreveld, M., Overmars, M., 2008. Computa-
469 tional Geometry: Algorithms and Applications. Springer-Verlag.

470 Delaunay, B., 1934. Sur la sphre vide. a la memoire de georges voronoi. Otde-
471 lenie Matematicheskikh i EstestvennyhNauk 7, 793–800.

472 Duchaineau, M., Wolinsky, M., Sigeti, D., Miller, M., Aldrich, C., Mineev-
473 Weinstein, M., 1997. Roaming terrain: real-time optimally adapting meshes.
474 In: VIS '97: Proceedings of the 8th conference on Visualization. pp. 81–88.

475 Elhassan, I., 2005. Fast texture downloads and
476 readbacks using pixel buffer objects in opengl.
477 [http://developer.nvidia.com/system/files/akamai/gamedev/docs/
478 Fast_Texture_Transfers.pdf?display=style-table&download=1](http://developer.nvidia.com/system/files/akamai/gamedev/docs/Fast_Texture_Transfers.pdf?display=style-table&download=1).

479 Guthe, M., Balázs, A., Klein, R., 2005. GPU-based trimming and tessellation
480 of nurbs and t-spline surfaces. *ACM Transactions on Graphics* 24 (3), 1016–
481 1023.

482 Hoppe, H., 1998. Smooth view-dependent level-of-detail control and its appli-
483 cation to terrain rendering. In: VIS '98: Proceedings of the conference on
484 Visualization '98. pp. 35–42.

485 Johanson, C., 2004. Real time water rendering-introducing the projected grid
486 concept. Tech. rep., Master of Science Thesis, Lund University.

487 Larsen, B. D., Christensen, N. J., 2003. Real-time terrain rendering using
488 smooth hardware optimized level of detail. *The Journal of WSCG* 11 (2),
489 282–289.

490 Levenberg, J., 2002. Fast view-dependent level-of-detail rendering using cached
491 geometry. In: VIS '02: Proceedings of the conference on Visualization '02.
492 pp. 259–266.

493 Lindstrom, P., Cohen, J. D., 2010. On-the-fly decompression and rendering
494 of multiresolution terrain. In: I3D '10: Proceedings of the 2010 ACM SIG-
495 GRAPH symposium on Interactive 3D Graphics and Games. pp. 65–73.

496 Lindstrom, P., Koller, D., Ribarsky, W., Hodges, L. F., Faust, N., Turner,
497 G. A., 1996. Real-time, continuous level of detail rendering of height fields.
498 In: SIGGRAPH '96. pp. 109–118.

499 Liu, X., Rokne, J. G., Gavrilova, M. L., 2010. A novel terrain rendering algo-
500 rithm based on quasi delaunay triangulation. *Visual Computer* 26, 697–706.

501 Livny, Y., Sokolovsky, N., Grinshpoun, T., El-Sana, J., 2008. A GPU persistent
502 grid mapping for terrain rendering. *The Visual Computer* 24 (2), 139–153.

503 Losasso, F., Hoppe, H., 2004. Geometry clipmaps: terrain rendering using
504 nested regular grids. *ACM Trans. Graph.* 23 (3), 769–776.

505 McConnell, J., 2006. *Computer Graphics: Theory Into Practice*. Jones and
506 Bartlett Publishers.

507 Pajarola, R., 1998. Large scale terrain visualization using the restricted
508 quadtree triangulation. In: *VIS '98: Proceedings of the conference on Vi-
509 sualization '98*. pp. 19–26.

510 Pajarola, R., Gobbetti, E., 2007. Survey of semi-regular multiresolution mod-
511 els for interactive terrain rendering. *Vis. Comput.* 23 (8), 583–605.

512 Pomeranz, A. A., 2000. Roam using surface triangle clusters (rustic). Tech.
513 rep., University of California at Davis.

514 Rabinovich, B., Gotsman, C., 1997. Visualization of large terrains in resource-
515 limited computing environments. In: *Proceedings of the 8th conference on
516 Visualization*. pp. 95–102.

517 Rebollo, C., Remolar, I., Chover, M., Ramos, J. F., 2004. A comparison of
518 multiresolution modelling in real-time terrain visualisation. In: *ICCSA* (2).
519 pp. 703–712.

