

# Graphics

in the game

# Third Degree

By: Chad Williams

California Polytechnic State University

Advisor: Zoë Wood

Winter 2011 - Spring 2011

# THIRD DEGREE

## TABLE OF CONTENTS

Introduction.....	4
Motivation .....	4
Team and Course Structure .....	4
Project Overview .....	4
Story.....	4
Genre & Setting .....	4
Game Mechanics .....	5
Look & Feel .....	5
Technical Specifics .....	5
Project Technologies .....	5
Related Works .....	6
Graphics.....	8
Shaders .....	8
Deferred Rendering .....	8
Glow Shader.....	11
Normal Mapping Shader.....	13
Respawn Effect Shader .....	16
Performance Enhancements .....	18
Vertex Buffer Objects .....	18
Texture and Mesh Instancing .....	19
Particle System Implementation .....	19
Overview.....	19









## Trine

The general mood/feel of the game was greatly influenced by this side-scrolling platformer, Frozenbyte's *Trine*. Visual inspirations, the overall feel of gameplay mechanics (such as movement and puzzle object interaction), and elements of the combat system helped in making decisions for *Third Degree*. The Figure below shows an in-game screenshot for *Trine*:



FIGURE 1 - FROZENBYTE'S TRINE

## Maya

The transformation tools were modeled after many 3D graphics software suites, particularly Autodesk Maya. The figure below shows an example of the transformation tools used in Maya:




FIGURE 2 - TRANSFORMATION TOOLS IN AUTODESK'S MAYA

### Doom 3

The animation in the game utilizes the MD5 format used in a number of 3D games, most notably in Activision's *Doom 3*. The MD5 structure created for *Doom 3* provides a robust and efficient way for representing animation. The figure below shows an example mesh from *Doom 3* modeled using MD5:



FIGURE 3 - EXAMPLE MD5 MESH FROM DOOM

## GRAPHICS

*Third Degree* relies heavily upon graphics to convey story and emotion during gameplay. There were many advanced graphical effects that the game could not employ due to target hardware which was not powerful enough. However, despite this, many effects were used that enabled the game to immerse the player in a very rich and detailed environment. Shaders are at the core of this, allowing for advanced lighting, post-processing, and offloading complex tasks that are inherently parallel from the CPU to the GPU.

### SHADERS

The game was built using OpenGL rather than DirectX, so the two shader languages available were GLSL (OpenGL Shader Language), or Cg. For this project, GLSL was chosen because it is OpenGL's official shader language and is being actively developed alongside the OpenGL specification.

### DEFERRED RENDERING

From the very beginning, *Third Degree* had a very ambitious visual style, not only trying to convey a Victorian-esque environment, but also a very futuristic environment as well. During the first quarter in development, it was clear that standard forward lighting would not be enough to convey the visual style that the game was targeting.



Deferred rendering, which allows for lighting to be done in its own pass, allows for an *essentially unlimited*<sup>1</sup> number of lights in the scene, allowing for a massive number of creative options to become available in the development process. *Third Degree* owes credit of its core deferred rendering engine to Ryan Schmitt, who very graciously allowed us to implement his deferred rendering code. Without the assistance Ryan provided, the advanced lighting that *Third Degree* employs would certainly not have been attainable within the two-quarter time frame that the game was developed in.

In order to focus the player's attention on the objects that are switching between historical and futuristic equivalents, the ambient light was set very low, and an abundance of point lights were used to highlight and showcase the switchable objects. Below is an exploration of the way deferred rendering works (i.e., how Ryan's core system works, and how many other implementations function), and later in the report, how this core system was modified to allow for advanced effects in *Third Degree*.

## HOW DEFERRED RENDERING WORKS

---

As the name implies, deferred rendering delays, or defers, the rendering/lighting of the scene until it is presented to the screen. What this means, essentially, is that a shader is "gathering" (or capturing) all the data resulting from OpenGL draw calls and placing this data in buffers. The three main buffers used in the core deferred rendering system are an albedo buffer, a depth buffer, and a normal buffer.

