libgraphics: Design and Implementation

Rodrigo G. López
rgl@antares-labs.eu

ABSTRACT

Libgraphics is a 3D computer graphics library that provides a way to set up a scene, fill it up with a bunch of models (with their own meshes and materials), lights and cameras, and start taking pictures at the user request. It implements a fully concurrent retained mode software renderer, with support for vertex and fragment/pixel shaders written in C (not GPU ones, at least for now), and featuring a z-buffer, front- and back-face culling, textures and skyboxes, directional and punctual lights, tangent-space normal mapping, among other things.

Introduction

Write the intro last.

1. The scene

A scene is a container, represented as a graph, that hosts the entities that make up the world. Each of these entities has a model made out of a series of meshes, which in turn are made out of geometric primitives (only *points*, *lines* and *triangles* are supported). Each model also stores a list of materials.

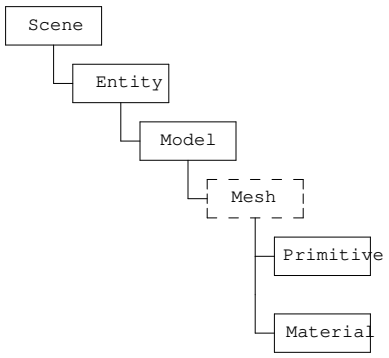

Figure 1 : The scene graph.

1.1. Entities

1.2. Models

1.3. Meshes

1.4. Primitives

1.5. Materials

2. Cameras

3. The renderer

The *renderer* is the core of the library. It follows a **retained mode** model, which means that the user won't get a picture until the entire scene has been rendered. Thanks to this we can also clear and swap the framebuffers without their intervention, they only need to concern themselves with shooting geometry and developing a camera.

It's implemented as a tree of concurrent processes connected by Channels as seen in **Figure 2**, spawned with a call to `initgraphics`, each representing a stage of the pipeline:

- The **renderer** process, the root of the tree, waits on a channel for a `Renderjob` sent by another user process, specifying a framebuffer, a scene, a camera and a shader table. It walks the scene and sends each `Entity` individually to the entityproc.
- The **entityproc** receives an entity and splits its geometry equitatively among the tilers, sending a batch for each of them to process.
- Next, each **tiler** gets to work on their subset of the geometry (potentially in parallel, see **Figure 3**). They walk the list of primitives, then for each of them apply the **vertex shader** to its vertices (which expects clip space coordinates in return), perform frustum culling and clipping, back-face culling, and then project them into the viewport (screen space). Following this step, they build a bounding box, used to allocate each primitive into a rasterization bucket, a **tile**, managed by one of the rasterizers, as illustrated in **Figure 4**. If it spans multiple tiles, it will be copied and sent to each of them.
- Finally, the **rasterizers** receive the primitive in screen space, slice it to fit their tile, and apply a rasterization routine based on its type. For each of the pixels, a **depth test** is performed, discarding fragments that are further away. Then a **fragment shader** is applied and the result written to the framebuffer after blending.

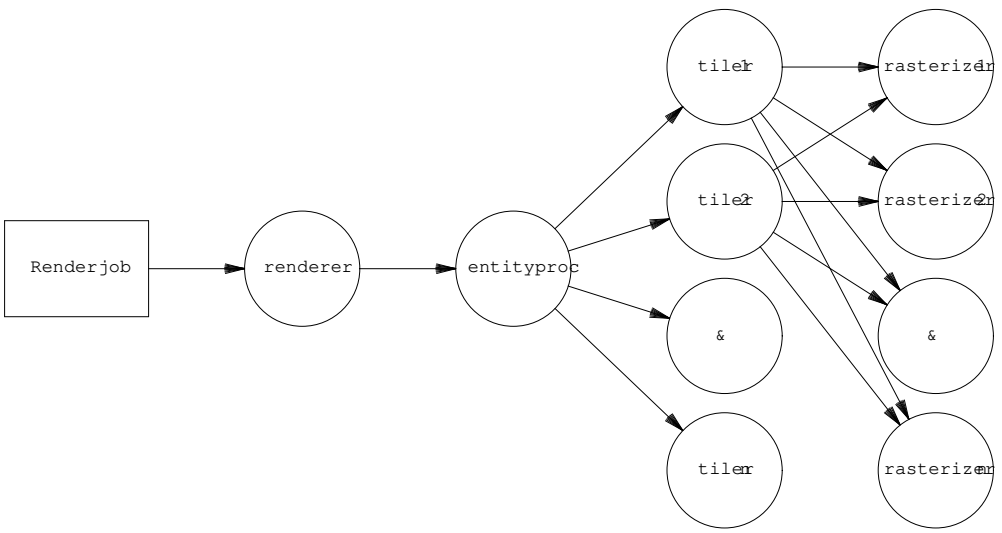

Figure 2 : The rendering graph for a 2n processor machine.