520 Sazid, S., Ramakrishnan, R., 2003. GeoTIFF - A standard image file format
521 for GIS applications. <http://www.geospatialworld.net/images/pdf/117.pdf>.

522 Schmiade, T., 2008. Adaptive GPU-based terrain rendering. Master's thesis,
523 Computer Graphics Group, University of Siegen.

524 Schneider, J., Boldte, T., Westermann, R., 2006. Real-time editing, synthesis,
525 and rendering of infinite landscapes on GPUs. In: *Vision, Modeling and*

526 Visualization 2006. pp. 145–152.

527 Schneider, J., Westermann, R., 2006. GPU-friendly high quality terrain ren-
528 dering. *The Journal of WSCG* 14 (1-3), 49–56.

529 Shiue, L., Jones, I., Peters, J., 2005. A real-time GPU subdivision kernel. *ACM*
530 *Transactions on Graphics* 24 (3), 1010–1015.

531 Silva, C. T., Chiang, Y. J., El-Sana, J., Lindstrom, P., 2002. Out-of-core
532 algorithms for scientific visualization and computer graphics. In: *IEEE Vi-*
533 *sualization Conference 2002*.

534 Szirmay-Kalos, L., Umenhoffer, T., 2008. Displacement mapping on the GPU
535 - State of the Art. *Computer Graphics Forum* 27 (1).

536 Tariq, S., 2009. D3D11 tessellation. *Game Developers Confer-*
537 *ence. Session: Advanced Visual Effects with Direct3D for PC,*
538 [http://developer.download.nvidia.com/presentations/2009/GDC/](http://developer.download.nvidia.com/presentations/2009/GDC/GDC09_D3D11Tessellation.pdf)
539 [GDC09_D3D11Tessellation.pdf](http://developer.download.nvidia.com/presentations/2009/GDC/GDC09_D3D11Tessellation.pdf).

540 Varadhan, G., Manocha, D., 2002. Out-of-core rendering of massive geometric
541 environments. In: *Proceedings of the conference on Visualization'02*. pp.
542 69–76.

543 **List of Figures**

544	1	Tessellation patterns (Schmiade, 2008). The red colour	
545		indicates the edges that need refinement.	29
546	2	Example of <i>t-vertex</i> (v_5) after a tessellation step.	30
547	3	Tessellation example with the <i>id</i> value of each triangle. The	
548		<i>patterinInfo</i> value of each triangle is also shown.	31
549	4	Workflow to process and render a terrain surface in the GPU	
550		by using our approach.	32
551	5	Graphical comparison between the conventional manner to	
552		load a texture into the graphics memory and the alternative	
553		one by using Pixel Buffer Objects, which enable us to perform	
554		asynchronous operations with no CPU intervention.	33
555	6	Sample tessellation using a heightmap to modify the terrain	
556		surface.	34
557	7	Sample tessellation guided by a simulated frustum and using	
558		a heightmap from Spain obtained from a public web service.	
559		Geometry is refined up to 2,348 triangles.	35
560	8	Performance obtained using a distance criterion.	36
561	9	Performance obtained when completely tessellating the mesh.	37

Number of triangles	Visualisation	Visualisation + Tessellation
4	1.45	2.28
16	1.46	2.29
64	1.48	2.28
256	1.58	2.45
1,024	1.71	2.75
3,644	1.83	3.16
5,756	1.88	3.95
3,644	1.83	3.39
1,024	1.71	2.75
256	1.58	2.45
64	1.48	2.28
16	1.46	2.28
4	1.45	2.28

Table 1

Comparison of time (in milliseconds) required for visualising and tessellating the input mesh using a distance criterion, by first increasing and then decreasing the detail following a smooth camera trajectory.

Number of triangles	Visualisation	Visualisation + Tessellation
4	1.45	2.29
16	1.46	2.29
64	1.48	2.44
256	1.58	2.44
1,024	1.71	2.76
4,096	1.80	3.17
16,384	2.08	4.56
65,536	2.71	6.42
262,144	4.89	9.31
562,500	7.05	15.23
262,144	4.89	10.96
65,536	2.71	6.83
16,384	2.08	5.16
4,096	1.80	3.59
1,024	1.71	2.76
256	1.58	2.8
64	1.48	2.44
16	1.46	2.29
4	1.45	2.29

Table 2

Comparison of time (in milliseconds) required for visualising and tessellating if completely tessellating the mesh, by first increasing and then decreasing the detail following a smooth camera trajectory.

