

EECS 468  
Project Review  
Study of Loop Perforation on GPUs

Alex Broad  
Akhil Guliani

# Aims and Objectives

- We propose to find the best method to apply loop perforation on a GPU like device with minimal loss in computational accuracy with reduction in both time taken to compute the results and reduction of power consumption.

# What is loop perforation ?

- Loop perforation transforms loops to execute a subset of their iterations.
- The goal is to reduce the amount of computational work
  - amount of time
  - resources such as power

# Ways of doing loop Perforation

- Criticality Testing :
  - Finding which methods produce unacceptable results
  - Throwing out a warp == not good
- Perforation Space Exploration
  - Finding all the various possible results
  - Throwing out threads selectively
    - Increment only one dimensions
    - Increment both dimensions
    - Skip certain loops based on divisibility by factor

# Benefits of Approximate computing

- Approximate computing is a new research direction that improves efficiency by carefully relaxing correctness constraints
- Quicker results
- Lower utilization of Resources

# Our Case

We Study effects of loop perforation on

- Filter like algorithm such as Image Blurring
- Feature classification algo's like : Local Binary Patterns
- Feature vector matching example Histogram Comparison

And compare the results with standard algorithms available on the CPU and GPU

# Method

- Implement the algorithms on a CPU and a GPU
- Optimize the performance on the GPU
- Find avenues for loop perforation and try various levels of perforation
- Obtain the various performance evaluation criteria
- Conclude Results

# Image Blurring

- We use Box Blur Algorithm
  - Spatial domain linear filter
  - Each pixel equal to average of its neighbors
  - low-pass ("blurring") filter
- Easily implementable on the GPU
- Ideal Example to study Perforation



# Image Blurring – CPU

```
27     bool perf = false;
28     int perf_rate = 5;
29
30     // We don't do any blurring around the boarder
31     for (int x = one_filter_dim; x < src.rows-one_filter_dim; x+=i) // should d
32     {
33         for (int y = one_filter_dim; y < src.cols-one_filter_dim; y+=j)
34         {
35             if (perf && ((y % perf_rate) && (x % perf_rate)) == 0)
36             {
37                 continue;
38             }
39             r_sum = 0.0; g_sum = 0.0; b_sum = 0.0;
40             for (int f_x = x-one_filter_dim; f_x <= x+one_filter_dim; f_x++)
41             {
42                 for (int f_y = y-one_filter_dim; f_y <= y+one_filter_dim; f_y++)
43                 {
44                     r_c = src.at<cv::Vec3b>(f_x,f_y)[0];
45                     g_c = src.at<cv::Vec3b>(f_x,f_y)[1];
46                     b_c = src.at<cv::Vec3b>(f_x,f_y)[2];
47                     r_sum += r_c; g_sum += g_c; b_sum += b_c;
48                 }
49             }
50             dst.at<cv::Vec3b>(x,y)[0] = r_sum/(filter_dims*filter_dims);
51             dst.at<cv::Vec3b>(x,y)[1] = g_sum/(filter_dims*filter_dims);
52             dst.at<cv::Vec3b>(x,y)[2] = b_sum/(filter_dims*filter_dims);
53         }
54     }
55
```

# Image Blurring – GPU

```
7  __global__
8  void blur(unsigned char* input_image, unsigned char* output_image, int width, int height) {
9
10     const unsigned int offset = blockIdx.x*blockDim.x + threadIdx.x;
11     int x = offset % width;
12     int y = (offset-x)/width;
13     int fsize = 5; // Filter size
14     if(offset < width*height) {
15
16         float output_red = 0;
17         float output_green = 0;
18         float output_blue = 0;
19         int hits = 0;
20         for(int ox = -fsize; ox < fsize+1; ++ox) {
21             for(int oy = -fsize; oy < fsize+1; ++oy) {
22                 if((x+ox) > -1 && (x+ox) < width && (y+oy) > -1 && (y+oy) < height) {
23                     const int currentoffset = (offset+ox+oy*width)*3;
24                     output_red += input_image[currentoffset];
25                     output_green += input_image[currentoffset+1];
26                     output_blue += input_image[currentoffset+2];
27                     hits++;
28                 }
29             }
30         }
31         output_image[offset*3] = output_red/hits;
32         output_image[offset*3+1] = output_green/hits;
33         output_image[offset*3+2] = output_blue/hits;
34     }
35 }
```



# Image Blurring – A comparison

## LOOP PERFORATION

---






### BLURING

#### CPU







No Perf - Clock Time (for 50 iterations) = 40.019  
loop+=2 - Clock Time (for 50 iterations) = 10.259  
loop+=5 - Clock Time (for 50 iterations) = 1.611

#### GPU -

No Perf - Clock Time (for 50 iterations) = 0.263085  
loop+=2 - Clock Time (for 50 iterations) = 0.155025  
loop+=5 - Clock Time (for 50 iterations) = 0.121848



# Different Types of Loop Perforation





gpu\_blur

gpu\_blur150310\_00...pture\_000.nvreport


CUDA Launches Hierarchy Flat

Filter

Viewing: 1 / 1

Function Name	Grid Dimensions	Block Dimensions	Start Time (µs)	Duration (µs)	Occupancy	Registers per Thread	Static Shared Memory per Block (bytes)	Dynamic Shared Memory per Block (bytes)	Cache Configuration Executed	Local Memory per Thread (bytes)	Device Name	Context ID	Stream ID	Process Name	Occupancy [0]: Allocated Warps Per Block	Occupancy [0]: Allocated Registers Per Block	Occupancy [0]: Allocated Shared Memory Per Block	Occupancy [0]: Max Block Limit Warps	Occupancy [0]: Max Block Limit Registers
1 blur	{1536, 1, 1}	{512, 1, 1}	326,832.236	4,779.552	100.00 %	21	0	0	PREFER_SHARED	0	GeForce GTX 870M	1	1	gpu_blur.exe	16	12288	0	4	5

# Without Perforation on GPU



With Perforation on GPU

# Local Binary Patterns

- The Algorithm

# Local Binary Patterns – CPU





# Local Binary Patterns – GPU

# Local Binary Patterns – A comparison


CPU (Change Loop in LBP.cpp file - line 319)

- No Perf - Example took 0.166892s
- loop+=2 - Example took 0.134939s
- loop+=3 - Example took 0.111867s



GPU -

- No Perf - Example took 0.063818s
- loop+=2 - Example took 0.104177s
- loop+=3 - Example took 0.0965s



# Conclusions And Discussion

- Loop Perforation studied with Image blurring and LBP algorithms
- Used Various Methods for Loop Perforation
  - Throwing out Warps Not a Good Idea
  - Doing it on the CPU produces patterned results
- GPU Implementations Generally faster
  - For Blurring we got 150x Speed up (without perforation)
  - We got 200x Speed up (with Perforation +2)
  - We got 66x Speed up (with Perforation +5)
- Even Faster with Perforation
  - 1.7x Speed up (with Perforation +2)
  - 2.1x Speed up (with Perforation +5)

# Future Work

- Try and add Power Consumption Statistics for the current results
- Try Perforation on a Feature Vector matching Algorithm
- Using perforated LBPs and Perforated Feature vector matching For face detection

Thank You

# Resources Used

Bad Result CPU

