

Real-Time Label Visualization in Massive Model Objects

Anonymous submission

Abstract

Virtual Labels are used in computer graphics applications to represent textual information arranged on geometric surfaces. Such information consists of names, numbering, or other relevant data that need to be noticed quickly when a user scans the objects in the scene. This paper focuses on the so-called massive models, which have a large number of geometric primitives whose rendering presents a high computational cost and for which conventional texturing techniques can extrapolate the available computational resources. In this work we have developed a way to view, in real time, virtual labels with different information on the surfaces of objects in massive models. The technique is implemented entirely on the GPU, and shows no significant loss of performance and low memory cost. CAD models' objects are the main focus of the work, although the solution can be used in other types of objects once their texture coordinates are adjusted correctly.

Keywords: Surface labeling, Real time visualization, Massive Models.

1. Introduction

Virtual labels are textual information displayed on the surfaces of virtual objects. Such information consists of names, numbers, and any relevant information that needs to be identified when the user scans the objects in the virtual scenario. Labels could be used in games, maps or any other graphical application which renders 2D or 3D objects and needs to show text on those objects' surfaces. This paper focuses on real-time 3D visualizers of models with a large number of objects, like CAD (Computer Aided Design) models used for scientific visualization such as an oil refinery.

In industrial plants there are structures that contain labels applied to their surfaces, which display the names of these structures or other information about them, in order to help professionals to easily identify them in fieldwork. The same way that these labels help engineers identify the real structures, virtual labels can help users in identifying objects in visualization software in real time.

The main goal when displaying labels on the surfaces of virtual objects is to help the user to quickly identify relevant information about the objects being displayed on the application, as well as to provide support for guidance in the virtual scenario. In CAD visualization systems, it is usual to have a list showing the names of all the objects present in the scene. Thus, to find the name of an object, the user needs to click on the ob-

ject and check in the list for the selected name, which is a more laborious process than reading labels directly on the objects. Moreover, large industrial plant models are typically massive, so called when they present great complexity, such as an oil refinery consisting of millions or billions of objects [1].

For this kind of model, there are some challenges for rendering virtual labels; one of them is related to the existence of a large number of objects with distinct names. To illustrate this problem, if a simple texturing method is used, it could be necessary to build a different texture for each object (because each object has a different name). Such an approach implies a large use of memory space that can exceed the capacity of the video card. To solve this problem, it is necessary to develop strategies to keep in video card memory a limited number of textures per frame. It is also possible that all textures could not be stored in main memory at the same time and thus they would need to be reconstructed to form the words of a label in real time, which can affect the performance of the application.

Even in a scenario where one can store a different texture for each object, there is still the need to perform multiple texture context switches at every frame, which can become the bottleneck of the graphical application [2]. Another challenge is the positioning of labels on objects with different geometries and sizes, making it necessary that the virtual labels are automatically po-

sitioned by the application in all cases.

These problems (excessive use of video memory, many texture context switches per frame, and the positioning of labels on different surfaces) encourage the exploration of solutions for the real time display of labels in massive models. Since the visualization of massive models has itself a high computational cost, it is necessary to devise a solution that does not worsen the performance and spends the least possible amount of memory.

This paper presents a technique for displaying virtual labels in massive models implemented entirely in the GPU. This technique does not cause significant performance loss regardless of the number of objects that display labels and consumes little video memory with labels coded and stored in a buffer in the graphics hardware. The technique relies on a two-pass algorithm. The first pass performs the rendering of the scene and stores texture coordinate information in off-screen buffers. The second pass uses the information from these buffers to map the labels on virtual objects, calculating the final color of each pixel of the screen individually by sampling a texture that contains all characters.

This paper is organized as follows. Section 2 presents work related to the rendering of textures with textual information. Section 3 describes the technique developed in this work and Section 4 discusses how to position labels on some CAD objects. Section 5 shows the results and Section 6 concludes the paper.

2. Related work

There are only a few works that directly address the problem of real time text rendering in massive models. This section briefly discusses some of these works and shows a few others that address more general problems regarding text rendering (such as label positioning, memory management, and aliasing) and that relate them to issues if applied to the case of virtual labels in massive models.

The commercial software Aveva Review [3] can display labels on the surfaces of objects from CAD models, which have unique names. However, only cylinders and boxes display labels and not all of them show their labels at the same time — there is a strategy to select which objects display their names at each frame; apparently, it prioritizes the ones with a larger size in screen space.

Something that can make an application limit the number of objects that display labels is the number of texture context switches that may be required at each

frame, since each object has a different name. To solve the problem of multiple texture context switches, NVIDIA [2] explains the concept of a “texture atlas” showing the benefits of grouping images into a single texture called an atlas, and using texture coordinates to access the relevant sub-rectangle of this atlas. In this paper we use an atlas that stores all the required characters, and the labels need to be built from this atlas.

