

Fragment shaders for agent animation using Finite State Machines

Isaac Rudomín, Erik Millán and Benjamín Hernandez

*Instituto Tecnológico y de Estudios Superiores Monterrey,
Campus Estado de México*

Abstract

In a previous paper we generated animated agents and their behavior using a combination of XML and Images. The behavior of agents was specified as a finite state machine (FSM) in XML. We used images to determine properties of the world that agents react to. While this is very flexible, it can be made much faster by using the power available in modern GPUs. In this paper we implement FSMs as fragment shaders using three kinds of images: world space images, agent space images and FSM table images. We show a simple example and compare performance of CPU and GPU implementations. Then we examine a more complex example involving more maps and two types of agents (predator-prey). Furthermore we explore how to render agents in 3D more efficiently by using a variation on pseudoinstancing.

Key words: virtual characters, GPU, shader, maps, crowds
PACS:

1 Introduction

In virtual environments and video games, it is common to find different animated characters that interact with the user, with other virtual characters, and with their surrounding environment. The behavior of these characters has been defined in many different ways: One first approach is emergent behavior (boids, by Reynolds [1]). A very common and simple way to specify behavior

Email addresses: rudomin@itesm.mx, emillan@itesm.mx, hbenjamin@itesm.mx (Isaac Rudomín, Erik Millán and Benjamín Hernandez).

URL: <http://rudomin.cem.itesm.mx/~rudomin> (Isaac Rudomín, Erik Millán and Benjamín Hernandez).

for virtual characters is by using finite state machines (FSMs). For instance, as done by Cremer [2], behavior is controlled through communicating, hierarchical and concurrent state machines. Work by Musse [3], has a hierarchical FSM architecture is used that models the behavior of crowds of characters. Devillers [4], develops a scenario language, similar to structured programming languages. This uses FSMs to specify most of the instructions. Becheiraz [5] includes more complex behaviors involving emotions in the FSMs.

In previous work [6], we created a system that specified behavior of by using XML scripting and images. While our approach proved useful for defining simple behaviors for semi-autonomous characters, we decided that, since it used images extensively, it could be useful, and deliver better performance, if we would implement agent behavior based on finite state machines as fragment shaders. Not much has happened in A.I. simulation in the GPU. There is some very interesting preliminary work for a class in UNC [7] that seems to take an approach similar to ours, based on Swarm Intelligence [8], but implemented in the GPU by using a modification of coupled map lattices [9] implemented as shaders. It is difficult to judge how far this effort has progressed.

2 FSM shaders

Modern GPUs allow using several textures and dependent texture lookups. Based on this idea, we implement a very simple system involving three kinds of maps:

- (1) world space maps
- (2) agent space maps
- (3) FSM maps

World space maps are what we normally consider maps. They codify some value for each location (pixel) in the map. This would cover many types of maps: texture maps, height maps, normal maps, collision maps, interest area maps, action maps. In general, all such maps can be considered maps labeled with some color on each pixel. The information represented by the colors in the maps, as well as the number of such maps that one could use in a particular simulation is up to the application designer.

Agent maps, on the other hand, have a pixel for each agent. They can be used to codify the state of the agent, the type of agent, the x,y position of the agent in the world map, the agent's velocity in said world map, or any other information of the agent.

A finite state machine can be represented as a table where given a certain state of the agent and a certain input, one can obtain a new state (and position) of the agent. This table can also be considered a map. Maps can be built for deterministic state machines. If there are several types of agents, there could be a portion of the table devoted to each type of agent. In a similar way one could describe probabilistic FSMs, hierarchical FSMs or layered FSMs. For the moment we consider a simple system.

- (1) A single labeled color texture represents the world where the agents will move. Suppose the image is size $X \times Y$ and there are L possible color labels.
- (2) Another texture represents our agents (with N agents represented as a RGB $N \times 1$ pixel image (R is state, G and B the xy position of the agent in the world). Each agent can have up to S different states.
- (3) Finally, there is a FSM image of size $S \times L$, indexed by state and input, as one can see in figure 1 a.

