A Low Power Hardware Implementation of Multi-Object DPM Detector for Autonomous Vehicles

Alaa Ali, Oladiran Olaleye, Bappaditya Dey, Kasem Khalil, Magdy Bayoumi

The Center for Advanced Computer Studies (CACS), University of Louisiana at Lafayette

a.ali@ieee.org, {ogo8842, bxd9836, kmk8148, mab0778} @louisiana.edu

Abstract

Object detection is a fundamental process in traffic management systems and self-driving cars. Deformable part model (DPM) is a popular and competitive detector for its high precision. This paper presents a programmable, low power hardware implementation of DPM based object detection for real-time applications. Our approach employs a very fast object detection pipeline with complementary techniques such as fast feature pyramid, Fast Fourier Transform (FFT) and early classification to accelerate DPM with a reasonable accuracy loss and achieves a speed-up of 50x and 6x over original DPM and cascade DPM respectively on single core CPU. The hardware circuit uses 65nm CMOS technology and consumes only 36.5mW (0.81 nJ/pixel) based on the post-layout simulation. The ASIC has an area of 3362 kgates and 295.5 KB on-chip memory and the design utilizes two simultaneous engines to process two independent object categories with 8 deformable parts per category.

Fast Feature Pyramid

Process flow to speed-up feature computation.

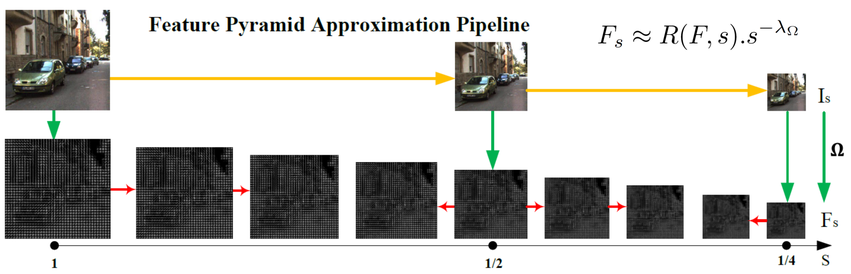

Look-up Table HOG

- **We use look-up tables to replace HOG complex operations based on concept of having finite combinations of gradient and orientation.**
- **It dramatically reduce the time bottleneck of HOG feature map computation,**
- **Gradient orientations and magnitude are pre-calculated.**

Conclusion

- **A low power and real-time hardware architecture to accelerate deformable part models and achieve a real-time detector with a very similar precision.**
- **The hardware circuit uses a 65 nm CMOS technology and processes 1920 × 1080 videos at 42 fps while consuming only 36.5 mW and resulting in an energy efficiency of 0.81 nJ/pixel.**
- **Our results on different benchmarks present DPM to real-time applications such as driver assistance and autonomous vehicles.**

Fast Fourier Transform

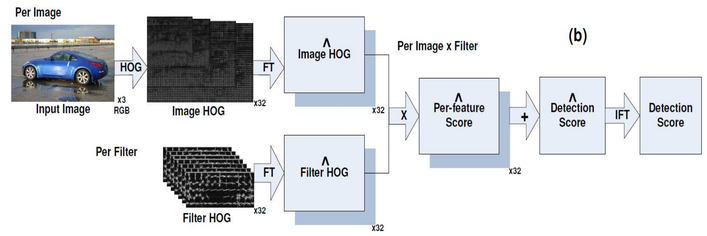

To convolve L HOG filters and sum them across K HOG features:

$$C_{fourier/img} = \underbrace{KC_{FFT}}_{forward\ FFTs} + \underbrace{LC_{FFT}}_{inverse\ FFTs} + \underbrace{KLC_{mul}}_{multiplications}$$

$$T_{fast} = \underbrace{K(L+R)v(MN)}_{reading} + \underbrace{KLRu(MN)}_{multiplications} + \underbrace{LRv(MN)}_{writing}$$

