



Semantic Web Programming

John Hebel
Matthew Fisher
Ryan Blace
Andrew Perez-Lopez



WILEY

Wiley Publishing, Inc.



Contents

Foreword	xxiii	
Introduction	xxv	
Part One	Introducing Semantic Web Programming	1
Chapter 1	Preparing to Program a Semantic Web of Data	3
	Defining the Semantic Web	4
	Identifying the Major Programming Components	10
	Determining Impacts on Programming	13
	Establishing a Web Data-Centric Perspective	13
	Expressing Semantic Data	14
	Sharing Data	16
	Making Data Dynamic and Flexible	18
	Avoiding the Roadblocks, Myths, and Hype	19
	Semantic Web Roadblocks	19
	Semantic Web Myths	21
	Semantic Web Hype	22
	Understanding Semantic Web Origins	23
	Exploring Semantic Web Examples	26
	Semantic Wikis (<i>semantic-mediawiki.org</i>)	26
	Twine (<i>www.twine.com</i>)	27
	The FOAF Project (<i>www.foaf-project.org</i>)	28
	RDFa and Microformats	30
	Semantic Query Endpoint (<i>dbpedia.org/sparql</i>)	32
	Semantic Search (<i>www.trueknowledge.com</i>)	32
	Summary and Onward	34
	Notes	34

Chapter 2	Hello Semantic Web World	35
	Setting Up Your Semantic Web Development Environment	36
	Programming the Hello Semantic Web World Application	38
	Summary	58
Part Two	Foundations of Semantic Web Programming	61
Chapter 3	Modeling Information	63
	Modeling Information in Software	64
	Sharing Information: Syntax and Semantics	65
	Serialized Objects	66
	Relational Databases	66
	Extensible Markup Language (XML)	66
	Metadata and Data in Information Sharing	67
	The Semantic Web Information Model: The Resource Description Framework (RDF)	68
	Nodes: Resources and Literals	69
	Edges: Predicates	71
	Exchanging Information with RDF	72
	Statements as Points	73
	RDF Serializations	74
	RDF/XML	74
	Terse RDF Triple Language (Turtle)	78
	N-Triples	81
	Quick Hack	82
	More RDF	84
	Blank Nodes	84
	Reification	88
	RDF Organizational Constructs	88
	Summary	91
Chapter 4	Incorporating Semantics	93
	Semantics on the Web	94
	Motivating Factors	94
	Understanding the World Wide Web	95
	Knowledge Domain Integration	97
	Expressing Semantics in RDF	98
	Vocabularies, Taxonomies, and Ontologies	99
	A Vocabulary Language for RDF	100
	An Ontology Language for the Web	101
	Introduction to Ontologies	102
	Distributed Knowledge	102
	Open World Assumption	103
	No Unique Names Assumption	104
	Overview of Ontology Elements	104

Ontology Header	105
Classes and Individuals	105
Properties	106
Annotations	106
Data types	106
Elements of an Ontology	107
OWL 2 Typing	107
Ontology Header	108
Annotations	109
Basic Classification	110
Classes and Individuals	110
<i>rdfs:SubClassOf</i>	111
Instance versus Subclass	112
<i>owl:Thing</i> and <i>owl:Nothing</i>	113
Defining and Using Properties	113
Property Domain and Range	114
Describing Properties	115
<i>rdfs:subPropertyOf</i>	115
Top and Bottom Properties	116
Inverse Properties	116
Disjoint Properties	117
Property Chains	118
Symmetric, Reflexive, and Transitive Properties	119
Functional and Inverse Functional Properties	120
Keys	121
Datatypes	122
Data type Restrictions	122
Defining Datatypes in Terms of Other Datatypes	124
Negative Property Assertions	126
Property Restrictions	127
Value Restrictions	127
Cardinality Restrictions	130
Qualified Cardinality Restrictions	132
Advanced Class Description	134
Enumerating Class Membership	134
Set Operators	134
Disjoint Classes	136
Equivalence in OWL	137
Equivalence among Individuals	138
Equivalence among Classes and Properties	139
Summary	139
Chapter 5 Modeling Knowledge in the Real World	141
Exploring the Components of the Semantic Web	141
Semantic Web Frameworks	143