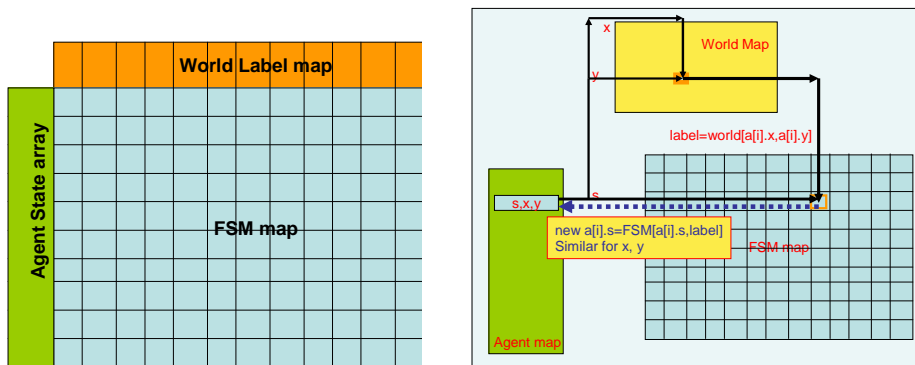


Fig. 1. a) FSM map indexed by state and input, b) Basic mechanism involving agent, world and FSM maps

All we must do then, is to consult the FSM map to update the agent map, and then render the world map with the agents overlaid. There is another step, however, since the FSM map is indexed by state, that can be obtained from the agent map, but also by input (or label color), that must be obtained from the agent's position in the world. Therefore, the basic mechanism, expressed as pseudocode, is extremely simple:

```

given agent i
  s=agent[i].s; x=agent[i].x; y=agent[i].y;
  l=world[x,y];
  agent[i].s=fsm[s,l];
  draw agent[i];

```

One would update the positions of the agent in a similar fashion. As you can see all we are doing are texture lookups. The basic mechanism can be seen in figure 1 b.

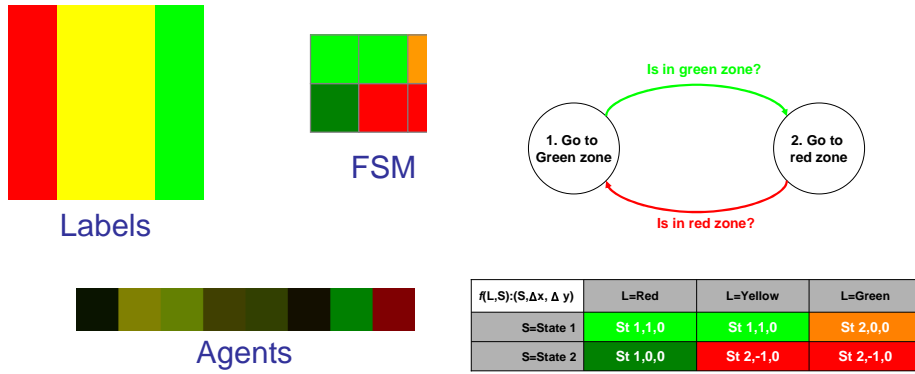


Fig. 2. a) Example agent, label and FSM maps b) Building the FSM map,

3 Results

The system has been implemented as fragment shaders using GLSL and runs on a PC with a 6800 Nvidia card. We will discuss two examples. A simple example that we will compare with a software implementation. Then we will discuss a more complex predator-prey example.

3.1 Simple example

The first example is extremely simple. Agents are generated at random positions. A very simple labeled world map is constructed with three vertical strips: an agent will consult the color of the labeled world map and follow a very simple FSM that causes it to move right until it reaches a the strip in the right, at which point the agent changes state and starts moving left until it gets to the strip in the left. The FSM map in this example was hand crafted from the table asociated with the diagram in figure figure 2 b. We have chosen the appropriate color for the table by using different R values for the states, different GB values for representing $\delta x, \delta y$. The resulting FSM map as well as the label and agent maps used for this example can be seen in figure 2 a.

Building the FSM map by hand seems simple enough for diagrams such as the one shown here, but one can see that it is probably not the best way to construct the maps for more complex FSMs. This is a problem that we will try to solve in the future.

To implement the basic mechanism in the case of this simple example is the following fragment shader in GLSL:

```
uniform sampler2DRect fsm;
uniform sampler2DRect agentsPos;
uniform sampler2DRect soilLabels;
```




Fig. 3. Frames from simple example

```

void main ()
{
  vec3 agentLookUp = texture2DRect (agentsPos, gl_TexCoord[0].st).rgb;
  vec2 agentPos = agentLookUp.rg;
  float agentState = agentLookUp.b;
  float soilLabelsLookUp = texture2DRect (soilLabels, agentPos).r;
  vec3 fsmOut = texture2DRect
    (fsm, vec2(soilLabelsLookUp.r, agentState)).rgb;
  gl_FragColor = vec4 (agentPos.x +
    (fsmOut.r-1), agentPos.y + fsmOut.g, fsmOut.b, 1);
}

```

Frames from a sample run of this example (for easier visualization a small number of agents was used) can be seen in figure 3. To do this, once we have applied the FSM using the fragment program shown above, we have to render these results so users could see the agents behavior. For testing purposes we have decided to render these agents as a points.