There are other works that do not directly address massive models, but address issues related to text rendering. Qin, McCool and Kaplan [4] present a way to perform text rendering on three-dimensional objects using a vector graphics image representation of text, so that characters show very little aliasing when magnified. They use a strategy similar to the one in “Texture Sprites” [5] to place the words on 3D objects. However, they use a quadtree in the CPU to store a table of characters with different resolutions depending on the complexity of each character and a “sprites table” to arrange the words. The table is encoded in the graphics hardware as a texture, so that it can be sampled directly in the GPU. As in “Texture Sprites”, the technique is presented for a single object at a time.

Cipriano and Gleicher [6] focus on text positioning on surfaces with sharp curvatures and even with holes. For this, they create meshes, called scaffolds, to store the letters and then place them so that they float over objects. A single atlas is used to store all the characters and the texture coordinates of the scaffolds are configured in order to obtain the relevant characters. One difficulty in using this solution is the introduction of new geometries in the scene, which may be undesirable in massive models rendering, especially when all objects need to display some text.

The use of virtual textures as an out-of-core solution for textures that occupy a very large space and cannot be entirely stored in RAM or VRAM was explained by Barret [7] and Mittring [8]. A “virtual texture” is stored on the hard drive and split into parts of equal size, called pages. Only the necessary pages for the graphics API are transferred to memory and stored in a “physical texture” as called by Barret. To map the virtual pages into physical pages, a page table is built.

In the case of virtual labels, they would need to be built as textures in pre-processing and arranged in a virtual texture, which could solve the problem of memory. The problem lies in how to organize the virtual texture, because the names of the objects vary in size and it would require that the page size (which is the same for all pages) be sufficient to contain the biggest name. This would imply pages with unused space, which is a waste of memory. Another way to organize the virtual texture

would be with a label occupying more than one page, and objects therefore needing more than one page to store their textual content. The solution adopted in the present work took another approach, creating only the necessary labels (those that will appear on the screen) at each frame and spending little memory to save them, as will be shown in the next section.

3. Applying labels to objects

This section explains the procedure for displaying labels on a large number of objects and the writing of words on geometric surfaces in real-time. To do this, a two-pass algorithm in the GPU was developed. For the implementation of this algorithm we used OpenGL 2.1 [9], OpenGL Shading Language 1.2 [10] and a few extensions.

The first pass of the algorithm renders the scene to an off-screen color buffer and also creates a second buffer with texture coordinates per pixel and an identifier of the label which must be displayed on the geometry to which the pixel belongs. This information will be useful for the second pass when the pixels' final colors are calculated and labels are added to the scene.

To calculate the final colors of the pixels, it is necessary to know whether they should show part of a label and which textual information that label contains. For this reason, before the rendering takes place, the ASCII codes of all labels must be stored and made available for the GPU. In this work we used a 2D texture for the label's ASCII codes that is referred to as a table, relating the label index with its text.

The number of texels that a 2D texture can store is hardware dependent, but in general there is enough space for a large quantity of textual information (e.g., 256MB for a GeForce 9600GT and even more for modern graphic cards). Of course, considering that the models to be visualized are massive and their geometries can consume plenty of space in the video memory, it is not be desirable to spend that much with labels. However, if each label has 20 characters and 2,000,000 distinct labels are needed, the memory consumption will be 38MB, what isn't high considering the large amount of information stored.

Also, a texture atlas that contains all the required characters must be available in the GPU, which is treated as a two-dimensional array of characters. The atlas is an image with the graphical representation of the characters arranged in the same order of the ASCII encoding, so that it is possible to find, in that image, the position of a character by its row and column initials.

Using the atlas with all characters and the texture with ASCII codes created by the application, two rendering passes are performed to render the scene with labels applied to the geometries. We are going to explain them in detail in the sub sections below.

3.1. First rendering pass

Before the drawing commands of the objects are called, the CPU informs the GPU of the label identifier that must be displayed on each object (index in the texture of ASCII codes) and the number of characters of that label.

The first pass is necessary for the generation of texture coordinates for the objects. We generated the texture coordinates in the vertex shader, as will be explained in section IV. The fragment shader writes information into two off-screen buffers that are textures attached to a Frame Buffer Object (FBO). In the first buffer, the color data of the scene are written, and in the second buffer, texture coordinates interpolated per fragment, the label identifier, and the label's text size are written.

3.2. Second rendering pass

In order to obtain the data per fragment generated by the first pass, it was necessary to draw a screen-sized quad, whose texture coordinates range from 0 to 1 in both dimensions. This makes it possible to sample the color buffer and the buffer with texture information.