Storing and Retrieving RDF	144
RDF Store Implementations	144
Retrieving Information in a Knowledgebase	146
Realizing the Semantics of OWL	147
Understanding Forward Chaining Inference	148
Understanding Backward Chaining Inference	152
Choosing the Right Inference Method	153
Common Frameworks and Components	153
RDF Store Implementations	154
Retrieval Components	156
Reasoning Engines	156
Knowledgebase Performance	157
Exploring the Profiles of OWL	158
OWL Full and OWL DL	159
The Profiles of OWL	160
OWL EL	160
OWL QL	161
OWL RL	162
Demonstrating OWL Inference	163
The Ontology	163
The Example Application	165
The Results	169
Performing No Inference	169
Performing RDFS Inference	170
Performing OWL Inference	171
Working with Ontologies	172
Decoupling the Knowledge Model from the Application	173
Sharing across Domain and Application Boundaries	174
What Is a Foundational Ontology?	175
Common Foundational Ontologies	177
BFO	177
Cyc and OpenCyc	178
DOLCE	179
SUMO	179
Dublin Core Metadata Initiative	179
FOAF	180
GeoRSS and OWL-Time	180
Finding Ontologies to Reuse or Extend	181
Choosing the Right Foundational Ontologies	183
Summary	184
Chapter 6 Discovering Information	185
Navigating the Semantic Web	186
Searching the Semantic Web	190

Querying the Semantic Web	192
Quickstart with SPARQL	192
Four Foundational Query Forms	196
SELECT Essentials	197
RDF Datasets, FROM and FROM NAMED	202
Query Modifiers	208
DISTINCT	209
REDUCED	210
ORDER BY	210
Data Streaming with OFFSET and LIMIT	211
Flexible Querying with FILTER and OPTIONAL	213
FILTER	213
OPTIONAL	215
UNION	219
CONSTRUCT Essentials	222
DESCRIBE Essentials	224
ASK Essentials	225
SPARQL Entailment	226
Unsupported Functionality	228
Data Modification	228
Subqueries	228
Aggregation	228
Summary	229
Chapter 7 Adding Rules	231
What Are Rules?	232
Reasons for Rules	232
No Support for Property Composition	232
Use of Built-ins	232
Ontological Mediation	233
Limiting Assumptions	233
Rule Languages	233
SWRL Essentials	234
The Abstract Syntax	235
The XML Concrete Syntax	237
var	238
imp	239
_rlab	239
_body	239
_head	240
classAtom	240
datarangeAtom	241
individualPropertyAtom	241
datavaluedPropertyAtom	242
sameIndividualAtom	242

	differentIndividualsAtom	243
	builtinAtom	243
	The RDF Concrete Syntax	243
	Built-ins	244
	Examples	244
	DL-Safe Rules	245
	Ontological Mediation	252
	Mapping Friends without Upsetting Any of Them	253
	The Power of Rules	256
	Jena Rules	257
	Rule Interchange Format	259
	Delving into the Details	260
	The Future of RIF	260
	Summary	261
Part Three	Building Semantic Web Applications	263
Chapter 8	Applying a Programming Framework	265
	Framing the Semantic Web	266
	The Jena Semantic Web Framework	269
	Defining Jena Programming Concepts	269
	Programming with Jena	273
	Establishing the Jena Development Environment	276
	Establishing the Knowledgebase: Setting Up the Model	276
	Populating the Model with Semantic Web Data	279
	Combining Semantic Web Data	282
	Interrogating Semantic Web Data	283
	Reasoning across Semantic Web Data	285
	Exporting Semantic Web Data	290
	Deallocating Semantic Web Data Resources	290
	Managing Semantic Web Data	291
	Getting Information Regarding Your Semantic Web Data	291
	Generating Events Based on Semantic Web Data	292
	Dealing with Concurrency and Your Semantic Web Data	293
	Customizing the Jena Framework	295
	Serializing Semantic Web Data	297
	Common App Overview: FriendTracker	298
	Summary	299
Chapter 9	Combining Information	301
	Combining Information	303
	Representing Information	303
	Translating between Representations	304
	Addressing the Challenges of Translation	305

