



# King Fahd University of Petroleum and Minerals

*Department of Computer Engineering (COE)*  
COE 405  
Design and Modeling of Digital Systems

## High-Speed Parallel HDL Implementation of K-Means Clustering for Image Segmentation

Semester 201

Ahmed Alharbi | 201636500  
Hamza Alsharif | 201642320

Dr Mohammed Elrabaa

# Contents

<b>1</b>	<b>Brief Overview</b>	<b>2</b>
<b>2</b>	<b>Algorithm</b>	<b>2</b>
<b>3</b>	<b>Datapath Diagrams</b>	<b>6</b>
<b>4</b>	<b>Verilog Modules</b>	<b>10</b>
4.1	Cluster Machine . . . . .	10
4.2	BEngine . . . . .	10
4.3	Cluster Engine . . . . .	10
4.4	Distance Comparator . . . . .	10
4.5	Sum Unit . . . . .	10
4.6	Mean Divider . . . . .	11
4.7	Mean File . . . . .	11
4.8	Pixel-to-Cluster . . . . .	11
<b>5</b>	<b>Results</b>	<b>11</b>

# 1 Brief Overview

The idea behind our design is that it uses a multi-engine, multi-thread parallel processing approach. Each engine corresponds to 16 threads that represent the means and K threads will be processing and enabled depending on the user's K input. In principle, we do this by only considering the enabled means when clustering the pixels. For our current configuration we are using 10 engines. First we store 5 individual parts of the input image into 10 separate dual-port (we only use the dual read port) memory blocks that each correspond to 2 engines. This allows us to process the whole image in chunks of a 10th of the image size, all in parallel. A thread takes in the pixels as inputs and accumulates their RGB values, but only if that thread is enabled. The accumulated thread values are then sent to the sum unit where they are all added together to get a final accumulated value for the clusters. These resultant centroid accumulated values are then divided by the number of pixels assigned to each of the respective clusters. The division is done through a set of ipcore generated dividers available in Xilinx. The new mean values are then compared with the previous mean values to see if the centroid means have stabilised. In our design we use a maximum allowed difference value of 6, meaning that if the new mean value is within an error range of  $\pm 6$  of the previous mean value then it is considered to be stable. Once all of the mean values have stabilised the circuit enters the outputting stage. In this stage, the image pixels are compared with the centroid means and the closest mean in terms of Manhattan distance is outputted in place of that pixel. The output pixels are also written to a file. After the outputs have been produced, the circuit is considered to have finished its operation. We believe that this parallel approach will give us a huge boost in performance.

# 2 Algorithm

The algorithm shown below is a generalised sequential implementation of the K-Means clustering algorithm on hardware. Our actual algorithm involves a module hierarchy and a parallel design which was not easy to represent using sequential-based code. For a more accurate representation of how our circuit works see the ASM diagram.

---

```
1 // Note: The algorithm has been written to be closer to verilog-style code even though it does not
2 // necessarily match its syntax since it is supposed to be pseudo-code
3
4 Core { // these operations happen in all of the enabled cores
5   if (reset) {
6     All registers and counters = 0;
7     goto Sleep;
8   }
9
10  Sleep:
11  if (Enable)
12    goto ReadPixel;
13
14  ReadPixel:
15  Pixel = PixelIn;
16  goto CalculateDistance;
17
```

```

18 CalculateDistance:
19 Distance = absolute_value(Pixel[23:16] - Mean[23:16]) + absolute_value(Pixel[15:8] - Mean[15:8]) +
20 absolute_value(Pixel[7:0] - Mean[7:0]);
21
22 if (CompareOnly) // if the core only needs to perform a distance comparison without updating its mean value
23     goto ReadPixel;
24 else if (isClosest) // if this core's mean is the closest to the pixel
25     goto GroupPixel;
26 else
27     goto Wait;
28
29 Wait:
30 remain idle and wait until the closest core to the pixel to perform its processing;
31
32 if (UpdateMean)
33     goto UpdateMean;
34 else
35     goto ReadPixel;
36
37 GroupPixel:
38 RedAccumulator = RedAccumulator + Pixel[23:16];
39 GreenAccumulator = GreenAccumulator + Pixel[15:8];
40 BlueAccumulator = BlueAccumulator + Pixel[7:0];
41 PixelCounter = PixelCounter + 1;
42
43 if (UpdateMean) // finished reading all pixels and ready to update mean
44     goto UpdateMean;
45 else
46     goto ReadPixel;
47
48 UpdateMean:
49 BusyDivide = 1; // this will be a status signal from the divider indicating that it has finished
50 Mean = RGB Accumulators / PixelCounter; // once the division has finished BusyDivide will become 0
51
52 if (!BusyDivide && PreviousMean == Mean)
53     MeanStable = 1;
54 else if (!BusyDivide && PreviousMean != Mean)
55     MeanStable = 0;

