

Contents

Preface	1
1 Introduction	7
1.1 Programming Models	7
1.2 Programming Parallel Architectures	7
1.3 A Metaphor for Analysis and Synthesis	9
1.4 The Need for Autotuning	13
1.5 The Need of Multi-objective Tuning	15
1.6 Motivations	15
1.7 Objectives	17
1.8 Outline	17
2 Program Analysis	19
Contribution 1: SSA and Type-inference in PHP 7.1	21
Contribution 2: CFG Vectorization on ARM Processors	32
Contribution 3: Limit of C++ Programming Models for Vectorization	42
3 High-level Programming Models	51
Contribution 4: libWater	53
Contribution 5: libWater on large-scale cluster	64
Contribution 6: CELERITY	76
4 Domain-Specific Languages	89
Contribution 7: OpenABL: A DSL for Agent-based Simulations	91
Contribution 8: OpenABL on Large-Scale Simulations	105
Contribution 9: OpenABL on Complex Hierarchical Models	131
5 Automatic Tuning	143
Contribution 10: Automatic Heterogeneous Task Partitioning	146
Contribution 11: Vectorization Cost Modeling	157
Contribution 12: Autotuning for Stencil Computations	159
Contribution 13: Multi-objective Tuning for Approximate Computing	169
Contribution 14: Multi-objective Tuning for Energy Consumption	179
6 Conclusions	189
6.1 Summary of the Results	189
6.2 Open Questions and Future Work	192
Bibliography	193