Note that because the only color information stored in these three buffers is in the albedo buffer, no lighting information is present (albedo is *just* color without any lighting information). Because of this, lighting can be deferred until all three buffers are completely filled with the draw information for a complete frame. Collectively these buffers are called a G-buffer, short for "geometry buffer." With the information in these buffers, ambient, directional, and point lighting (as well as other lighting models) can be calculated. Because the lighting is done after the entire scene has been rendered to buffers, lighting is calculated per-pixel in view space. This has the distinct advantage of only lighting pixels that are visible after the scene has been rendered. In the case of the forward rendering model, having a large number of lights in the scene becomes very expensive, because for each light in the scene, OpenGL has to compute the light for each object. **Below is a summary of the complexity of deferred and forward rendering models:**

**Forward rendering:**  $O(M * L)$

*...where M and L represent the number of meshes and lights present in the frame, respectively*

**Deferred rendering:**  $O(M + L)$

*The following three figures are a snapshot of the three core buffers that make deferred rendering possible:*

---

<sup>1</sup> "Essentially unlimited" means that almost any reasonable number of lights in the scene would have perfectly decent performance.



FIGURE 4—DEFERRED RENDERING DEPTH BUFFER (NON-LINEARIZED)



FIGURE 5—DEFERRED RENDERING ALBEDO BUFFER




FIGURE 6—DEFERRED RENDERING NORMAL BUFFER (NON-NORMAL-MAPPED)

A clear disadvantage of deferred rendering is the large memory footprint required, since at least three buffers are required to perform lighting (albedo, depth, and normals). If additional effects are desired, even more buffers will be used, such as velocity, roughness, and specular buffers. Another disadvantage occurs when only a couple of lights are in the scene – in this case deferred rendering is most likely more complex than forward lighting. In *Third Degree*, however, each frame contains a large number of lights, so this disadvantage does not apply. Other important disadvantages include difficulty rendering transparent objects in the scene (like glass), and anti-aliasing problems. There are many techniques for anti-aliasing in the deferred rendering pipeline, but none of them rely on traditional anti-aliasing provided by the hardware.

## GLOW SHADER

In order to completely differentiate futuristic objects from historical objects, glow was applied to most future objects. In order to support glow, a fourth MRT (multiple render target) was added to the deferred rendering system. This fourth buffer contains color information which is then fed into the post-processing shader described below. In a nutshell, objects that support glow have an associated glow mask (Figure 7), the gather shader then masks off the albedo buffer with this glow mask, and places this masked albedo into the fourth MRT, which is the glow buffer (Figure 8).

## HOW THE GLOW ALGORITHM WORKS

The first step, as outlined above, is to write the albedo modulated by the glow mask into the glow buffer which will subsequently be used in the glow shader. This step is done in the gather shader which is part of the deferred rendering process:

```
gl_FragData[2] = texture2D(albedoTexture, Texcoord) * texture2D(glowTexture, Texcoord);
```

The above code simply samples the glow mask (`glowTexture`) and albedo buffer (`albedoTexture`), multiplies the samples together, and puts it into the fourth MRT. This works because the glow mask is either white or black, so the glow buffer will be populated with the albedo color when the glow mask is white (represented in color as `RGBA = 1.0, 1.0, 1.0, 1.0`) and black when the glow mask is black. Below are samples of an individual glow mask (Figure 7), the glow MRT (Figure 8), and the final rendered image (Figure 9):



FIGURE 7 - GLOW MASK USED FOR THE FUTURE LAMP PROP




FIGURE 8 - GLOW MRT



FIGURE 9 - FINAL IMAGE WITH GLOW

Once the glow MRT has been populated, the post-processing shader is run before the final image is presented to the screen. Below is the main algorithm used in the glow portion of the post-processing shader:

```
for(i = 0; i < KERNEL_SIZE; i++)
{
    texcoordOffset.x = (mod(i, 3)) * step_w - step_w;
    texcoordOffset.y = (i / 3) * step_w - step_w;

    vec4 sample = texture2D(GlowMap, gl_TexCoord[0].st + texcoordOffset);
    sum += sample * ONE_NINTH;
}
```

The fundamental idea behind this is the concept of a *kernel*. A kernel forms a matrix to tell the filter being applied to the image how to multiply, or weight, the pixels around the pixel being sampled. In this case, the kernel is a 9x9 matrix that has an equal weighting for each neighbor. In other words, the kernel is a 9x9 matrix of all ones. Because the matrix is all ones, the end result is an *averaging filter* which results in a blur. Other image filters are possible with kernels that consist of different weights, such as an edge-detection filter.

