
DESIGN AND ANALYSIS OF ARCHITECTURES FOR STEREO VISION

Von der Fakultät für Elektrotechnik und Informatik
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades
Doktor-Ingenieur
genehmigte Dissertation von

Dipl.-Ing. Christian Banz

geboren am 31. Januar 1983 in Hannover, Deutschland

2013

1. Referent
2. Referent
Tag der Promotion

Prof. Dr.-Ing. Peter Pirsch
Prof. Dr. Tech. Jarmo Takala
12. September 2013

Berichte aus der Informationstechnik

Christian Banz

Design and Analysis of Architectures for Stereo Vision

Shaker Verlag
Aachen 2013

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.

Zugl.: Hannover, Leibniz Univ., Diss., 2013

Copyright Shaker Verlag 2013

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publishers.

Printed in Germany.

ISBN 978-3-8440-2401-2

ISSN 1610-9406

Shaker Verlag GmbH • P.O. BOX 101818 • D-52018 Aachen

Phone: 0049/2407/9596-0 • Telefax: 0049/2407/9596-9

Internet: www.shaker.de • e-mail: info@shaker.de

ACKNOWLEDGMENTS

This dissertation could never have been completed successfully without the contribution and support of many people. Firstly, I would like to thank my doctoral advisor Prof. Dr.-Ing. Peter Pirsch for his excellent support and guidance from my first days at the Institute of Microelectronic Systems all the way through to my dissertation. Without his trust, several great opportunities would not have been available to me.

I would also like to thank Prof. Dr. Tech. Jarmo Takala for his generous acceptance of the role of Coexaminer. Special thanks goes also to Prof. Dr.-Ing. Holger Blume for being a splendid sparring partner in many discussions, for providing excellent working conditions at the Institute of Microelectronic Systems, and for agreeing to act as the Examining Committee's Chairman.

I am grateful for the funding of the Hans L. Merkle Stiftung. The remarkable programme of the Trust has given me many interesting insights which will stay with me for life. Many thanks also go to the Trust Programme's mentors at Robert Bosch GmbH. Dr. Wolfgang Niehsen, Jochen Wingbermühle, Arne Zender, and Henning von Zitzewitz have given me great scientific and practical advise and always inspired me with new ideas.

I would like to thank all my former colleges at the Institute of Microelectronic Systems for the collegiate atmosphere. Special thanks also go to Prof. Dr. Guillermo Payá Vayá for our collaboration. Further thanks go to Steffen Blume, Daniel Pefferkorn, Sebastian Hesselbarth, Dr. Javier Martín Langerwerf, Nico Mentzer, and Dr. Carsten Dolar for many insightful discussions.

Finally, I would like to thank my parents, Hildegard and Harald Banz, for their constant support and always being there for me; my brother Dr. Lothar Banz for mathematical and other advice at all times; Sheila Isabell White for being there for me when it matters. My deepest thanks go to Dr. Wiebke Schaper for her patience, encouragement, constant support, and for listening until she knew what an FPGA is.

Christian Banz

Hannover, September 2013

ABSTRACT

The demand for high performance image processing systems, driven by ever-evolving image processing algorithms, has begun to grow immensely in recent years. For the implementation of such systems, the designer faces a vast number of choices and has to confront a set of conflicting constraints comprising performance, power consumption, and flexibility, to name just a few. In order to identify optimal solutions there is a prevailing need to constantly evaluate and quantify existing and emerging architectures using relevant, state-of-the-art image processing algorithms. This issue is especially apparent in advanced driver assistance systems, where for many relevant algorithms efficient implementations for demanding throughput constraints are unavailable on any architecture. One of these algorithms is semi-global matching (SGM) which enables 3D perception using inexpensive, passive stereo cameras.

This work presents the design, optimization and analysis of high efficiency implementations in the hardware and software domain of the semi-global matching algorithm. The target is a heterogeneous set of architectures comprising a wide spectrum of current architectural concepts for image processing: two ASPs (one RISC processor and one VLIW processor), a GPU, a multi-core GPP, and a dedicated architecture mapped onto FPGA and standard cell ASIC. A design space exploration (DSE) framework is introduced to assess architectures quantitatively in terms of silicon area, energy consumption, and throughput performance. It contains a specifically designed FPGA-based SoC framework for image processing tasks enabling emulation-based analysis and rapid prototyping. Using the DSE framework, detailed analyses of all respective implementation options for each architecture (intra-architecture analysis) and across architectures (inter-architecture analysis) are performed. The SoC framework and FPGA results are further applied to compose a fully functional real-time stereo vision system.

The evaluation shows that there is an intra-architecture optimization potential in terms of area efficiency and power efficiency of factor five to ten. Across the individually optimized implementations, there are inter-architecture efficiency differences of typically one to two orders of magnitude. Moreover, maximum throughput performance and absolute power consumption both cover two orders of magnitude. Consequently, the choice of architecture for a particular application scenario depends on the importance attributed to the architectures' various properties.

The architecture-tailored parallelization concepts required for highest performance emphasize the mutual dependency between software and hardware the need for communication between both engineering domains.