After setting new information of color and texture to a fragment of the quad, the fragment shader performs calculations to add labels to the scene. If the fragment has no label, it receives, in the first pass, a label identifier equal to zero and texture coordinates equal to zero, so that no calculations for labels are done in the second pass. The calculation for mapping labels on an object is done as follows for each fragment that has a label.

3.2.1. Find character index

The algorithm considers that the label will be written parallel to the s axis of the object's texture space. This space can vary from 0 to 1, or to a value greater than 1, to repeat the label along s . For each interval of size 1 at the s axis, one must consider that the interval will be divided in a total of columns equal to the number of characters in the label. The purpose of this step is to figure out in which column the current fragment is contained, which indicates the index of the character in the object's label.

For example, in Fig. 1, for a label with 4 characters ($n_{chars} = 4$), given a fragment with texture coordinate

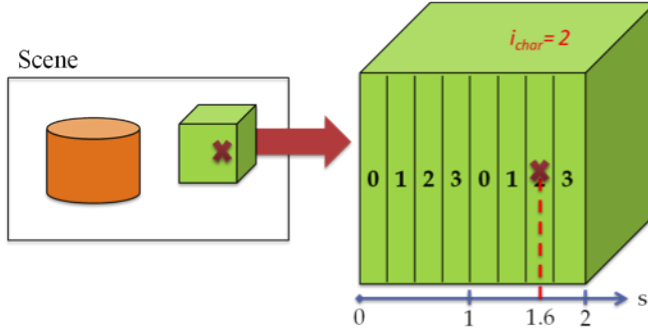


Figure 1: Fragment with character index = 2.

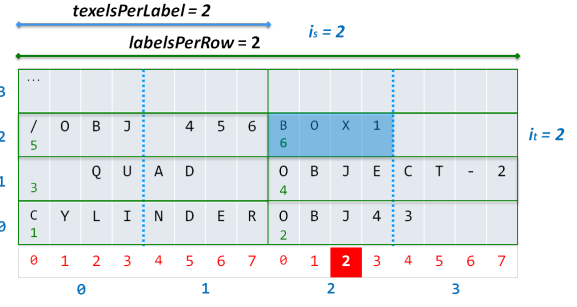


Figure 2: ASCII codes stored in a 2D texture.

255 $s = 1.6$, the expected result is a character index $i_{char} = 2$.
 256 For that, we use equation (1), where s_{frac} indicates the
 257 fractional part of texture coordinate s . This fractional
 258 part indicates the offset of the fragment within one in-
 259 stance of the label, and is important for dealing with
 260 label repetitions.

$$i_{char} = \text{floor}(s_{frac} * n_{chars}) \quad (1)$$

261 3.2.2. Find the ASCII code of the character

262 As explained before, there is a texture that stores the
 263 ASCII codes of the characters of all labels in the scene.
 264 The goal of this step is, with the index of the character
 265 found in the previous step and the label identifier of the
 266 current fragment, to find, within this texture, the ASCII
 267 code of the character sought.

268 Fig. 2 shows an example of how the charac-
 269 ters are stored in the texture with textual information
 270 that, in this case, uses two texels to store each label
 271 ($\text{texelsPerLabel} = 2$). Since a texel has four color com-
 272 ponents, in two texels it is possible to store up to eight
 273 characters. The values in green on Fig. 2 are the la-
 274 bel identifiers and in red are the indices of each char-
 275 acter within its labels (highlighted in the figure, the in-
 276 dex found on the previous step $i_{char} = 2$). The values
 277 in blue show that the texture in the given example pos-
 278 sesses four texels horizontally and four vertically and,
 279 at the same time, represent an index for each texel along
 280 the s and t axes. As each label in the example occupies
 281 2 texels, it is possible to store 2 labels per row of the
 282 texture ($\text{labelsPerRow} = 2$).

283 To obtain the character's ASCII code, it is first nec-
 284 essary to find the texel containing the character. In the
 285 example of Fig. 2, as highlighted, for a fragment in
 286 which label identifier (id) is 6 and $i_{char} = 2$, we will find
 287 the texel with horizontal index $i_s = 2$ and vertical index
 288 $i_t = 2$ (values in blue on the figure). These calculations
 289 are done according to equations (2-3) below:

$$i_s = ((id - 1) \bmod \text{labelsPerRow}) * \text{texelsPerLabel} + \text{floor}(i_{char}/4) \quad (2)$$

$$i_t = \text{floor}\left(\frac{id - 1}{\text{labelsPerRow}}\right) \quad (3)$$

290 Using the indices above, the texture coordinates of
 291 the texel must be found so that the 2D texture can be
 292 sampled, according to equations (4-5), where w and h
 293 are, respectively, the width and height of the texture, in
 294 texels.