Maintaining Fidelity	305
Tracking Provenance Information	305
Reversing the Process	306
Handling Varying Data	306
Managing Data Volume	306
Introducing the FriendTracker Data Sources	307
Facebook XML Web Service	307
Jabber Java Client	308
Upcoming.org XML Web Service	308
WordPress Relational Database	308
Exposing XML-Based Web Services as RDF	309
Introducing the Weather.gov XML Feed	310
Exposing XML Using XSL Transformations	311
Traversing XML Documents with XPath	313
Applying XSLT to a Simple Example	315
Processing XML and XSLT Programmatically	319
Applying XSLT to the Facebook Data Source	322
Weighing the Benefits and the Costs of XSLT	326
Exposing XML Using XML Bindings and Velocity	328
Generating Java Bindings for XML Data	329
Unmarshalling XML Data into Java	331
Introducing the Velocity Template Engine	333
Generating RDF with Velocity	336
Weighing the Benefits and the Costs	336
Exposing Relational Databases as RDF	337
Exposing a WordPress Blog Using D2RQ	338
Creating D2RQ Mappings for the WordPress Database	339
Wrapping the D2RQ Instance in a Jena Model	341
Querying the D2RQ Exposed WordPress Database	342
Weighing the Benefits and the Costs of D2RQ	345
Exposing Other Sources of Data	346
Exposing Jabber with a Custom Streaming RDF Writer	346
Exposing Java Objects Using Reflection	352
Applying the RDF Generator to the Weather.gov XML Feed	356
Applying the RDF Generator to the Upcoming.org XML Feed	358
Summary	359
Chapter 10 Aligning Information	361
Data Source, Domain, and Application Ontologies	361
Aligning Ontologies	362
Ontology Constructs	363
Translation via Rules	365

Explicit Translation	365
Ad Hoc Approaches to Translation	366
FriendTracker	366
Aligning Ontologies with OWL and SWRL	372
Aligning Ontologies with XSLT	376
Aligning Ontologies with Code	381
Aligning Simple Ontologies with RDFS	382
Record Linkage	385
Summary	388
Chapter 11 Sharing Information	389
Microformats	390
eRDF	392
RDFa	395
Supported Attributes	396
xmlns	396
rel	396
rev	398
content	398
href	399
src	399
about	399
property	399
resource	400
datatype	401
typeof	401
Blank Nodes	402
Language Support	403
Tools and Frameworks	404
RDF Transformational Tools	404
SPARQL Endpoints	404
Joseki Installation and Operation	405
xOperator	408
Installation and Operation	409
Example Query	410
FriendTracker in RDFa	411
Summary	417
Part Four Expanding Semantic Web Programming	419
Chapter 12 Developing and Using Semantic Services	421
Background	422
Discovery	424
Invocation	424
Negotiation	425

Error Handling	425
Monitoring	425
Composition	425
Implementing Semantic Services	426
Semantic Markup for Web Services	427
ServiceProfile	428
ServiceModel	428
ServiceGrounding	428
Web Service Modeling Ontology	429
Semantic Annotations for WSDL	432
SAWSDL Example	433
SAWSDL Tools	434
Summary	435
Chapter 13 Managing Space and Time	437
Space and Time in Software	437
Spatial Information	438
Temporal Information	440
Representing Spatiotemporal Data on the Semantic Web	441
Spatial and Temporal Software with Jena	450
Working with Spatial Data	452
Example: Spatial Queries	453
Framing the Problem	453
Approach and Rationale	454
Components	454
Example: Transaction Time-Bounded Queries	461
Framing the Problem	461
Approach and Rationale	462
Components	462
Summary	465
Notes	465
Chapter 14 Semantic Web Patterns and Best Practices	467
Aggregating Disparate Data Sources	468
Exposing Data Sources as RDF	469
Bringing Data into the Domain Knowledge Model	470
Storing Information in the Knowledgebase	472
Initiating the Flow of Data	472
Annotating Unstructured Data	472
Annotation Management	474
Ontology Management	475
Unstructured Data Sources and the Client Application	476
Coordinating Semantic Services	478
Applying Semantic Web Best Practices	480

Creating URIs	480
Making URIs Unique	480
Making URIs Consistent	481
Making URIs Resolvable	483
Specifying Units of Measurement	483
Unit-Specific Properties and Datatypes	484
Statement Reification	485
Value Containers	485
Representing N-ary Relationships	486
Managing Bad Data	487
Summary	487
Chapter 15 Moving Forward	489
Advancing Ontologies	491
Ontology Repositories and Registries	491
Linked Data	492
Versioning	493
Ontology Metrics	495
Advancing Integration	497
Semantic Pipes	497
Distributed Queries	498
Alignment	500
Advancing Reasoning	506
Rule Interchange Format (RIF)	506
Probabilistic Reasoning	507
Trust: Proof Markup Language	510
LarKC: The Large Knowledge Collider	512
Advancing Visualization	513
Summary	516
Appendix A RDF	519
Appendix B The OWL Web Ontology Language	525
Appendix C SWRL	533
Appendix D SPARQL	555
Appendix E Jena Reference Guide	567
Appendix F Installation Reference Guide	577
Index	585