

# Contents

<b>Contents</b>	<b>13</b>
<b>1 Introduction</b>	<b>15</b>
1.1 The Semantic Web . . . . .	15
1.2 RDF data management . . . . .	16
1.3 RDF and graph benchmarks . . . . .	23
1.4 Thesis Outline and Contributions . . . . .	24
<b>2 Background and Related Work</b>	<b>27</b>
2.1 Semantic Web Technologies . . . . .	27
2.2 RDF storage . . . . .	42
<b>3 Deriving an Emergent Relational Schema from RDF Data</b>	<b>59</b>
3.1 Introduction . . . . .	59
3.2 Emerging A Relational Schema . . . . .	61
3.3 Experimental Evaluation . . . . .	76
3.4 Related Work . . . . .	82
3.5 Conclusions . . . . .	83
<b>4 Exploiting Emergent Schemas to make RDF systems more efficient</b>	<b>85</b>
4.1 Emergent Schema Introduction . . . . .	85
4.2 Emergent Schema Aware RDF Storage . . . . .	87
4.3 Emergent Schema Aware SPARQL Optimization . . . . .	92
4.4 Emergent Schema Aware SPARQL Execution . . . . .	93
4.5 Performance Evaluation . . . . .	98
4.6 Related Work . . . . .	100
4.7 Conclusion . . . . .	102
<b>5 Benchmarking RDF stores</b>	<b>105</b>
5.1 S3G2: A Scalable Structure-correlated Social Graph Generator . .	105
5.2 LDBC Social Network Benchmark (SNB) . . . . .	121
<b>6 Conclusions</b>	<b>125</b>
6.1 Contributions . . . . .	125
6.2 Future research directions . . . . .	131
6.3 Summary . . . . .	132

<b>List of Figures</b>	<b>133</b>
<b>List of Tables</b>	<b>134</b>
<b>A Query plan transformation for star pattern</b>	<b>135</b>
<b>B DBpedia queries</b>	<b>139</b>
<b>C LDBC Datagen</b>	<b>143</b>
<b>Bibliography</b>	<b>149</b>