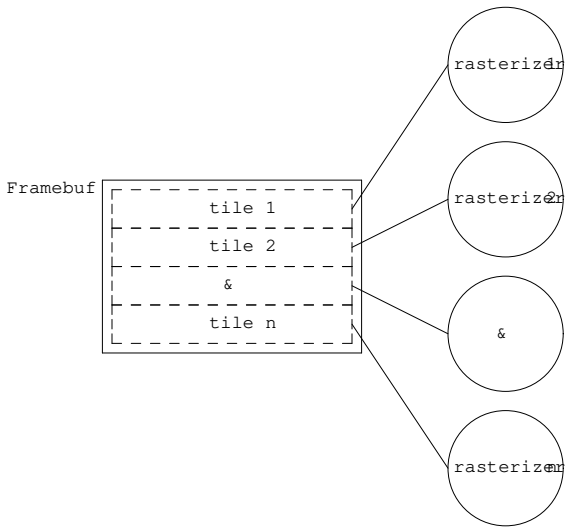

Figure 3 : Per tile rasterizers.

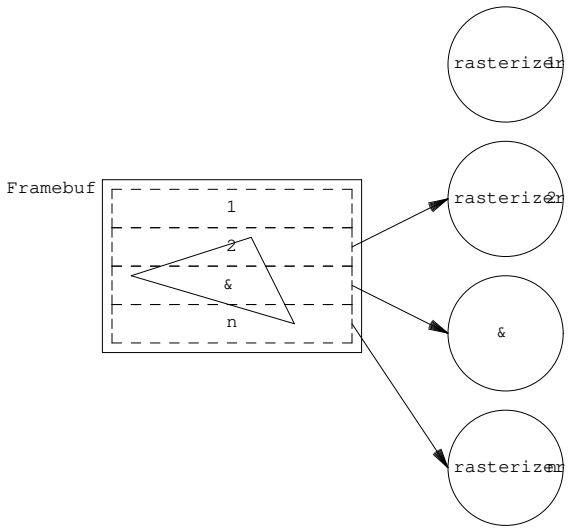

Figure 4 : Raster task scheduling.

4. Frames of reference

Frames are right-handed throughout every stage.

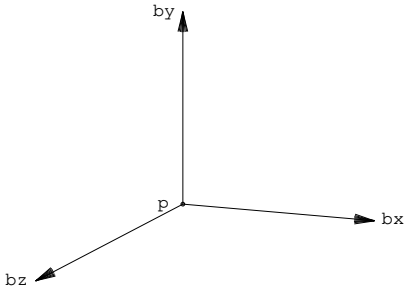

Figure 5 : Example right-handed rframe.

5. Viewports

A *viewport* is a sort of virtual framebuffer, a device that lets users configure the way they visualize a framebuffer, which changes the resulting *image* (6) after a call to its draw or memdraw methods. As of now, the only feature available is upscaling, which includes user-defined filters for specific ratios, such as the family of pixel art filters *Scale[234]x*, used for 2x2, 3x3 and 4x4 scaling respectively [?].

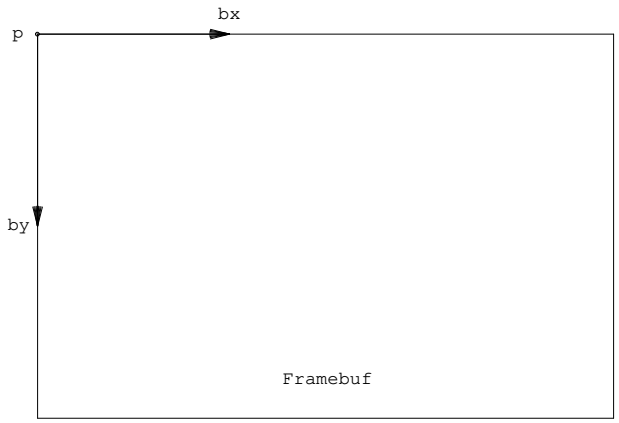

Figure 6 : Illustration of a 3:2 viewport.

References

[?] https://www.scale2x.it/