As we have mentioned, the FSM results are stored in a color texture or more precisely in a Pixel Buffer Object (PBO) for the GPU implementation. We take advantage of the fact that the PBO is resident in video memory so we can use the technique called "render to vertex array", were the color information of the PBO is directly converted to vertex position information using a shader program and Vertex Buffers Object (VBO).

For the CPU case, since the FSM results are stored in main memory, i.e. an array, we cannot use the same technique used in the GPU case. Nevertheless we still are using the Vertex Buffers Object. These VBOs are loaded with the array information so render performance is just affected by the intrinsic transfer of information from main memory to video memory. This makes the rendering in both implementations basically the same, and so simulation results we will see shortly, comparable.

If we run the simulation the simulation runs at interactive rates even with a large number of agents. We compare with a software FSM implementation. Frames per second resulting with different numbers of agents can be seen in the following table and in figure 4 (the same rendering algorithm is used in both cases).

It is clear from these numbers that even for larger numbers of agents (a mil-

agents	CPU	GPU
16384	428	713
65536	85	356
262144	11.8	137
1048576	3	33.9




Fig. 4. CPU vs GPU fps

lion), using the GPU can give us a significant advantage. Using the GPU stays within the interactive range above 30 fps while the same algorithm running entirely on the CPU is now down to 3 fps.

3.2 Predator-prey example

GPU agent animation is not restricted only to the generation of the simple character behaviors in our first example. More complex behaviors can be achieved through the use of static and dynamic maps that can modify the outcome of the FSM. One example to illustrate this point is a predator-prey simulation.

The general idea of the simulation of a predator is described in Figure 5. Figure 5.a, contains a general description of the FSM: the predator is wandering within its environment until it locates a nearby prey. When this happens, it starts following its prey until it escapes, and then goes on wandering.

This FSM is encoded in the map shown in Figure 5.b. Here, the red component is used, as in the first example, to encode the current state of the automata. However, rather than containing the actual motion for the agent, the FSM contains a reference for this motion in its green component. According to the current state, the agent will decide either to stay in its own location or to check the most appropriate direction of motion.

In the wandering state, there is no clear hint on where to locate a prey. In this case, the agent uses a couple of static maps, shown in Figure 5.c as *wander* maps. The agent selects a pixel from each map according to its current location. The values from these pixels are added to the current location of the agent.

The maps used for the *hunt* state are used in a similar way to those used in the wandering state, but they are calculated dynamically; after each simulation step, these maps should be recalculated to produce dynamic behavior.

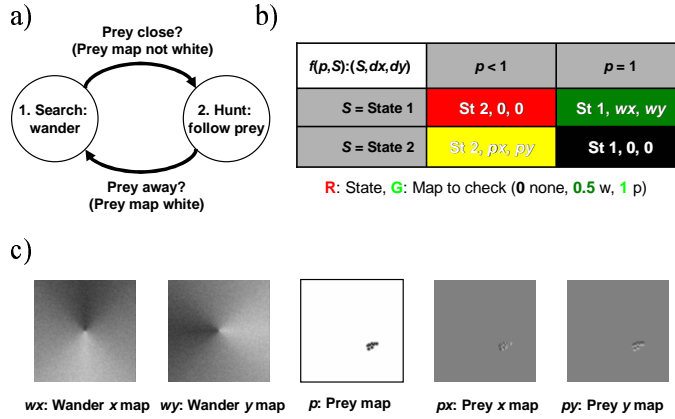


Fig. 5. Maps and FSM for predator-prey example

Initially, the location of all prey is marked in a collision map similar to the one used in [6], where each agent corresponds to one pixel according to its location in the map. A gaussian blur filter is applied to this collision map in order to obtain the prey map. In this way, a predator can locate only prey within its vision range. When a predator is located on a white pixel, no prey can be seen; but in a darker pixel, prey are seen and predators may follow them.

Another advantage of using the gaussian blur filter to obtain the visibility regions is that it is simple to infer the direction in which one would follow this prey. By applying horizontal and vertical sobel filters to the prey map we obtain the direction to follow prey, shown in Prey x and Prey y maps in Figure 5.c. These maps are interpreted just as wander maps are, simplifying the behavior implemented in the shader.

Back in the shader, transitions are selected according to the value of the prey map. If a predator is over a white area of the map, no prey is near, so the wander maps are used to obtain agent displacement. When prey are close, prey maps are used to select agent displacement. This modifies slightly the script from the simple example by querying a different map for the displacement instead of obtaining it from the FSM map, but the basic algorithm remains very similar and simple.

The results are similar to the ones in the simple example.