```

```

56 goto CoreReady;
57
58 CoreReady:
59 Mean = PreviousMean; // mean value is now constant
60
61 if (AllMeansStable) // if all of the means from all cores have stabilized
62     goto CoreReady;
63 else
64     goto ReadPixel; // it might be the case that this mean core has stabilized but other mean cores have not
65 }
66
67 System { // This is the overall K-Means system process that uses the 16 cores as components within it
68 if (reset) {
69     All registers and counters = 0;
70     goto Initial;
71 }
72
73 Initial:
74 if (Start)
75     goto ReadImageSize;
76 else
77     goto Initial;
78
79 ReadImageSize:
80 ImageSize = SerialIn;
81 goto ReadK;
82
83 ReadK:
84 K = SerialIn;
85 goto StoreImage;
86
87 StoreImage: // store the image serially from DataIn input
88 Mem[Address] = DataIn;
89 Source = 0;
90
91 if (Address == ImageSize) {
92     Address = 0;
93     goto PutPixelOnBus;

```

```

94 }
95 else {
96     Address = Address + 1;
97     goto StoreImage;
98 }
99
100 PutPixelOnBus: // finished storing to memory, now we read from it
101 DataOut = Mem[Address];
102 Core(DataOut); // cores receive pixels and perform their processing
103
104 if (ReadNextPixel for all cores && Address != ImageSize) { // all cores are ready to receive a new pixel
105     Address = Address + 1;
106     goto PutPixelOnBus;
107 }
108 else if (ReadNextPixel for all cores && Address == ImageSize) {
109     Address = 0;
110     goto PutPixelOnBus;
111 }
112 else if (MoveToWrite) {
113     Address = 0;
114     goto Write;
115 }
116 else
117     goto PutPixelOnBus;
118
119 Write: // writing clusters back into the memory in place of the pixels
120 Mem[Address] = DataIn;
121 Source = 1;
122
123 if(Address == ImageSize)
124     goto Stop;
125 else {
126     Address = Address + 1;
127     goto Write;
128 }
129
130 Stop:
131 Wait for a new image to be received;

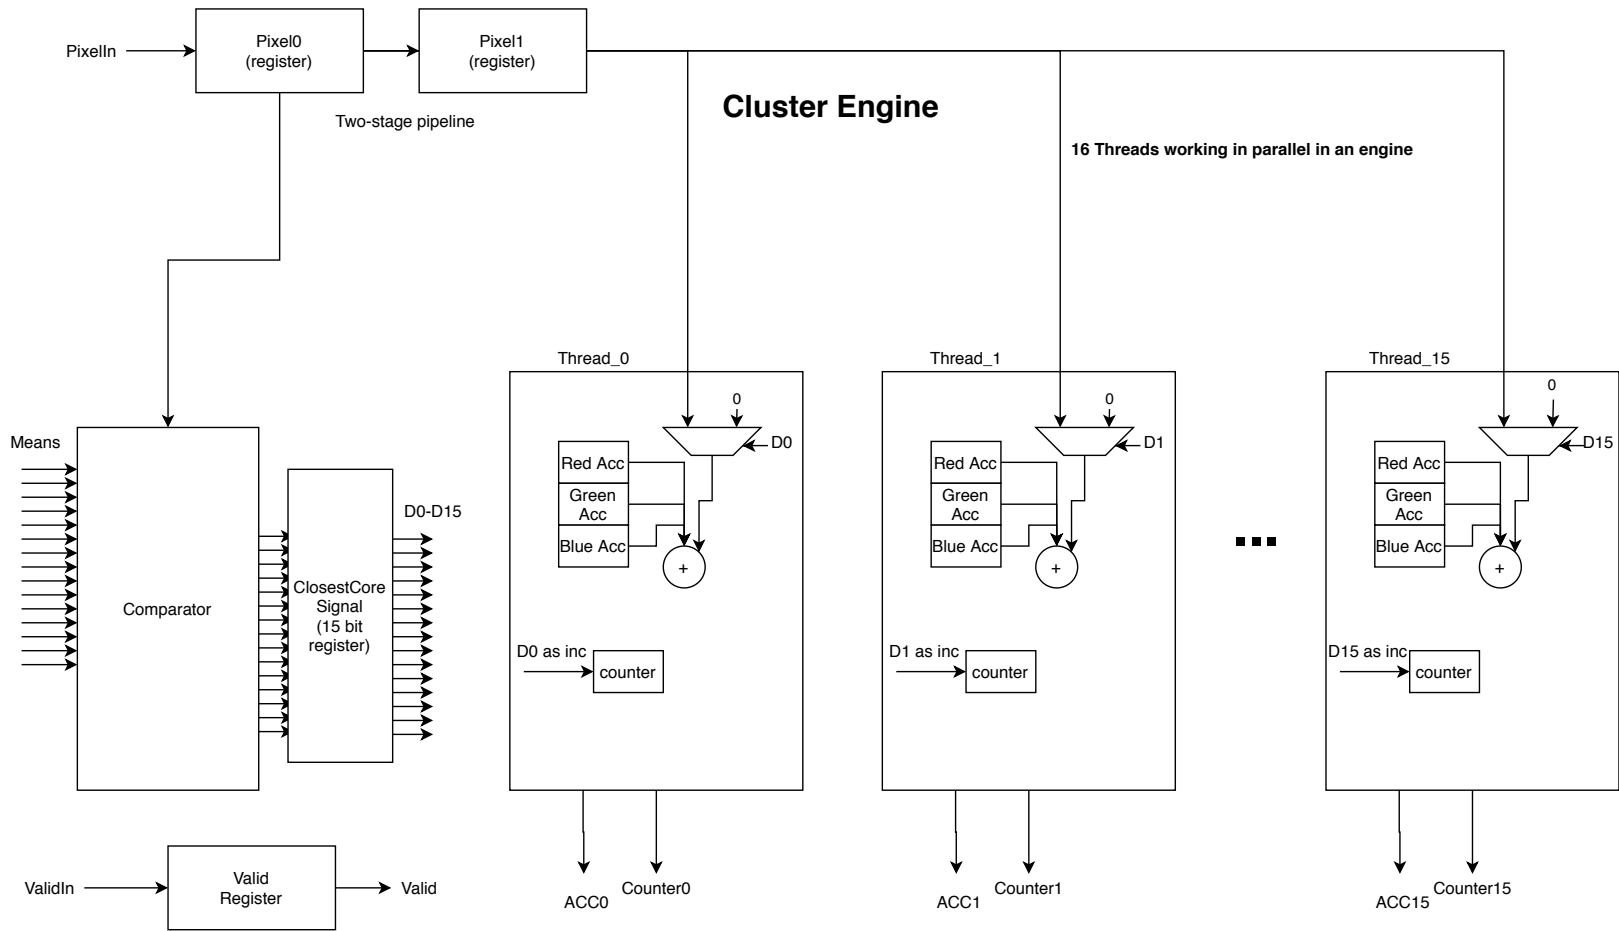
```

```
132 if (NewImage) // signal received that there is a new image
133     goto Initial;
134 else
135     goto Stop; // remain here so that the K-Means segmented image remains in the memory;
136 }
```