$$s_{temp} = \frac{i_s + 0.5}{w} \quad (4)$$

$$t_{temp} = \frac{i_t + 0.5}{h} \quad (5)$$

295 The purpose of equations (4) and (5) is to find the
 296 texture coordinates whose values represent the texel's
 297 center, and with a nearest neighbor filter, the texel can
 298 be correctly retrieved from the texture. The obtained
 299 texel contains the correct character in one of its RGBA
 300 components. To finally find the ASCII code in the cor-
 301 rect component, the calculation shown in equation (6)
 302 is done, with $i_{char} \bmod 4$ an index that goes from 0 to
 303 3 due to the RGBA components, and $char$, the ASCII
 304 code we are looking for:

$$char = \text{texel}[i_{char} \bmod 4] \quad (6)$$

305 In the example shown in Fig. 2, the expected result
 306 will be the code of the 'X' character.

3.2.3. Sample the character atlas

Knowing the character's code, it is then necessary to find its graphical representation in the texture atlas containing all characters. More specifically, the correct texel within the bounds of a character's image in the atlas must be found. For that, the offset of the fragment in the geometry's texture space must be calculated and related to an offset in the atlas, as shown in Fig. 3.

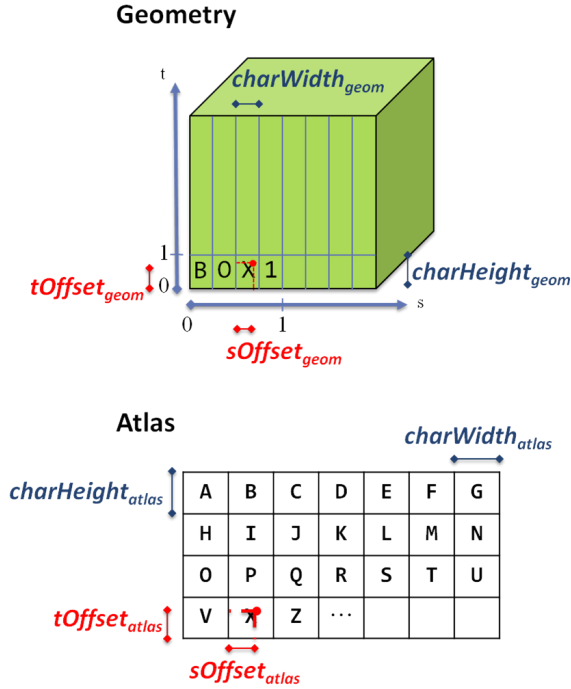


Figure 3: On top, a fragment and its offsets from the beginning of the character on the object. Below, the related offsets in the atlas.

To calculate the offset along the s axis in the geometry texture space ($sOffset_{geom}$), at first, one must calculate the width of the character in a 0 to 1 range. This value ($charWidth_{geom}$) is found simply by dividing 1 by the number of characters in the label. The offset is calculated with equation (7), using once again, the fractional part of the s coordinate of the fragment (s_{frac}) and the character index obtained previously (i_{char}).

$$sOffset_{geom} = s_{frac} - (i_{char} * charWidth_{geom}) \quad (7)$$

The calculation of the offset along the t axis ($tOffset_{geom}$) is simpler, since the height of the character ($charHeight_{geom}$) is always 1, even considering the vertical repetition of text. This way, $tOffset_{geom}$ is simply the fractional part of the t coordinate of the fragment.

To find the related offsets in the atlas's texture space, it is necessary only to multiply the offsets found in the geometry's texture space by the ratio of the character's size in the geometry's texture space to the size of the same in the atlas's texture space, a value known since the creation of the atlas. This way:

$$sOffset_{atlas} = sOffset_{geom} * \frac{charWidth_{atlas}}{charWidth_{geom}} \quad (8)$$

$$tOffset_{atlas} = sOffset_{geom} * charHeight_{atlas} \quad (9)$$

In equation (9), the denominator of the fraction doesn't appear because, as we mentioned earlier, $charHeight_{geom}$ is always 1. The offsets calculated with equations (8-9) must be added to the character's initial position in the atlas, to find the texture coordinate and then to sample the atlas in the correct position to obtain a color for the current fragment.

With the color obtained from the atlas and the color of the object (which is stored in the color buffer created in the first pass) it is possible to calculate the fragment's final color treating the label as a sticker.

With the presented technique, it is noticeable that, independently of the number of objects loaded by the application, label calculation is needed only for the pixels of the screen, which justifies the use of the technique for massive models. In the second pass, some equations with simple operations are calculated, and a few texture accesses are needed. In relation to these accesses, the algorithm benefits from hardware resources for texturing, since the fragments next to each other, that are processed in parallel, in most cases, need to access near parts of the textures. The technique for displaying labels did not present any significant performance impact with massive models, as will be shown in Section 5. The next section discusses the way to position labels in objects of CAD models.