Number of Triangles	Coherence Exploitation	No Coherence Exploitation
16	2.29	2.49
64	2.29	3.63
256	2.28	4.89
1,024	2.75	5.46
3,644	3.16	7.55
5,756	3.95	8.04

Table 3

Performance comparison (visualisation and tessellation) with and without exploiting coherence (in milliseconds).




Fig. 1. Tessellation patterns ([Schmiade, 2008](#)). The red colour indicates the edges that need refinement.



Fig. 2. Example of t -vertex (v_5) after a tessellation step.




Fig. 3. Tessellation example with the *id* value of each triangle. The *patterinInfo* value of each triangle is also shown.

Fig. 4. Workflow to process and render a terrain surface in the GPU by using our approach.

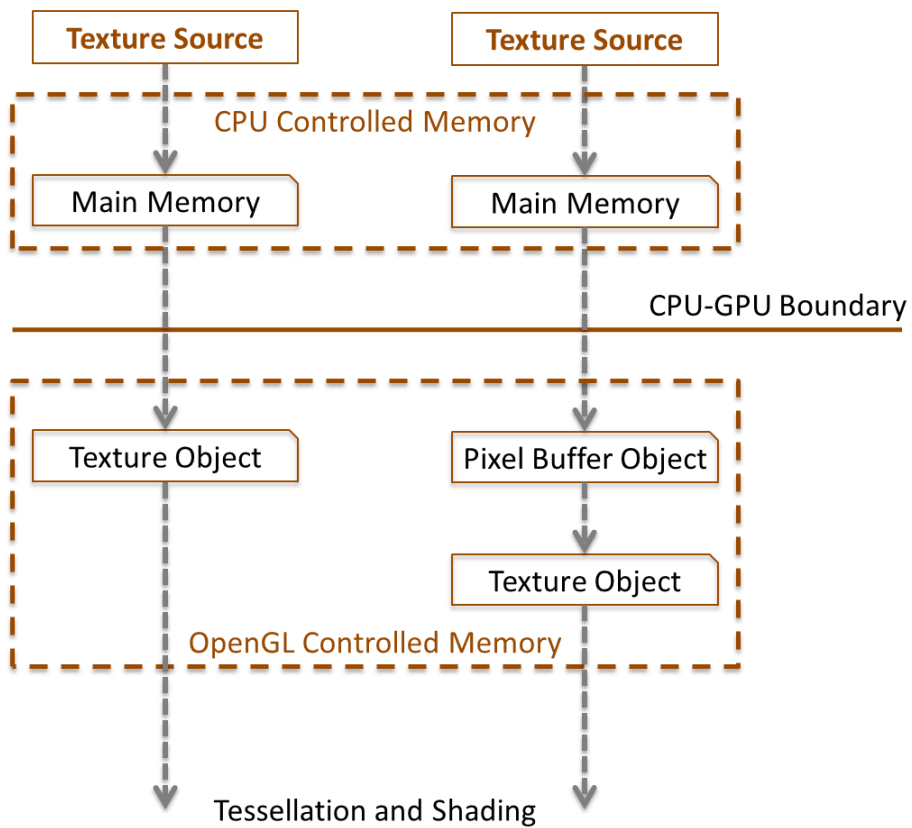


Fig. 5. Graphical comparison between the conventional manner to load a texture into the graphics memory and the alternative one by using Pixel Buffer Objects, which enable us to perform asynchronous operations with no CPU intervention.



Fig. 6. Sample tessellation using a heightmap to modify the terrain surface.



Fig. 7. Sample tessellation guided by a simulated frustum and using a heightmap from Spain obtained from a public web service. Geometry is refined up to 2,348 triangles.



Fig. 8. Performance obtained using a distance criterion.




Fig. 9. Performance obtained when completely tessellating the mesh.