Referencing the above code, the `KERNEL_SIZE` is nine, as stated before. The glow map (which is the glow MRT) is sampled and multiplied by 1/9 (because there are 9 pixels in the kernel), and added to the sum. The final result stored in `sum` is the average of the 9 pixels being sampled in the glow MRT. This blurred glow is then overlaid on top of the image that resulted from the deferred rendering process, as seen in the below code:

```
gl_FragColor = sum * 4.0 + NewColor;
```

This is the final step in the post-process shader, which combines `sum`, multiplies it by 4 (in order to intensify the glow), and adds in `NewColor`, which is the result of any other post-processing (and will contain some form of the final deferred render).

## NORMAL MAPPING SHADER

The process of normal mapping is fairly straight-forward in a non-deferred rendering setting. Since *Third Degree* doesn't use forward rendering however, the transforms involved were more complex. Below is a diagram of the transform that is required to take *tangent space* normal maps and convert them into *view space* normals which are used in the normal MRT:

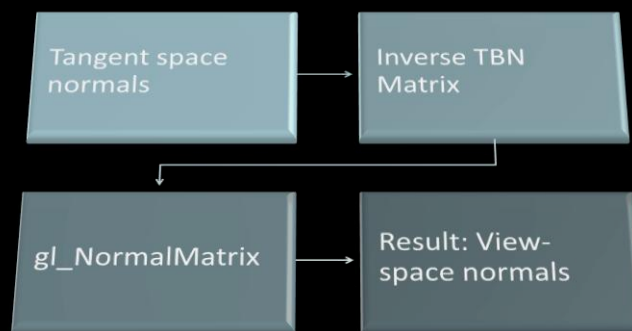


FIGURE 10 - NORMAL TRANSFORMATION PROCESS, FROM TANGENT SPACE TO VIEW SPACE








FIGURE 11 - A TANGENT-SPACE NORMAL MAP USED ON THE CRATE PROP




FIGURE 12 - VIEW-SPACE NORMALS (CONTAINED IN THE NORMAL MRT)



FIGURE 13 - FINAL LIT IMAGE WITH NORMAL MAPPING

## RESPAWN EFFECT SHADER

When the player in *Third Degree* dies (and respawns at the last checkpoint), or is spawned for the first time into the world, an overwhelming “brightness” occurs which is called the “respawn effect.” The effect is composed of two parts. The first part processes the final image to create a sepia effect; the second part of the effect involves significantly brightening up the image.

```
OldColor = texture2D(DeferredScene, gl_TexCoord[0].st);

NewColor.r = OldColor.r * 0.393 + OldColor.g * 0.769 + OldColor.b * 0.189;
NewColor.g = OldColor.r * 0.349 + OldColor.g * 0.686 + OldColor.b * 0.168;
NewColor.b = OldColor.r * 0.272 + OldColor.g * 0.534 + OldColor.b * 0.131;
NewColor.a = 1.0;

Intensity = (20 - 19 * FadeLevel);

NewColor = Intensity * mix(NewColor, OldColor, FadeLevel);
```

The first step in the sepia effect is sampling the final image generated by the deferred rendering, `DeferredScene`. The individual R, G, and B channels are then modulated by specific values to generate a sepia effect (see Figure 14 below). The color generated by the sepia effect, stored in `NewColor`, is then modulated by an intensity value to brighten up the image. The values selected in the intensity function are somewhat arbitrary and were derived simply on a trial-and-error basis. `FadeLevel` is a value that varies from 0.0 to 1.0, so the end result is that the image will be brightened by 20x when the `FadeLevel` is 0, and 1x (or no added brightness) when the `FadeLevel` is 1. After the intensity is generated, `NewColor` and `OldColor` are simply linearly interpolated via `mix` by the `FadeLevel`, and multiplied by the generated intensity.