---

### 3 Datapath Diagrams

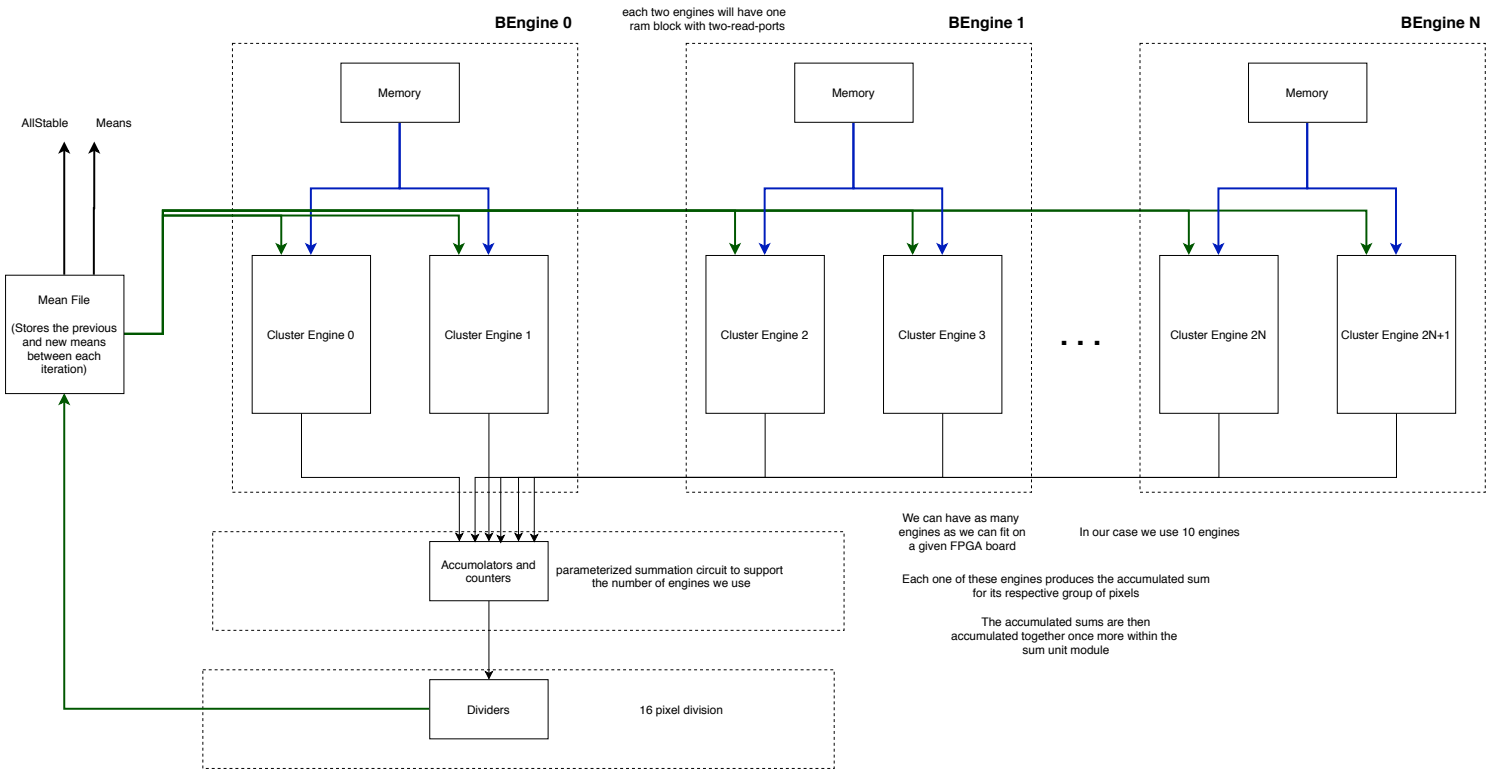

The following pages show the main datapath diagrams for our circuit. First we present the Engine module which is comprised of 16 threads. After that we show the Cluster Machine datapath which is our top level module and is the overall system. For this specific datapath we show a general design that can use  $N$  number of Engines to run  $2N$  engines. The limitation here is how much the FPGA can actually fit in terms of area. For our implementation we use a value of  $N = 5$  for 10 engines as this was the maximum allowable number of engines that fit for our given FPGA.

