

---

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	Claims of this Thesis . . . . .	4
1.3	Contributions of this Thesis . . . . .	5
1.4	Thesis Structure . . . . .	6
<b>2</b>	<b>Background</b>	<b>9</b>
2.1	The Emergence of Heterogeneous Computing . . . . .	9
2.1.1	Semiconductor Design - CPU Design Reaching Limits . . . . .	9
2.1.2	From Parallel Computing to Heterogeneous Computing . . . . .	12
2.1.3	Application Diversity . . . . .	14
2.1.4	Diverse Hardware Requires Diverse Software . . . . .	15
2.2	Heterogeneous Computing Today . . . . .	18
2.2.1	Heterogeneous Systems . . . . .	18
2.2.2	Programming Heterogeneous Applications . . . . .	19
2.2.3	Runtime Systems . . . . .	20
2.3	Potentials and Objectives of Heterogeneous Computing . . . . .	21
2.4	Challenges and Limitations of Heterogeneous Computing . . . . .	22
2.5	Chapter Conclusion . . . . .	24
<b>3</b>	<b>Basic Concepts and Ideas</b>	<b>25</b>
3.1	Towards the Efficient Use of Heterogeneous Resources . . . . .	25
3.1.1	Heterogeneous Computing is Fast . . . . .	27
3.1.2	Heterogeneous Computing Saves Energy . . . . .	27
3.1.3	Heterogeneous Scheduling May Achieve More . . . . .	28
3.2	Heterogeneous System Model . . . . .	29
3.2.1	Responsibilities of a Heterogeneous Scheduler . . . . .	30
3.2.2	Operating System Integration of a Heterogeneous Scheduler . . . . .	31
3.3	Heterogeneous Task Model . . . . .	33
3.3.1	Specification of a Task . . . . .	33
3.3.2	Affinity Model . . . . .	35
3.4	Heterogeneous Thread Execution Model . . . . .	37
3.4.1	Completely Fair Scheduling . . . . .	37

3.4.2	Preemptive Multitasking Versus Run-to-completion . . . . .	38
3.4.3	Heterogeneous Task Migration . . . . .	40
3.4.4	Cooperative Multitasking . . . . .	44
3.4.5	Scheduler Evaluation Scenario . . . . .	48
3.5	Related Work . . . . .	49
3.6	Chapter Conclusion . . . . .	51
<b>4</b>	<b>Programming Pattern for Heterogeneous Task Scheduling</b>	<b>53</b>
4.1	Lifecycle of a Task . . . . .	53
4.2	Programming Pattern for Applications . . . . .	55
4.2.1	Generic Application Example . . . . .	55
4.2.2	Checkpoint Example . . . . .	57
4.2.3	Discussion of Practicability . . . . .	59
4.3	Application to a Heterogeneous Scheduler Scenario . . . . .	60
4.3.1	Affinity Information on Tasks . . . . .	60
4.3.2	Delegate Threads . . . . .	61
4.4	Related Work . . . . .	62
4.5	Chapter Conclusion . . . . .	64
<b>5</b>	<b>Algorithms for a Heterogeneous Scheduler</b>	<b>67</b>
5.1	Metrics For Evaluating Schedulers . . . . .	67
5.2	Scheduling Parameters . . . . .	68
5.2.1	CFS Characteristics . . . . .	68
5.2.2	Heterogeneous Parameters . . . . .	70
5.2.3	CPU Load . . . . .	71
5.2.4	Fairness . . . . .	76
5.2.5	Hardware Parameters . . . . .	77
5.3	Scheduling Policies . . . . .	77
5.3.1	Task Granularity . . . . .	79
5.3.2	Task registration . . . . .	80
5.3.3	Load Balancing Policies . . . . .	81
5.4	Related Work . . . . .	84
5.5	Chapter Conclusion . . . . .	86
<b>6</b>	<b>Thread Level Acceleration</b>	<b>89</b>
6.1	Towards a Heterogeneous Task Scheduler . . . . .	89
6.2	Kernel-Space Task-Scheduling Framework . . . . .	90
6.2.1	Data Structures . . . . .	92
6.2.2	Scheduler API . . . . .	94
6.2.3	Control API . . . . .	96
6.2.4	Scheduler Interface . . . . .	97
6.2.5	Experimental Setup . . . . .	97
6.2.6	Experimental Results . . . . .	98
6.3	User-Space Task-Scheduling Framework . . . . .	103
6.3.1	Differences to the Kernel-Space Task-Scheduling Framework . . . . .	103

6.3.2	Scheduler Framework Architecture . . . . .	105
6.3.3	Scheduler Library . . . . .	106
6.3.4	Scheduler Implementation . . . . .	107
6.3.5	Experimental Setup . . . . .	110
6.3.6	Experimental Results . . . . .	113
6.4	Comparison of Presented Scheduling Frameworks . . . . .	125
6.5	Related Work . . . . .	125
6.6	Chapter Conclusion . . . . .	127
<b>7</b>	<b>Function Level Acceleration</b>	<b>129</b>
7.1	Function Execution Model . . . . .	129
7.2	Approaches to Shared Library Interposing . . . . .	132
7.2.1	Dynamic Linker Configuration . . . . .	133
7.2.2	Exchanging the Dynamic Linker . . . . .	134
7.2.3	Static Binary Infection . . . . .	134
7.2.4	Binary Instrumentation with Pin . . . . .	135
7.2.5	Summary . . . . .	136
7.3	Library-Interposing Framework for Transparent Application Acceleration .	136
7.3.1	Accelerating BLAS Libraries . . . . .	137
7.3.2	Plugins for Shared Library Interposing . . . . .	137
7.3.3	Selector Policy . . . . .	139
7.3.4	Profiling Support . . . . .	141
7.3.5	Automatic Plugin Generation Using XSLT . . . . .	142
7.4	Liftracc Experimental Results . . . . .	143
7.4.1	Runtimes . . . . .	143
7.4.2	Overheads . . . . .	145
7.5	Related Work . . . . .	146
7.6	Chapter Conclusion . . . . .	148
<b>8</b>	<b>Summary and Outlook</b>	<b>151</b>
8.1	Contributions . . . . .	151
8.2	Conclusions and Lessons Learned . . . . .	154
8.3	Future Directions . . . . .	156
	<b>Acronyms</b>	<b>159</b>
	<b>List of Figures</b>	<b>161</b>
	<b>List of Tables</b>	<b>163</b>
	<b>Listings</b>	<b>165</b>
	<b>Author's Publications</b>	<b>167</b>
	<b>Bibliography</b>	<b>169</b>