FIGURE 14 - THE FIRST STAGE IN THE RESPAWN EFFECT (SEPIA)



FIGURE 15 - THE RESPAWN EFFECT MIDWAY THROUGH COMPLETION. THE FINAL IMAGE IS GENERATED BY MULTIPLYING THE SEPIA FILTER BY THE INTENSITY VALUE.








FIGURE 16- CLASS DIAGRAM OF PARTICLE SYSTEM IMPLEMENTATION

## CONFIGURATION

There are many ways that the particle system can be configured to achieve the desired effect or look. Any of the system properties can be changed during gameplay to create a dynamically-changing effect. Although the particle system can be configured in many ways, there is still room for improvement. The main drawback for the approach that was taken was due to the discrete nature of each particle. In the future, creating a more procedural system that allows for both discrete particles, as well as effects that allow for cohesion of particles, would be worth looking into.

Figure 17 describes how the particle system can be configured, including emitter and affector properties. In addition, rather than hard-coding the particle system properties and configurations, the main map file (which contains everything in the scene) simply specifies positional properties and a “script” for each particle system. The map loader then examines the specified particle script and sets the system properties, and adds emitters and affectors.

It is also worth noting that, again, particles are stored in VBOs to speed up performance, and all particle manipulation is done on the GPU (such as physics calculations) to speed up rendering. The only process carried out on the CPU is particle generation, and in a future revision of the particle system, geometry shaders will be used to generate particles on the GPU, making the system nearly completely reside on the GPU.

### Particle System Instance

- Position
- Scale

### Particle Emitter

- Type (Point, Box, Circle)
- Enabled

### Color Fade Affector

- Target Color
- Fade Type (Linear, Quadratic, Smooth Step)
- Inverted Fade

### Gravity Affector

- Gravity Strength

- Particles per Second
- Burst Enabled
- Burst Amount
- Burst Looping Enabled
- Start Time
- Particle Start Color
- Minimum & Maximum Scale
- Minimum & Maximum Particle Lifetime
- Minimum & Maximum Initial Velocity

**Swirl Affector**

- Swirl Speed
- Swirl Plane (XY, YZ, XZ)

**Move Affector**

- Enable Move X, Y, and Z

**Omni Affector**

- Omni Velocity

FIGURE 17 - PARTICLE SYSTEM PROPERTIES

EXAMPLES



FIGURE 18 - "SURREAL WHITE" PARTICLE EFFECT



FIGURE 19 - "FIREFLIES" PARTICLE EFFECT




FIGURE 20 - "SHOCKWAVE" FOCUS PARTICLE EFFECT



# MAP EDITOR

*Third Degree* takes place in a world that is composed of a very large number of meshes, lights, “triggers”, and other game support elements, so creating a map editor specifically for the game was deemed critical before the game even began development. Because I have prior experience developing game-related editors, I took up the task of creating the editor, using Nokia’s Qt for the user interface, and *Third Degree*’s engine for the 3D rendering. Credit goes to Joshua Marcelo for object movement in the editor, and Jon Moorman for modifications to the editor’s save and load functionality.

Because this write-up is chiefly concerned with the graphical elements of *Third Degree*, only a brief overview of its functionality will be given. The majority of the editor’s functionality is simply achieved by giving commands to *Third Degree*’s rendering engine via the Qt interface.




FIGURE 21—OVERVIEW OF THE MAP EDITOR

As Figure 21 demonstrates, there are four main areas to the editor. The first, and arguably most important area, is the preview window in the middle of the editor. This window gives a real-time preview of what the game will look like, but *without* any post-processing effects (to keep rendering as fast as possible). In addition, special cues, such as particle system selection handles, object manipulators, spawn points, and other editor-only helpers are visible.

