

# Contents

<b>Acknowledgements</b>	<b>i</b>
<b>Abstract (English/Français)</b>	<b>v</b>
<b>List of figures</b>	<b>xiii</b>
<b>List of tables</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 High-Throughput and Efficient Data Lookups . . . . .	4
1.1.1 Hiding Memory Access Latency . . . . .	4
1.1.2 Specialized Hardware . . . . .	6
1.2 Thesis and Dissertation Goals . . . . .	7
1.3 Asynchronous Memory Access Chaining . . . . .	7
1.4 On-chip Accelerator for Index Traversals . . . . .	8
1.5 Memory Subsystem Bottlenecks . . . . .	9
1.6 Contributions . . . . .	9
<b>2 Background</b>	<b>13</b>
2.1 DBMS Basics . . . . .	13
2.2 Pointer-Intensive Data Structures in Database Systems . . . . .	14
2.2.1 Hash Tables . . . . .	14
2.2.2 Tree Search . . . . .	16
	<b>ix</b>

<b>3</b>	<b>AMAC: Asynchronous Memory Access Chaining</b>	<b>19</b>
3.1	Hiding Memory Access Latency . . . . .	20
3.1.1	Software Prefetching Techniques for Pointer-Chasing Database Operations	20
3.1.2	Performance Analysis of Software Prefetching . . . . .	23
3.2	Asynchronous Memory Access Chaining . . . . .	25
3.2.1	Design Overview . . . . .	26
3.2.2	Handling Read/Write Dependencies . . . . .	30
3.3	Methodology . . . . .	31
3.4	Evaluation . . . . .	33
3.4.1	Hash Join . . . . .	33
3.4.2	Scalability Analysis . . . . .	37
3.4.3	Group-By . . . . .	39
3.4.4	Tree Search . . . . .	41
3.5	Discussion . . . . .	42
3.6	AMAC Summary . . . . .	43
<b>4</b>	<b>Widx: On-chip Accelerator for Index Traversals</b>	<b>45</b>
4.1	Profiling Analysis of a Modern DBMS . . . . .	46
4.2	Database Index Traversal Acceleration . . . . .	48
4.2.1	Overview . . . . .	48
4.2.2	First-Order Performance Model . . . . .	50
4.3	Widx . . . . .	55
4.3.1	Architecture Overview . . . . .	55
4.3.2	Programming API . . . . .	56
4.3.3	Additional Details . . . . .	57
4.4	Methodology . . . . .	59
4.5	Evaluation . . . . .	61
4.5.1	Performance on Hash Join Kernel . . . . .	62
4.5.2	Case study on MonetDB . . . . .	64

4.5.3 Area and Energy Efficiency . . . . .	66
4.6 Discussion . . . . .	68
4.7 Widx Summary . . . . .	69
<b>5 Quantifying the Impact of Memory Subsystem on Acceleration</b>	<b>71</b>
5.1 Quantifying the Overhead . . . . .	72
5.1.1 Experimental Setup . . . . .	73
5.2 Data Structure Padding . . . . .	74
5.3 Virtual Memory Page Size . . . . .	74
5.4 Miss Status Handling Registers . . . . .	75
5.5 Throughput Scalability and Bottleneck Analysis . . . . .	76
5.6 Putting It All Together . . . . .	78
5.7 Combining Hardware and Software Acceleration . . . . .	81
5.8 Discussion . . . . .	83
<b>6 Related Work</b>	<b>85</b>
6.1 Software and Hardware Prefetching . . . . .	85
6.2 Hardware-Conscious Algorithms . . . . .	86
6.3 Decoupled Architectures . . . . .	86
6.4 Specialized Hardware and Accelerators . . . . .	87
<b>7 Concluding Remarks</b>	<b>89</b>
7.1 Future Directions . . . . .	90
<b>Bibliography</b>	<b>93</b>
<b>Curriculum Vitae</b>	<b>103</b>