Keywords — design space exploration, image processing architectures, stereo vision, semi-global matching

KURZFASSUNG

Die Einsatzmöglichkeiten von Bildverarbeitungssystemen zur maschinellen Umfeldinterpretation werden durch den rapiden technologischen und algorithmischen Fortschritt fortwährend erweitert. Jedoch erfordert der Entwurfsprozess solcher Systeme im zunehmenden Maße die Justierung einer Vielzahl an Designoptionen, die einen komplexen Einfluss auf die zum Teil gegenläufigen Entwurfskriterien wie z.B. Durchsatzrate, Verlustleitung, und Flexibilität haben. Die beständige Evaluierung aufkommender sowie bestehender Architekturen mit relevanten Bildverarbeitungsalgorithmen ist elementar, um frühzeitig die optimale Lösung bestimmen zu können. Dieses gilt insbesondere für Fahrerassistenzsysteme, bei denen für viele relevante Algorithmen noch keine schritthaltenenden Implementierungen für eingebettete Systeme bekannt sind. Ein Algorithmus dieser Kategorie ist das Semi-Global Matching zur 3D-Umfelderfassung mittels passiver, kostengünstiger Stereokameras.

Die vorliegende Arbeit präsentiert Entwurf, Optimierung und Analyse von hocheffizienten Hardware- und Software-Implementierungen des Semi-Global-Matching-Algorithmus für ein diversitäres Set aktuell verbreiteter Architekturkonzepte. Die Evaluierung umfasst zwei ASPs (RISC und VLIW), eine GPU, einen Universalmehrkernprozessor, sowie eine dedizierte Architektur, die auf FPGAs und Standardzell-ASICs abgebildet wird.

Um eine einheitliche quantitative Bewertung der Architekturen hinsichtlich Siliziumfläche, Verlustleistungsaufnahme und Durchsatzrate durchführen zu können wird eine Entwurfsraumexplorationsumgebung eingeführt. Diese beinhaltet eine FPGA-basierte SoC-Entwicklungsumgebung, welche emulationsbasierte Analysen ermöglicht. Darauf basierend wird eine detaillierte Untersuchung der Implementierungsalternativen innerhalb einer Architektur (Intraarchitekturanalyse) sowie der Architekturen zueinander (Interarchitekturanalyse) vorgestellt. Die praktische Anwendung der SoC-Umgebung wird weiterhin durch den Aufbau eines echtzeitfähigen Stereovideosystems gezeigt.

Die Auswertung zeigt, dass ein Intraarchitekturoptimierungspotential von Faktor fünf bis zehn hinsichtlich Flächen- und Verlustleistungseffizienz besteht. Das Optimierungspotential zwischen den individuell optimierten Implementierungen der untersuchten Architekturkonzepte beträgt typischerweise ein bis zwei Größenordnungen. Weiterhin umspannen Durchsatzrate und absolute Verlustleistung jeweils einen Bereich von ca. zwei Größenordnungen. Konsequenterweise kann eine Festlegung der Architektur nur unter zusätzlicher Berücksichtigung der anwendungsspezifischen Bedeutung der jeweiligen Eigenschaften vorgenommen werden.

*Schlagworte — Entwurfsraumexploration, Architekturen zur Bildverarbeitung,
Stereobildverarbeitung, Semi-Global Matching*

CONTENTS

1	Introduction	1
1.1	Research Objectives	3
1.2	Overview	4
2	Design Space Exploration	7
2.1	Concepts, Methods, Studies	8
2.2	Illampu: Framework for Exploring the Design Space	15
2.2.1	Application	16
2.2.2	Implementation	16
2.2.3	Architectures	17
2.2.4	Evaluation	19
2.2.5	Summary	24
3	Illimani: A Development Framework for System-on-Chips	25
3.1	Concept	26

3.2	System Architecture	28
3.3	Hardware and Software Libraries	30
3.4	Evaluation	33
3.5	Implementation	34
3.6	Further Work: Stream-Enabled Data Flow	35
4	Stereo Vision	37
4.1	Principles of Stereo Vision	40
4.1.1	Epipolar Geometry and Rectification	40
4.1.2	Stereo Correspondence	42
4.2	Complexity Analysis of Semi-Global Matching	46
4.3	Quality Analysis and Algorithm-Architecture Co-optimization	50
4.3.1	Methodology and Results	52
4.3.2	Discussion	60
4.4	Conclusion	62
5	Dedicated Architectures for Stereo Matching	63
5.1	Existing Architectures	64
5.2	A novel VLSI architecture	65
5.2.1	Parallelization	65
5.2.2	Systolic Architecture	68
5.3	Performance	70
5.4	Optimization and Analysis for FPGAs	73
5.5	Optimization and Analysis for ASICs	77
5.6	Discussion	80
6	Programmable Architectures for Stereo Matching	85
6.1	Existing Implementations	86
6.2	A novel GPU Implementation	87
6.2.1	Background on the GPU	88
6.2.2	Performance Bounds and Optimization Techniques	89
6.2.3	Parallelization, Implementation, and Optimization	90
6.2.4	Evaluation and Discussion	97
6.3	Instruction Set Extension for RISC Processor Templates	99
6.3.1	Optimization Techniques	100
6.3.2	Instruction Set Extension	101

6.3.3	Results and Evaluation	106
6.4	Architecture Optimizations for VLIW Processor Templates	111
6.4.1	The Moai VLIW Processor	112
6.4.2	VLIW Architecture Optimizations	113
6.4.3	Results and Evaluation	117
6.5	A Generic Multi-Core Implementation	123
6.5.1	Parallelization for Multi-Core Architectures	123
6.5.2	Results and Evaluation	126
6.6	Discussion	129
6.6.1	Application Specific Processor Comparison	130
7	Analysis of Architectural Trade-Offs	133
7.1	Key Results of Examined Architectures	134
7.2	Architecture Application Scenarios	135
7.3	Intra-Architecture Optimization Potential	139
7.4	Inter-Architecture Optimization Potential	142
7.5	Discussion	143
8	Extended Application	149
8.1	A real-world Stereo Vision System	150
8.2	Dynamic Partial Reconfiguration	153
9	Conclusion	159
A	Definition Benchmark Operations	161
References		163