To the left of the preview area is the “Scene Manager” which contains the “Scene Hierarchy,” a list of all objects in the scene, and the “Object Browser,” a list of meshes, particle systems, and puzzle objects that are available to be inserted into the scene. The Scene Manager makes it very easy to insert any available item into the scene quickly and without having to go through any other windows. In a previous iteration of the editor, feedback obtained indicated that having to go through file browsing dialogs slowed down map building considerably, which is why the Object Browser was put into the editor. The Scene Manager also contains a separate panel near the bottom that

allows the user to specify translation, scale, and rotation per-object, object-switching properties, light properties, and more.

At the bottom of the editor, the console gives information back to the user regarding actions that have been completed, errors that occurred, and any warnings that may indicate a potential problem has arisen. The console was not in the first iteration of the editor, and has proven invaluable in debugging and giving feedback to the artists as to what is happening in the scene.

Lastly, the toolbar (at the top of the editor) provides quick access to actions that the artist may frequently utilize, such as toggling the map between 100% future and 100% historical preview, manipulating the ambient light settings for easy editing, duplicating an object to the left or right (good for tiling road objects), etc. The menu bar (not to be confused with the toolbar that has large buttons) exposes all the actions available in the editor (showing available hotkeys), including full undo/redo.

## RESULTS

The end result of *Third Degree*'s two-quarter development cycle was a solid proof-of-concept that many team members plan on taking further and possibly selling on Steam. The original premise of the game was extremely ambitious, and *Third Degree* accomplished a lot, including excellent graphics, gameplay prototyping, a strong backend engine (including support for physics), a developed story, and a robust editor. A screenshot from the final game is shown below:




FIGURE 22 - THE FINISHED GAME PROTOTYPE





## CREDITS

*Third Degree* relied on a large amount of outside help to accomplish its ambitious goals. The following people contributed their time to the project in the specified areas:

<b>Josh Holland</b>	Art lead, 2D artwork
<b>Ben Funderberg</b>	3D modeling
<b>Tom Funderberg</b>	3D modeling
<b>Hector Zhu</b>	Splash screen, game modes
<b>Mikkel Sandberg</b>	3D modeling
<b>Mitch Epeneter</b>	Voice acting
<b>Ryan Schmitt</b>	Core deferred rendering system

## LIST OF FIGURES

Figure 1 - Frozenbyte's Trine .....	7
Figure 2 - Transformation tools in Autodesk's Maya .....	7
Figure 3 - Example MD5 mesh from Doom .....	8
Figure 4 – Deferred rendering depth buffer (non-linearized) .....	10
Figure 5 – Deferred rendering albedo buffer .....	10
Figure 6 – Deferred rendering normal buffer (non-normal-mapped) .....	11
Figure 7 - Glow mask used for the future lamp prop .....	12
Figure 8 - Glow MRT .....	12
Figure 9 - Final image with glow .....	12
Figure 10 - Normal transformation process, from tangent space to view space .....	13
Figure 11 - A tangent-space normal map used on the crate prop .....	15
Figure 12 - View-space normals (contained in the normal MRT) .....	15
Figure 13 - Final lit image with normal mapping .....	16
Figure 14 - The first stage in the respawn effect (sepia) .....	17
Figure 15 - The respawn effect midway through completion. The final image is generated by multiplying the sepia filter by the intensity value. ....	17
Figure 16- Class diagram of particle system implementation .....	20
Figure 17 - Particle system properties .....	21
Figure 18 - "Surreal white" particle effect .....	21
Figure 19 - "Fireflies" particle effect .....	21
Figure 20 - "Shockwave" focus particle effect .....	21
Figure 21 – Overview of the map editor .....	22
Figure 22 - The finished game prototype .....	23

## REFERENCES

- Guinot, J. (2006, December 30). *Bump Mapping using GLSL*. Retrieved from oZone3D:  
[http://www.ozone3d.net/tutorials/bump\\_mapping.php](http://www.ozone3d.net/tutorials/bump_mapping.php)
- Guinot, J., & Riccio, C. (2007, January 7). *OpenGL Vertex Buffer Objects*. Retrieved from oZone3D:  
[http://www.ozone3d.net/tutorials/opengl\\_vbo.php](http://www.ozone3d.net/tutorials/opengl_vbo.php)