4. Label Positioning

In this section we are going to explain the procedure to position labels on some surfaces of CAD models, with the goal of providing to the user a good way to visualize the text, independently of the angle of vision. We are going to show that, with the correct definition of texture coordinates, the algorithm for displaying labels works for objects of different shapes and sizes. The proposed algorithm considers that the texture coordinates of the geometries must be defined during the first pass.

Those coordinates can vary from 0 to 1, or from 0 to any value greater than 1 to repeat the text along the geometry.

For all presented primitives, the texture coordinates are generated in the vertex shader of the first pass. Such procedure can be done because the developed graphics engine uses a primitive instantiation technique, which shapes the different geometries in the vertex program from a grid of vertices built in the CPU. For that, the vertex program needs to know information regarding the dimension, scale, position and orientation of the object and those data need to be informed by the CPU before the drawing of each object. The scale and dimension data are essential to automatic generation of texture coordinates for the geometries with different shapes and sizes, and for having labels with a varying character quantity.

Firstly, the technique was tested in simple geometries, such as quads — since it was easy to extend the idea to other objects. Texture coordinates can be defined in several ways, but it was concluded that it would be better to centralize the labels vertically on the objects so that the visualization of the 3D models wouldn't be polluted. Horizontally, it would be better to repeat the text as many times as necessary, so that the user is able to visualize the text or easily knows how to find it from any position from which he/she is examining an object.

To centralize the text vertically, only a determined area of the object should have coordinates varying from 0 to 1. Thus, the fragments whose texture coordinates are outside that interval are identified in the fragment shader of the first pass and associated to a label identifier equal to zero — which makes the second pass to calculate the final color of those fragments without the mapping of labels. The procedure to centralize the text calculates the ideal number of times that a label must be repeated along the surface (considering the object's size). Initially an ideal number of label repetitions must be defined based on a visual observation of a unit quad, and then, those parameters will provide a base for the automatic texture coordinate generation for quads of any size — for instance, if it was noticed that vertically it was ideal to repeat the text 3 times so that the label wasn't stretched or squeezed, then in a quad of size 5 the text could be repeated 15 times. However, text should be displayed only at the center of the object; therefore, the texture coordinates should be shifted from 0 to 15 to -7 to 8, totaling 15 repetitions, in such a way that the area which goes from 0 to 1 stays exactly in the center of the object.

To repeat the text horizontally it is also necessary to visually identify how many characters fit in a unit quad

to obtain a good visual effect. For instance, if it is identified that a unitary quad fits 6 characters, then a quad of size 5 fits 30 characters. If a label has, for example, the text “Quad 0” of 6 characters, then it is possible to repeat the text 5 times along the s axis. Therefore, knowing a base for unit quads, it is possible to create simple equations for quads of any size.

One thing that has also to be considered for the calculation of the number of repetitions is a relation between the width and height of the objects. If the scale to be applied to the texture coordinates is calculated separately for each dimension, the bigger the geometry, the bigger is the disproportion between the object's size and the label's size. If a quad has dimensions 1×1 or 100×100 the screen area occupied by the label is the same, but it could be larger in the second case when an equal increase in both dimensions of the text's area is desirable, and no label repetition is actually needed for either dimension. Repetition of text is only needed when one dimension of the object is bigger than the other. Therefore, the ratio between the object's dimensions must be considered for a better label display. If the ratio between width and height is 1, then a scale of 1 must be applied to the s and t coordinates. However, if the width is 100 times bigger than the height, the scale that should be applied to s is around 100 times greater than the one that should be applied to t . By taking into consideration the ratio between the object's dimensions, there is no disproportion between the label and the object's size (Fig. 4).

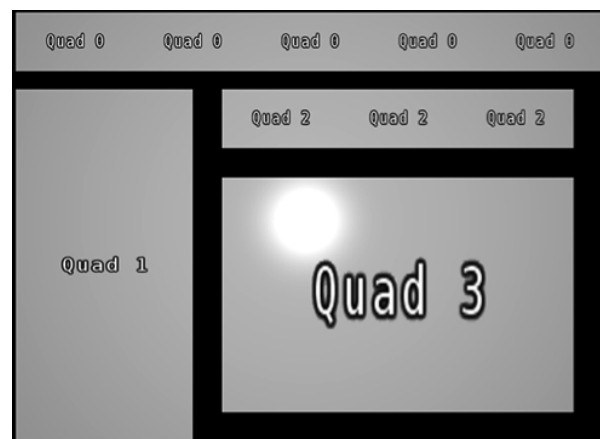


Figure 4: Quads of different sizes and automatic positioning of labels.