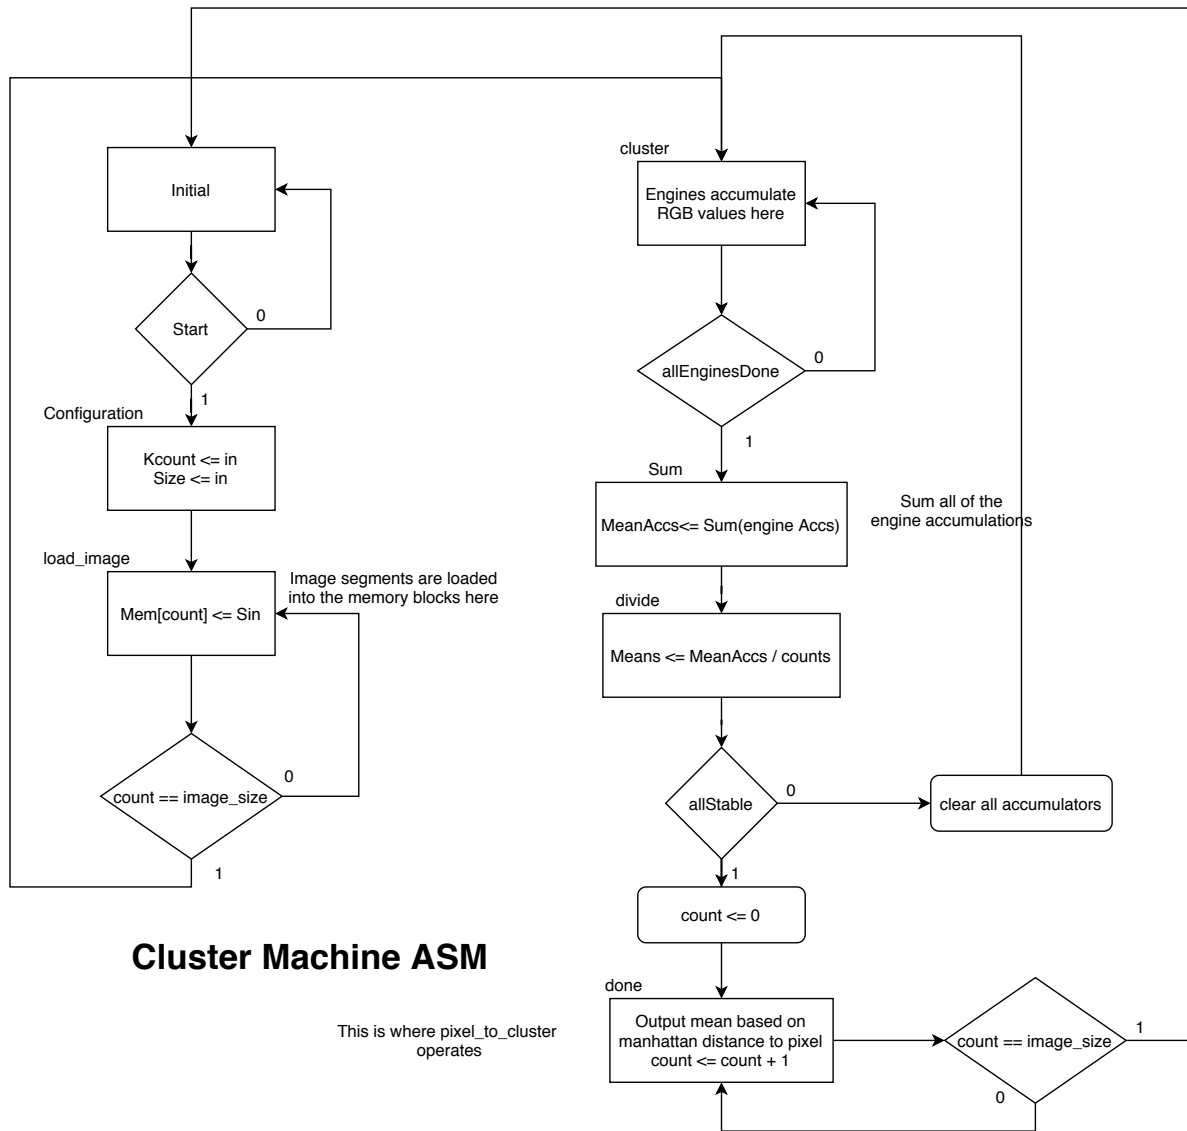


D0-D15 are ANDed with Valid to distinguish between the cases in which a PixelIn's value is 0 or if there is no input



## Cluster Machine (Overall System)





## Cluster Machine ASM

This is where pixel\_to\_cluster operates

## 4 Verilog Modules

### 4.1 Cluster Machine

As mentioned earlier, the cluster machine is the top-level module and is our parallel implementation of K-Means Clustering for image segmentation in hardware. It contains 5 BEngines that perform the accumulation of the parts of the image. The accumulations from the BEngines are accumulated once more within the sum unit into RGB components for each of the respective clusters. The sum unit accumulations are then divided using a set of dividers and their mean outputs for each iteration is stored in the mean file. The mean file compares the new means with the means from the previous iteration and checks if they are stable, and if all of them are stable then the means computation is complete. Now we output the new segmented image through mean\_out using the pixel\_to\_cluster module. This module compares the image pixels from the main image memory block and outputs the new pixel based on the Manhattan distances of the computed means to the original image pixel. The Cluster Machine also has a control unit as a state machine to synchronize the transitions between the different parallel stages.

### 4.2 BEngine

A BEngine contains a single dual-port memory block and two engines. Each engine processes half of what is contained of the image in the memory block, so one engine would be operating on the pixels at even addresses and the other would be operating on the pixels at odd addresses.

### 4.3 Cluster Engine

Cluster engines contain the accumulators and counters for each of the cluster threads for a set of pixels belonging to that engine, as well as the distance comparator and Manhattan modules. These engines cluster all of the RGB values and increment the thread pixel counters based on the Manhattan distances which are compared using the distance comparator module. The engine is two-stage pipelined to reduce the overhead of the Distance Comparator and Manhattan modules.

