

# EECS150: Spring 2009 Project Checkpoint 4, Line Engine

UC Berkeley College of Engineering  
Department of Electrical Engineering and Computer Science

Revision B

## 1 Time Table

---

<b>ASSIGNED</b>	Friday, April 10 <sup>th</sup>
<b>DUE</b>	Week 15: April 28 <sup>th</sup> – 30 <sup>th</sup> , 10 minutes after your lab section starts

---

## 2 Overview

Now that the frame buffer, address generator, and color map are complete, we can display an entire frame of video on the screen through the video interface. However interesting the frame is, however, it is still just a single frame (a static image!). In order to make the video subsystem useful, we must have a means of changing the image that is to be displayed.

What immediately comes to mind is to just change the contents of the frame buffer through the MIPS150 processor. This definitely works: for example, to clear the entire screen, we could perform 786,432 `sw` operations (one for each pixel on the screen) and write the same color into all locations. As it turns out, however, this approach has several shortcomings. First, it is memory intensive: we have to allocate space for the `sw` instructions<sup>1</sup> Second, it wastes CPU time: the MIPS150 processor actually has to perform all of the `sw` operations! Regardless of how we organize our code, given that the MIPS150 can only perform 1 `sw` per cycle, it will always take *at least* 786,432 CPU cycles to redraw an entire frame. It will take less cycles if we only want to change a part of the frame, but it will consume CPU cycles regardless.

In order to free up CPU time and resources when performing common graphics routines, we will build a **Line Drawing Engine** (or “Vector Accelerator Engine”) that works alongside the processor and frame buffer. The line engine is the deliverable in checkpoint 4. The job of the line engine is to quickly draw a line from one set of  $x, y$  coordinates to another. This block will not only speed up the line drawing process but also allow for the CPU to be doing other work while lines are being drawn.

We will be using the **Bresenham Line Drawing Algorithm** to implement the line engine. The algorithm itself can be explained as a software routine (in a high-level language such as C). We will provide the software algorithm for you. Your job will be to create a hardware implementation for the software routine.

## 3 Checkpoint Components

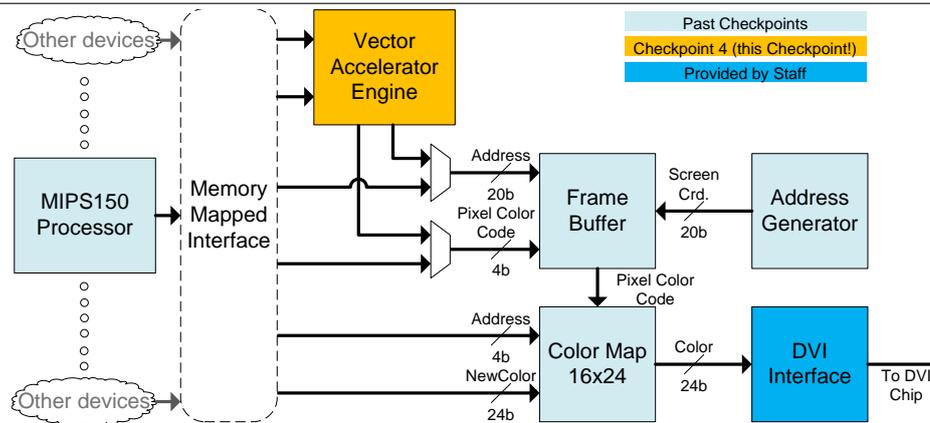
Your project, factoring in work done for checkpoint 3, now looks like what is shown in Figure 1.

Your goal in this checkpoint will be to implement the line engine in hardware. As mentioned in Section 2, we have provided the software routine, that represents what your hardware implementation will do, in Program 1.

---

<sup>1</sup>Depending on how we code the “clear screen” operation, it would take more or less space in instruction memory to store all of the instructions necessary. For this example, we assume that all of the `sw` operations are in memory back-to-back to optimize performance (like what was shown in [Lecture 14](#), slide 7).