Other primitives that have their texture coordinates adjusted to display labels are boxes, cylinders, spheres and dishes. In the case of boxes, the calculations are about the same as the quads, however, different texture

457 coordinates must be defined for each of the box's faces.

458 For cylinders, the texture coordinates of the vertices
459 are set so that the s-axis of the texture space is parallel
460 to the y-axis of the object space — axis along the
461 height of the cylinder — and the t-axis is parallel to the
462 x-axis. With this, the text is repeated several times along
463 s depending on the height of the object and, around the
464 cylinder, it was agreed that the text should be repeated
465 four times, so it could be easily visualized no matter
466 which side of the cylinder the user is looking at. The
467 calculation for finding scales, which must be applied to
468 the texture coordinates, and depends on the cylinder's
469 size and the number of characters of the labels, is all
470 based on the calculations presented for quads.

471 In the case of spheres, text is displayed only once over
472 t, on the object's center, and it can be repeated several
473 times along s depending on the sphere's diameter. The
474 primitive called dish, as found in CAD models, has the
475 shape of a half-ellipsoid and, with its information of radius
476 and height, it is possible to parameterize its surface.
477 Fig. ?? shows boxes, cylinders, spheres and dishes with
478 labels.

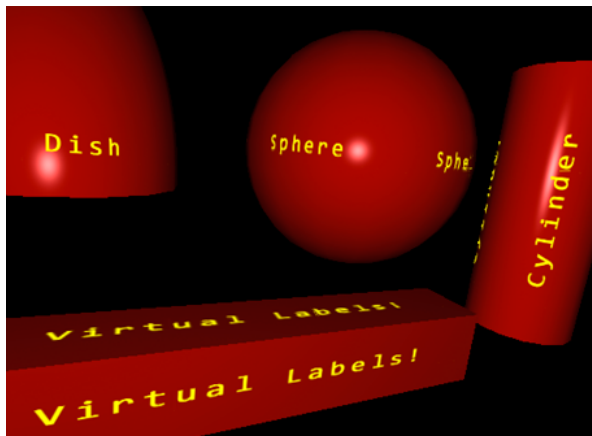


Figure 5: CAD objects with labels.

479 The positioning of labels on objects like cones and
480 pyramids could be a bit more complex due to the different
481 dimensions of the bases of those objects. With
482 these, the texture coordinates must be calculated in a
483 way to compensate that difference, so that the text won't
484 be stretched proportionally to the size increase of one
485 base in relation to the other. Despite that, the calculation
486 has not been done until the present moment, and,
487 with the technique explained in this work, the positioning
488 of labels is possible with more primitives than in the
489 related works [3] — which allow the labels to be displayed
490 on cylinders and boxes only.

491 5. Results

492 The results are going to be presented in two parts. Initially,
493 we are going to show the result of the technique applied to CAD
494 models. Then, we are going to present performance tests to show
495 the effectiveness of the proposed technique.

497 5.1. Labels in CAD Models

498 The display of text with the technique presented in this work
499 can be used for different purposes, but the motivation to create
500 it was to display labels on massive model objects and more
501 specifically, on CAD models.

502 For the rendering of CAD models we built an engine that uses
503 a primitive instantiation technique, explained in the previous
504 section. The engine renders all CAD model objects, placing
505 labels on the previously mentioned objects: cylinders, boxes,
506 spheres and dishes.

507 The models used are based in files, which contain information
508 for each object to be drawn (position, orientation, scale and
509 parameters specific to each primitive). Those object parameters
510 are sent to the GPU as texture coordinates, since that was the
511 fastest way found, considering the limitations of OpenGL 2.1 and
512 GLSL 1.2. Besides that instantiation technique, the engine uses
513 no other optimizations, but, with the built engine we can render
514 a large number of objects at interactive rates, as will be shown
515 in the next section.

517 Fig. 6 shows labels applied to the objects of an oil refinery
518 CAD model. As it can be noticed, it's possible to display labels
519 on a large number of CAD model objects and, also, that most of
520 the geometries that exist in that model are boxes or cylinders.



Figure 6: CAD model of an oil refinery with labels.

522 The character atlas can have several variations to facilitate
523 the reading of labels. For instance, in Fig. 6,

white letters with black borders were used, which can be good to contrast the labels with different background colors. It is also possible to use an atlas with a black color for the characters, making it simple to modify that color within the fragment shader of the second rendering pass when the texture is sampled. Fig. 7 shows colored labels on a CAD model, from an atlas with black characters with an alpha component, whose color was changed in the GPU.