### 4.4 Distance Comparator

The distance comparator module takes in the thread distances computed using the Manhattan module and then compares them all through a sequential stage based comparison scheme to find the thread that corresponds to that smallest Manhattan distance. In the first stage we compare 8 distances with 8 other distances and produce the minimum out these 8. These 8 minimums are split into two sets of 4 minimums that are compared with each other in the second stage. In the same way, the resulting 4 minimums are then split into two sets of 2 minimums. The final stage compares these last two minimums and produces a signal based on the final minimum that corresponds to one of the inputs to a priority encoder. The priority encoder outputs the index of the nearest cluster so that a pixel may be accumulated to it within a thread.

### 4.5 Sum Unit

The sum unit received the accumulated values and counters for all of the cluster threads for all of the engines and accumulates them into a single value for each of the respective means. It does this

by sequentially adding the threads using a single accumulator for the RGB values and pixel counts for each of the threads. The output of this module is the final accumulated values for each of the means that are ready for division in the next stage.

## 4.6 Mean Divider

The mean divider module is simply an encapsulation of 3 dividers corresponding to the RGB values of the pixels. The actual dividers are standard generated Xilinx ipcore dividers. Throughout testing we found that the Xilinx ipcore dividers take 26 cycles to finish 1 division operation. The mean divider module also includes the signal `rgb_ready` that indicates that all three of the divisions have finished. In the Cluster Machine we use 16 of these mean dividers to divide the cluster accumulators by the number of pixels assigned to the clusters.