3.3 Optimizing 3D Render

If we render each 3D agent in the scene in the usual fashion, the frame rate decreases, but remains interactive for a reduced, but still significant, number of agents. Techniques such as pseudoinstancing allow us to increase this number back up significantly. This technique consists of

- (1) generating a vertex buffer with the agent positions
- (2) animating the vertex positions with a shader

To do this, then, we use the character position in the agent map to display a set of more complex virtual animated characters (instead of points as we did above). Our technique is similar to that used in [10]. First, we update the animation of a virtual character. While we used the same animated character for all agents, different animation instances could be used to increase diversity in character appearance.

The current frame of this animated character is sent to the graphics hardware as a set of vertex arrays and compiled into a display list. Then, for each agent instance, current and previous agent positions are sent as parameters for a vertex shader as color values. Using this vertex shader, the display list is rendered. This vertex shader uses the color value, which is the agent position, to translate the character to the position of the agent, and the difference between the current and the old agent position to rotate the character to follow its current motion direction. See the vertex shader next:

```
void main(void){
    vec2 delta = normalize(vec2(gl_Color.rg - gl_Color.ba));
    float cosine = delta.x;
    float sine   = delta.y;
    mat4 rot = mat4 (cosine, 0,  -sine, 0,
                    0,      1,   0, 0,
                    sine,  0, cosine, 0,
                    0,      0,   0, 1);
    vec4 pos = rot*(gl_Vertex);
    gl_Vertex.x = pos.x + gl_Color.r;
    gl_Vertex.z = pos.z - gl_Color.g;
    gl_TexCoord[0] = gl_MultiTexCoord0;
    gl_Position = ftransform();
}
```

The simple simulation discussed previously, with 3D rendering using pseudoinstancing and 16384 agents (512×512) runs at 5.45 frames per second without any further optimization. A frame from this simulation with the above mentioned pseudoinstancing, with 16384 agents, can be seen in figure 6 a.

A different and more efficient option could be reading the agent position directly from the texture where this information is stored, avoiding unnecessary traffic from the GPU to the CPU. While texture lookups within vertex shaders are not particularly efficient in current graphics hardware, at least from GLSL, next generation graphic cards will privilege this option. Another option is to use the instancing facilities provided by DirectX.




Fig. 6. 3D Frame from examples

A frame from the predator-prey simulation with 3D rendering can be seen in figure 6 b.

4 Conclusions

We have presented a simple approach to generating agent behavior implementing Finite State Machines on the GPU. The basic idea is to use dependent texture lookups of three kinds of maps: agent space, world space and FSM maps within a fragment shader. This works at interactive rates for even a large number of agents in current generation GPUs. We have tried a simple and a more complex example. In the future we will attempt even more complex situations using several maps and more FSMs (probabilistic, layered, hierarchical). We are also working on a simpler, better way for users to construct FSM maps.

We have also implemented pseudoinstancing. While interesting, the performance is not yet good enough or general enough for what is needed. Further optimizations (culling, for example) will be attempted. While simple culling is viable, we are working on a version within the map and shader paradigm.

References

- [1] Craig W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [2] James Cremer, Joseph Kearney, and Yiannis E. Papelis. HCSM: A framework for behavior and scenario in virtual environments. *Modeling and Computer Simulation*, 5.(3):242–267, 1995.
- [3] Soraia Raupp Musse and Daniel Thalmann. Hierarchical model for real time simulation of virtual human crowds. *IEEE Transactions on Visualization and*

Computer Graphics, 7(2):152–164, 2001.

- [4] Frédéric Devillers and Stéphane Donikian. A scenario language to orchestrate virtual world evolution. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer animation*, pages 265–275. Eurographics Association, 2003.
- [5] Pascal Becheiraz and Daniel Thalmann. A behavioral animation system for autonomous actors personified by emotions. In *Proc. First Workshop on Embodied Conversational Characters (WECC '98)*, pages 57–65, 1998.
- [6] Isaac Rudomín and Erik Millán. Xml scripting and images for specifying behavior of virtual characters and crowds. In *proceedings CASA 2004*, pages 121–128. University Press, 2004.
- [7] Hantak, Chad. Comparison of Parallel Hardware Based and Graphics Hardware Based Platforms for Swarm Intelligence Simulations. Class report: <http://www.cs.unc.edu/hantak/ip.html>
- [8] Theraulaz, G. and Deneubourg, J. On Formal Constraints in Swarm Dynamics. *Proceedings of the 1992 IEEE International Symposium on Intelligent Control*, IEEE Computer Society Press, 225-233. 1992.
- [9] Harris, M., Coombe, G., Scheuermann, T. and Lastra A. Physically-Based Visual Simulation on Graphics Hardware. *Proceedings of the ACM SIGGRAPH / EUROGRAPHICS Conference on Graphics Hardware*, September 01-02, 2002, Saarbuckten, Germany.
- [10] Jeremy Zelsnack. GLSL Pseudo-Instancing. Technical Report, *NVIDIA Corporation*, 2004.