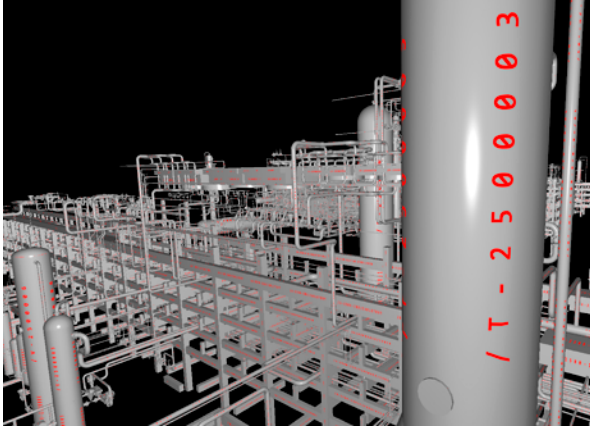


Figure 7: CAD model with colored labels.

5.2. Performance Tests

The technique presented in this work, for the display of labels on geometric surfaces, uses two rendering passes, and the first adds little complexity to the CAD graphics engine — in the first pass, calculations are added only to generate the texture coordinates and for the fragment shader to write information on two buffers instead of one. In the second pass several per pixel operations are executed and also some accesses to different textures. Therefore, it's important to evaluate the performance of the proposed technique, especially when the application utilizes higher resolutions. For that, tests were done and are presented below.

For the first tests, a scene with quads with random dimensions and positions was created, with texture coordinates generated in the simplest way, so that labels occupy the entire area of the objects and, therefore, calculations related to labeling were done for all the quads' pixels. The object's labels had different characters and sizes. The same scene was used to compare the performance of the versions with and without labels, whereas in the first case the algorithm does just the first pass and doesn't perform any calculation related to labeling, not even the generation of texture coordinates for the ge-

ometries. One of the created scenes, with 100,000 objects and resolution of 1920x1200 is shown in Fig. 8. It is worth mentioning that labels have whitespace characters before and after the text, and that the entire surface of each quad has a label.



Figure 8: Test scene with 100,000 quads. Version with labels (1920x1200).

For all tests we used a computer with the operational system Windows 7 64 bits, an Intel Core I7 870 processor with 2.93 GHz, 8GB of RAM and a GeForce GTX 580 graphics card.

Table 1 presents the comparison of rates in frames per second (FPS) obtained for a fixed camera position, for scenes with different number of objects, at a resolution of 1920x1200. One can note that the difference of FPS is small in scenes with fewer objects and virtually nonexistent in scenes with more objects.

Table 1: Comparative performance tests for a fixed camera position without and with the use of labels (resolution: 1920X1200).

Number of Objects	FPS (without labels)	FPS (with labels)
10,000	457	439
50,000	106	105
100,000	54	54
200,000	27	27
500,000	11	11
1,000,000	6	6

Table 2 shows the result of another random scene of quads, but this time we measured the average frame rate in FPS obtained when navigating throughout the scene for 50 seconds. The exact same camera path was used for the scenes with and without labels. Tests were performed in two resolutions: 600x400 and 1920x1200 and it is once again shown that the performance difference tends to be null for a scene with a larger number of

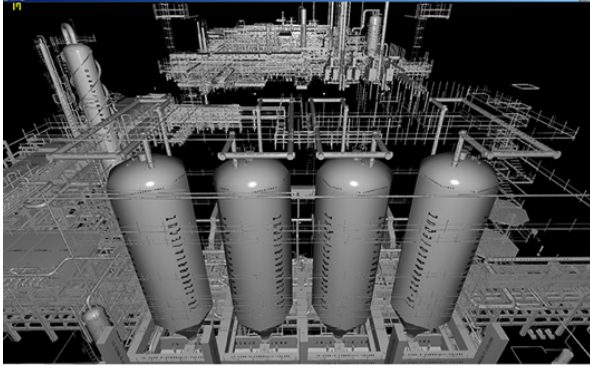


Figure 9: Test scene with CAD model.

objects. In scenes with fewer objects, the difference reaches 5% but the rates are still very high. There have been some cases with the 600x400 resolution in which the resultant average was slightly higher when using labels. Such a difference is insignificant and can be explained by the interference of events external to the program, like processes of the operating system.

Table 2: Comparative performance tests for camera trajectory without and with the use of labels.

