Technische Universiteit Eindhoven University of Technology
Department of Mathematics and Computer Science
The System Architecture and Networking group

Accelerating SIFT Feature extraction with a vector DSP- A feasibility study

Master Thesis

Arjun Sathyanarayana Prasad

ERICSSON

Supervisors:
Prof. Dr. Ir. C. H. (Kees) Van Berkel
Ir. D. (David) Van Kampen
Dr. Ir. W. E. H. (Wim) Kloosterhuis
Prof. Dr. K. G. W. (Kees) Goossens

This thesis work has been carried out at Ericsson. The content of this report is confidential until August 29, 2015

Acknowledgments

I would like to express my deepest gratitude to Professor Dr. Ir. Kees Van Berkel from TU, Eindhoven for providing me an opportunity to conduct this thesis work under him. He has been motivating and has guided me to carry out this thesis work in a fruitful manner.

I am extremely grateful to David van Kampen from Ericsson for tutoring me throughout this thesis work. He has given me valuable suggestions and feedback in each and every stage of my work. His willingness to give time so generously is immensely appreciated.

Special thanks to Dr. Ir. Wim Kloosterhuis from Ericsson for advising me in the critical stages of the thesis work and providing me the constructive feedback. He has made time to review my work even with a busy schedule which is much appreciated.

Lastly, I would also like to extend my thanks to Ericsson and specially the EVP team for being helpful and accommodating for this period. This has been a great learning experience.

Arjun Sathyanarayana Prasad

Table of Contents

1. Introduction

Over the past decade we have seen rapid adoption of the mobile smartphones among users worldwide. It is estimated that the number of smartphones sold in the year 2013 itself was close to 1 billion [1]. It is not only that these devices are affordable but the reason for this evolution of the smartphone business can also be attributed to tremendous advancements in the processor, battery and memory technologies. These devices can handle tasks that were unimaginable few years ago and hence the name Smartphone. The number of applications being built into these mobile devices is ever increasing but so is the computational capability of mobile processors. One of the surveys [2] show that the majority of users purchase smartphones based on its ability to support personal applications. The report also indicates that users are most likely to use their phone for multimedia applications like playing games, photo editing, streaming audio/video content and web surfing. Also more and more intriguing applications like augmented reality is just around the corner and will soon find its way into our lives [3]. Hence, the ability to handle multimedia content is very crucial and this is not true just for smartphones but also for other technologies that are in use today such as tablets, high tech glasses, play stations etc. But, it is well known that processing of multimedia content is highly resource intensive and has a direct impact on the devices` performance.

Of particular relevance in this report is the field of Image analysis and computer vision. Image analysis involves a series mathematical transformations performed on digital images to extract meaningful information. One of the ultimate goals that is sought out to be achieved through Image analysis is to use computers for emulating human vision. A computer needs to be capable of learning from its visual inputs, make inferences and take the necessary action. Computer vision is a field of study that encompasses wide variety of algorithms that are aimed at mimicking functionality of human visual system. Computer vision and Image analysis concepts are applied in numerous areas spanning from a simple photo editor to complex augmented reality applications.

Sophisticated computer vision applications require performing transformations on images at very high frame rates. An image consists of huge chunk of data (pixels) in the orders of hundreds of KBs (Kilobytes) or few MBs (Megabytes) and has direct impact on performance of systems when analyzing them. Running Image analysis algorithms on GPPs (General purpose processors) is often not recommended as they are very inefficient. Especially in Embedded systems with tight constraints on timing, it is difficult to achieve real time performance on resource intensive algorithms by using only GPPs (CPU). Hence different alternatives have been explored in the form of Co-processors, DSPs (Digital signal processors), GPUs (Graphic processing units), SIMD/MIMD (Single instruction multiple data/ Multiple instructions multiple data) processors and hardware accelerators to work in tandem with CPUs to share the workload. GPUs are by far the most preferred choice in the state of the art mobile platforms while it is no surprise that custom accelerators are the most computationally efficient alternative. GPUs offer high degree of flexibility which can be exploited to improve performance of variety of algorithms, but they are power hungry. Accelerators, though very efficient, offer least degree of flexibility and hence are not desirable for varying standards. SIMD processors occupy a sweet spot in this hierarchy of processing units as they provide a better degree of flexibility when compared to accelerators and also they can be computationally efficient when compared to GPUs if the algorithms are vectorizable. In many of the image transformations this is often the case. Since SIMD processors have similar architectures when compared to GPUs, considering SIMD processors as a design alternative in mobile platforms for Image analysis is highly relevant in this context.

In this report we focus on implementation of a widely popular and robust feature extraction algorithm called Scale invariant feature transform (SIFT) [4] that is applied in variety of computer vision applications. Some of the transformations used in SIFT exhibit high degree of data level parallelism and these are mapped on a SIMD processor developed by Ericsson called the Embedded vector Processor (EVP).

The structure of this chapter is as follows: in section.1.1 we briefly discuss basics of a digital image and its characteristics. Section.1.2 gives a conceptual description of the architecture of embedded vector processor (EVP)

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.