

Contents

Acronyms	xv
1 Introduction	1
1.1 Towards Parallel Systems	2
1.2 Towards Approximate Computing	3
1.3 Motivation	5
1.4 Research Questions	13
1.5 Major Contributions	16
1.6 List of Publications	17
1.7 Thesis Organization	18
2 Background	19
2.1 Approximate Computing	20
2.1.1 Hardware-based Techniques	21
2.1.2 Software-based Techniques	23
2.2 GPUs	25
2.2.1 GPU Architecture	26
2.2.2 GPU Vendors	29
2.2.3 Programming Models	30
2.2.4 Software Optimization Techniques for GPUs	31
2.3 Polyhedral Model	37
2.3.1 Iteration Domain	39
2.3.2 Memory Access Functions	39
2.3.3 Polyhedral Transformations	41
2.3.4 Polyhedral Compilers	41
2.4 Summary	42
3 Related Work	43
3.1 Background in Approximate Computing	44
3.2 Approximate Computing on GPUs	46
3.3 Approximation on Embedded Systems	49

3.4	Compiler-based Approximation	51
3.5	Advances Over the State of the Art	55
3.6	Summary	56
4	Local Memory-Aware Kernel Perforation	57
4.1	Introduction	58
4.2	Overview	60
4.3	Kernel Perforation	61
4.3.1	From Loop to Kernel Perforation	61
4.3.2	Local Input Perforation	63
4.3.3	Perforation Schemes	64
4.3.4	GPU Perforation Schemes and Halos	65
4.4	Reconstruction	67
4.4.1	Reconstruction Techniques	68
4.5	Experimental Evaluation	69
4.5.1	Benchmarks	69
4.5.2	Input Data Sensitivity	70
4.5.3	Parametrization	73
4.5.4	Pareto Optimality	78
4.6	Summary	80
5	Kernel Perforation on Mobile GPUs	83
5.1	Introduction	84
5.2	Kernel Perforation on Mobile GPUs	86
5.3	Experimental Evaluation	87
5.3.1	Experimental Platforms	88
5.3.2	Benchmark Applications	89
5.4	Results	90
5.4.1	Performance on Three Platforms	90
5.4.2	Application Performance on Different Platforms	94
5.4.3	Execution Times	94
5.4.4	Analysis of the Error	96
5.4.5	Interpretation of the Experimental Results	96
5.5	Summary	97

CONTENTS

6	Advanced Techniques for Automatic Loop Nest Perforation	99
6.1	Introduction	100
6.2	ALONA	103
6.2.1	Polyhedral Loop Nest Perforation	104
6.2.2	Polyhedral Model	104
6.2.3	Polyhedral Loop Perforation	105
6.2.4	Extensions to Classical Loop Perforation	107
6.3	Reconstruction	107
6.4	Approximation Space Pruning	110
6.4.1	Approximation Ordering	111
6.4.2	Barvinok Counting Based Pruning	112
6.5	Experimental Evaluation	112
6.5.1	Benchmark Setup	113
6.5.2	Discovered Solutions	114
6.5.3	Reconstruction	115
6.5.4	Evaluation of Perforation Configurations	116
6.5.5	Approximation Space Pruning	117
6.6	Summary	118
7	Conclusions	121
7.1	Summary	122
7.2	Conclusions	124
7.3	Answers to Research Questions	127
7.4	Future Work	129