## 4.7 Mean File

The mean file is the module that stores the values of the means. Initially, we store constant values for each mean that are equally spaced between the range of 255 (the max value that a pixel component can have). We store the mean values between each iteration so that we can compare the previous iteration's mean value with that of the new iteration to check for their stability. This is done through the stability checker module, which is a modified comparator that checks to see if the difference of the previous iteration means and the new iteration means is less than 6. The stability checker produces a stable signal for the respective mean if the threshold difference condition is satisfied. If all of the means are stable, then an output signal `allStable` is produced indicating that the iterations can stop.

## 4.8 Pixel-to-Cluster

This module is used to output the final segmented pixel values after the means have been computed. This happens in a similar fashion to how the pixels are assigned to each cluster except that in this case we read from the main (whole) image memory block rather than separate engine memory blocks. We also re-use to Manhattan and Distance Comparator modules to be able to assign a mean value to an input pixel. In the test bench, the output mean values are also written to a file so that we may be able to compare our image outputs with the golden model.

# 5 Results




Figure 1: A resized version of the original Wikipedia article image, used for testing.

The final results that we obtained through both simulation and synthesis were highly satisfactory. The algorithm efficiently works as required for different values of  $K$ . We tested the algorithm on



Figure 2:  $K = 4$ , Cycles spent processing means = 1701.



Figure 3:  $K = 8$ , Cycles spent processing means = 1276.



Figure 4:  $K = 12$ , Cycles spent processing means = 2127.




Figure 5:  $K = 16$ , Cycles spent processing means = 2551.

a resized version of the image given on the Image Segmentation article on Wikipedia (shown in figure 1) for values of  $K = 4, 8, 12,$  and  $16$  (shown in figures 2, 3, 4, 5). The number of cycles taken for computing 4, 8, 12, and 16 means were 1701, 1276, 2127, and 2551 respectively. It should be noted that the image size was 3800 pixels, so to calculate the total number of cycles you would also have to add  $2 \cdot 3800$  to the cycles spent processing because the circuit first spends 3800 cycles to read the image, and then another 3800 cycles to output the segmented image. The circuit also synthesizes and maps on our given board (xc7a100t-3fgg484) and is within the area constraints of the board. The minimum period in which our circuit operates at is 11.471 ns (Max Frequency = 87.180 MHz). The timing report for the circuit is shown below in figure 6. We have also shown that our design meets the area constraints through the device utilization report given in figure 7. For the given clock period, if we compute the total time spent processing it would be clock period \* total clock cycles, where the total clock cycles is the sum of the reading, processing, and writing cycles. For our tests above, we can find that for 4, 8, 12, and 16 means the total time for the circuit would be 106.68  $\mu\text{s}$ , 101.82  $\mu\text{s}$ , 111.58  $\mu\text{s}$ , and 116.44  $\mu\text{s}$  respectively. If we only consider the cycles spent processing (neglecting the cycles spent reading and writing), we will find that the total processing times for 4, 8, 12, and 16 means are 19.51  $\mu\text{s}$ , 14.64  $\mu\text{s}$ , 24.40  $\mu\text{s}$ , and 29.26  $\mu\text{s}$  respectively. It must be stated however, that the FPGA choice given to us was disappointing as we believe that we can

run at much less clock cycles if we were able to fit more engines on the board, but unfortunately we were limited to using 10. For an in-depth analysis of our module's implementation on an actual board, see the Xilinx reports submitted with this project.

```

Timing Summary:
-----
Speed Grade: -3

Minimum period: 11.471ns (Maximum Frequency: 87.180MHz)
Minimum input arrival time before clock: 2.164ns
Maximum output required time after clock: 16.085ns
Maximum combinational path delay: No path found
=====

```

Figure 6: Timing Report

Device Utilization Summary (estimated values)				[1]
Logic Utilization	Used	Available	Utilization	
Number of Slice Registers	76069	126800	59%	
Number of Slice LUTs	58276	63400	91%	
Number of fully used LUT-FF pairs	30991	103354	29%	
Number of bonded IOBs	52	285	18%	
Number of Block RAM/FIFO	18	135	13%	
Number of BUFG/BUFGCTRLs	2	32	6%	

Figure 7: Device Utilization Report