Feature Pyramid Generation

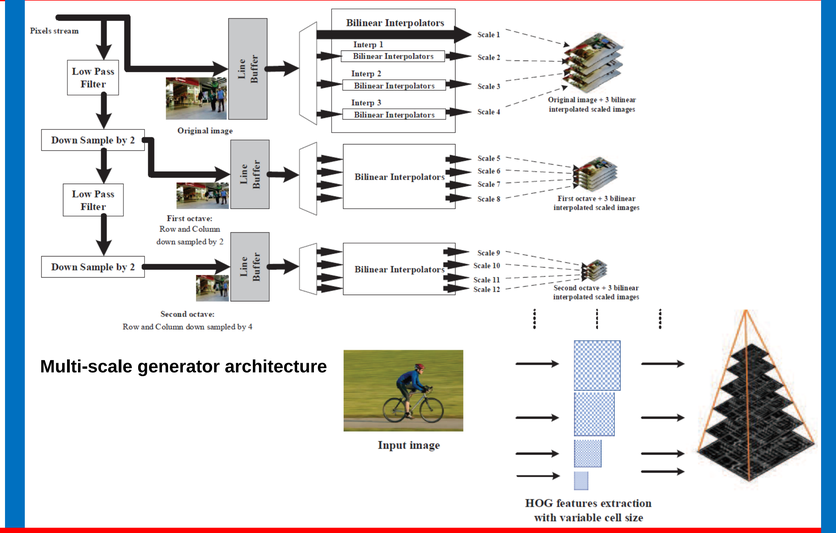

Multi-scale generator architecture

Hardware Architecture

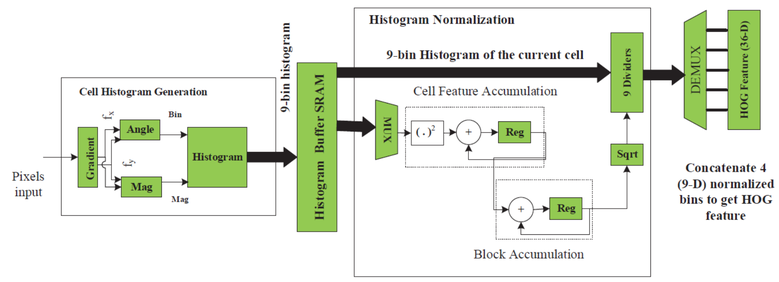

Block diagram of HOG feature extraction

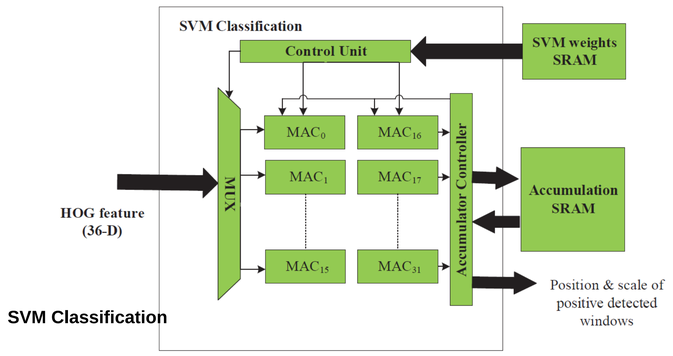

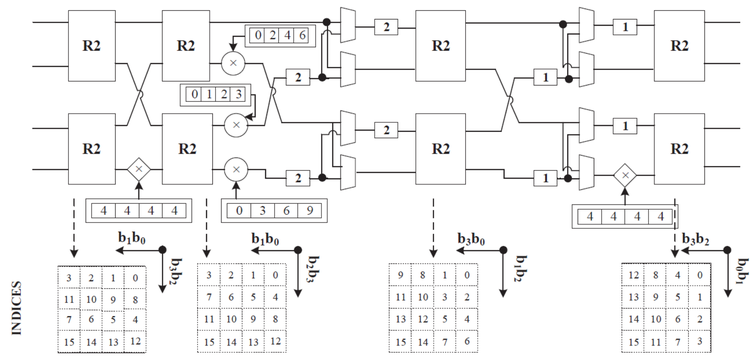

4-parallel radix-2^2 feedforward architecture for the computation of the 16-point FFT

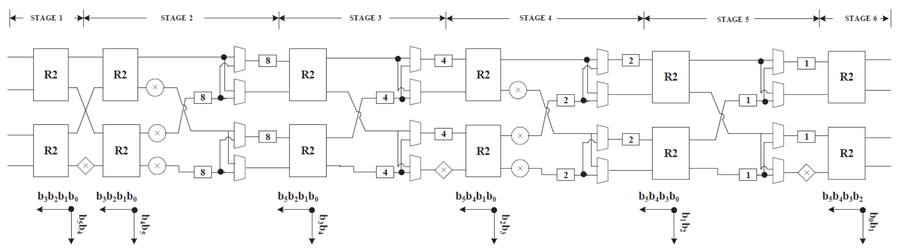

4-parallel radix-2^2 feedforward 64-point DIT FFT architecture

Experimental Results

Table 1. Comparison with ODROID-XU3 board processing Full HD 1920 × 1080 frames

Platform	Cortex-A7 [16]		Cortex-A15 [16]		DPM [16]		Our Work	
	1 core	4 cores	1 core	4 cores	0.77V	1.11 V	0.77V	1.11 V
Process Technology	28nm HKMG		28nm HKMG		65nm CMOS		65nm CMOS	
Throughput(fps)	0.04	0.10	0.11	0.24	30	60	42	74
Power(mW)	155.6	383.5	1,703.8	3,575.6	58.6	216.5	36.5	182.4
Energy(nJ/pixel)	1,881.8	1,849.0	7,301.2	7,165.7	0.94	1.74	0.81	1.48

Table 2. Performance comparison on PASCAL VOC 2007 dataset for multiple hardware implementations.

	HOG [14]	DPM [16]	Our Work
Process	65nm	65nm	65nm
Chip size	4.2 × 2.1 mm^2	4.0×4.0 mm^2	4.2×4.0 mm^2
Input resolution	1920 × 1080	1920 × 1080	1920 × 1080
Multi-scale	one scale	12 scales	12 scales
Deformable parts	No	8	8
Object classes	2	2	2
Frame rate	30	30	42
Frequency	84.3 MHz	62.5 MHz	62.5 MHz
Power	84 mW	58.6 mW	36.5 mW
Energy/pixel	1.35 nJ	0.94 nJ	0.81 nJ
Mean AP	18.5%	26%	31.4%