**Figure 1** Your project circa checkpoint 4.



**Program 1** Bresenham Line Drawing Algorithm in C.

```

1#define SWAP(x, y) (x ^= y ^= x ^= y)
2#define ABS(x) (((x)<0) ? -(x) : (x))
3
4void line(int x0, int y0, int x1, int y1) {
5    char steep = (ABS(y1 - y0) > ABS(x1 - x0)) ? 1 : 0;
6    if (steep) {
7        SWAP(x0, y0);
8        SWAP(x1, y1);
9    }
10   if (x0 > x1) {
11       SWAP(x0, x1);
12       SWAP(y0, y1);
13   }
14   int deltax = x1 - x0;
15   int deltay = ABS(y1 - y0);
16   int error = deltax / 2;
17   int ystep;
18   int y = y0
19   int x;
20   ystep = (y0 < y1) ? 1 : -1;
21   for (x = x0; x <= x1; x++) {
22       if (steep)
23           plot(y,x);
24       else
25           plot(x,y);
26       error = error - deltay;
27       if (error < 0) {
28           y += ystep;
29           error += deltax;
30       }
31   }
32}

```

The implementation shown in Program 1 works for drawing lines of any slope and for lines in any quadrant of the 2D plane. Looking at the software implementation, think about how the MIPS150 processor would perform the line drawing algorithm if it had to in software. How long would drawing an entire line take?

Given all of this, there is still a lingering issue: how the line engine will be told to start drawing a line (and which line to draw). As with the other devices talking to MIPS150, the line engine communicates over the memory mapped interface. Specifically, it occupies three addresses as shown in Table 1.

**Table 1** Map of the line engine address space.

Addresses	Read/Write	Composed of ...
0x8040_0048	R	Ready
0x8040_0044	W	{y1, x1}
0x8040_0040	W	{color,y0, x0}

In Table 1, `Ready` is a single bit control signal that the MIPS150 CPU will poll to see if the line engine is ready to receive a new set of `x0,x1,y0,y1` signals. Of these, `x0`, `x1`, `y1` are 16-bit values. `y0` is a 12-bit value. This leaves 4 bits for `color`.

Thus, the operation of the line engine is as follows:

1. On system reset, the line engine will set the `Ready` bit.
2. When the CPU is ready to draw a line, it will poll the `Ready` bit.
3. Since the `Ready` bit was high, the CPU will write `x0,x1,y0,y1` and `color` into their appropriate locations.
  - (a) Writing to `0x8040_0044` (corresponding to `y1` and `x1`) will tell the line engine about that endpoint.
  - (b) Writing to `0x8040_0040` (`y0`, `x0` and `color`) will tell the line engine about that endpoint **and tell the line engine to start drawing the line.**
4. The line engine will clear `Ready` while it is drawing the line (the CPU must be sure to not modify any of the line engine's memory map locations during this time).
5. When the line engine is finished, it will set `Ready` again and the CPU will be able to write more coordinates.

In summary:

1. Writing to `0x8040_0044` has one meaning: to tell the line engine about that point.
2. Writing to `0x8040_0040` has two meanings: to tell the engine about that point and to start the drawing operation.

This scheme enables some optimizations on the software side. If you want to draw multiple lines with a single endpoint, simply write to `0x8040_0040` multiple times in a row without writing to `0x8040_0044`. Think about what happens and what gets drawn when you do this!

## 4 Architectural Concerns

Your goal after implementing the line drawing algorithm shown in Program 1 is to write one new pixel to the frame buffer every cycle. Contrast this performance with that which you would be able to obtain from the MIPS150 processor running the algorithm in pure software. Know that in order to meet timing, you may have to pipeline your implementation.

Notice that in Figure 1, there is a mux that chooses between the output of the MIPS150 processor and the output of the line engine. This mux is necessary because both of these components want to write

to the frame buffer, but the frame buffer only has one write port (remember that its other port is a dedicated read port for the address generator). This implies that we have to establish a priority scheme to handle the case when the processor and line engine both want to write to the frame buffer during the same cycle. **To avoid stalling the MIPS150 processor, we will always give priority to the processor when multiple write requests happen in the same cycle. This means that the line engine will have to stall when it and the processor both want to perform a write at the same time.** Only when the processor is finished can the line engine resume operation.

Rev.	Name	Date	Description
B	<a href="#">Chris Fletcher</a>	4/21/2009	Clarified how the line engine starts an operation (added information about the difference between writing to address 0x8040_0040 and address 0x8040_0044.
A	<a href="#">Chris Fletcher</a> <a href="#">Ilia Lebedev</a>	3/26/2009	Wrote new Document