Number of Objects	Resolution	FPS (without labels)	FPS (with labels)
10,000	600 x 400	511.6	485.3
	1920 x 1200	488.1	470.1
50,000	600 x 400	106.7	107.9
	1920 x 1200	102.0	99.9
500,000	600 x 400	10.6	10.6
	1920 x 1200	10.6	10.5
1,000,000	600 x 400	4.5	5.3
	1920 x 1200	5.4	5.4
1,200,000	600 x 400	4.6	4.5
	1920 x 1200	4.5	4.5

The last tests were performed with a CAD model of a real oil refinery, rendering all the existing primitives in the model, except for triangle meshes. Cylinders, boxes, spheres and dishes exhibit labels and the total number of rendered objects was 300,000, mostly boxes and cylinders, as shown in Fig. 9. The texture coordinates of the objects were defined as explained in section IV, being generated automatically depending on the size of the objects and the number of characters of the labels. For a fixed camera position, shown in Fig. 9 we obtained 17 FPS for a scene with and without labels. The results of the tests when navigating the scene with the same camera path for a scene with and without labels can be found in Table 3, which shows that the performance loss was little or none with the use of labels.

Table 3: Comparative performance tests for camera trajectory without and with the use of labels in a CAD model.

Number of Objects	Resolution	FPS (without labels)	FPS (with labels)
300,000	600 x 400	15.8	15.5
	1920 x 1200	16.1	16.1

6. Conclusion

In this paper we developed an algorithm entirely in the GPU to display, in real time, labels on the surfaces of geometric primitives of massive models, and found that it presents good performance and visual results. One of the challenges identified was that, when dealing with massive models, each of its objects could have a different textual content needing distinct labels. Therefore, if we used a different texture for each label, multiple texture context switches would be needed by frame, which would certainly severely compromise the performance of the application. Likewise, the use of one texture per object would imply a high memory cost and the development of strategies to maintain, in video memory, a limited number of textures per frame.

With the algorithm developed in this paper, only the labels that need to be shown on screen are built per frame, in the GPU, and the label mapping is done for each fragment by sampling an atlas with the graphical representation of the characters and a buffer with the ASCII codes that contain textual information of the labels. Therefore, we don't need to deal with problems related to texture context switches and also the problem of building and storing textures with the images of each label in video memory before the objects are rendered. We need only to store the atlas and the buffer with ASCII codes. The storage of ASCII codes uses much less memory than storing the images of labels.

The problem of positioning labels on objects of CAD models with different sizes and shapes was solved for cylinders, boxes, dishes and spheres. It was identified that, for the user to be able to see the labels independently of how the camera points to an object, text repetition was a good choice.

As a future work we want to add labels to other kinds of objects from CAD models, such as cones and pyramids, and also to improve the visual quality of labels with the use of anti-aliasing techniques. We intend to study the use of dynamic labeling, to display operational data such as temperature and pressure directly on the objects' surfaces. We also want to explore the use of labels in domains other than CAD, once the technique presented here for the generation of virtual labels is not restricted to the CAD domain.

646 References

- 647 [1] S. Yoon, E. Gobbetti, D. Kasik, D. Manocha, Real-time Massive
648 Model Rendering, Vol. 2 of Synthesis Lectures on Computer
649 Graphics and Animation, Morgan and Claypool, 2008.
- 650 [2] NVIDIA Corporation, Improve batching using texture atlases,
651 [https://developer.nvidia.com/sites/default/files/](https://developer.nvidia.com/sites/default/files/akamai/tools/files/TextureAtlasWhitepaper.pdf)
652 [akamai/tools/files/TextureAtlasWhitepaper.pdf](https://developer.nvidia.com/sites/default/files/akamai/tools/files/TextureAtlasWhitepaper.pdf)
653 (2004).
- 654 [3] Aveva Review Home Page, <http://www.aveva.com> (2010).
- 655 [4] Z. Qin, M. D. McCool, C. S. Kaplan, Real-time texture-mapped
656 vector glyphs, in: Proceedings of the 2006 symposium on Inter-
657 active 3D graphics and games, I3D '06, 2006, pp. 125–132.
- 658 [5] S. Lefebvre, S. Hornus, F. Neyret, Texture sprites: Texture ele-
659 ments splatted on surfaces (April 2005).
- 660 [6] G. Cipriano, M. Gleicher, Text scaffolds for effective surface
661 labeling, IEEE Transactions on Visualization and Computer
662 Graphics 14 (6) (2008) 1675–1682.
- 663 [7] S. Barret, Sparse virtual textures,
664 <http://silverspaceship.com/src/svt> (2008).
- 665 [8] M. Mittring, Advanced virtual texture topics, in: ACM SIG-
666 GRAPH 2008 Games, SIGGRAPH '08, 2008, pp. 23–51.
- 667 [9] OpenGL 2.1 Reference Pages,
668 <http://www.opengl.org/sdk/docs/man2/> (2006).
- 669 [10] The OpenGL Shading Language 1.2,
670 [http://www.opengl.org/registry/doc/GLSLangSpec.](http://www.opengl.org/registry/doc/GLSLangSpec.Full.1.20.8.pdf)
671 [Full.1.20.8.pdf](http://www.opengl.org/registry/doc/GLSLangSpec.Full.1.20.8.pdf) (2006).