

Dragan Gašević
Dragan Djurić
Vladan Devedžić

Model Driven Architecture and Ontology Development

 Springer

Model Driven Architecture and Ontology Development

Dragan Gašević · Dragan Djurić
Vladan Devedžić

Model Driven Architecture and Ontology Development

Foreword by Bran Selic

With 153 Figures and 7 Tables

 Springer

Authors

Dragan Gašević
School of Interactive Arts
and Technology
Simon Fraser University Surrey
13450 102 Ave.
Surrey, BC V3T 5X3
Canada
dgasevic@sfu.ca
and dgasevic@acm.org

Dragan Djurić
Vladan Devedžić
FON – School of
Business Administration
University of Belgrade
Jove Ilića 154
11000 Belgrade
Serbia and Montenegro
dragandj@gmail.com
devedzic@fon.bg.ac.yu

Library of Congress Control Number: 2006924633

ACM Computing Classification (1998): D.2.10, D.2.11, H.4.0, I.2.4

ISBN-10 3-540-32180-2 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-32180-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typeset by the Authors
Production: LE-TeX Jelonek, Schmidt & Vöckler GbR, Leipzig
Cover design: KünkellOpka Werbeagentur, Heidelberg

Printed on acid-free paper 45/3100/YL - 5 4 3 2 1 0

To our families

Foreword

The first time I paid attention to the term “ontology” was in the late 1980s when I was part of an engineering team that was responsible for defining what we would now call a domain-specific modeling language. In our case, the domain was telecommunications software and the purpose of our language was to give system architects the ability to describe the high-level structure of their software in the most direct and most expressive manner possible.

The team members were all experienced designers with deep knowledge of the domain so that we had no trouble putting together the initial list of key language concepts. We knew that we needed to include standard architectural modeling constructs such as components, ports, connectors, and the like. We also wanted our language to be object-oriented, so notions such as class, objects, and inheritance were added to the list. However, soon after this very promising start, all progress ground to a halt. Somehow, the definition of the seemingly trivial fine-grain details of these constructs kept eluding us despite long, passionate, and occasionally acrimonious discussions that can only be compared to medieval theological debates.

It was our good fortune that at that point we met Professor Doug Skuce of the University of Ottawa. He had a method and a tool that helped us develop an explicit ontology for our domain. From that exercise we learned that our difficulties stemmed from the fact that, although we shared a general intuition for the chosen constructs of our language, there were numerous subtle and *unstated* differences in our individual conceptualizations that were a barrier to mutual understanding. Furthermore, we discovered that certain commonly used terms had multiple meanings – all equally valid – but which we had not differentiated adequately, leading to much confusion. Only after we had defined our ontology, which included semi-formal definitions of all key terms and their relationships, were we able to finish our task successfully.

Ever since I’ve felt that defining a formal domain ontology is a useful and often necessary step in almost any software project. This is because software deals principally with ideas rather than physical artifacts. Whereas the nature of physical artifacts is generally self-evident, this is not

the case with conceptual entities, which are products of the mind. As we all know, different minds see the same thing differently.

The definition and application of ontologies for developing software systems is a central theme of this book. However, the book is about much more than that. It explains, in a clear and didactic manner, how a variety of recent buzzword developments in software theory and practice (intelligent agents, Model Driven Architecture, metamodeling, etc.) can be combined, and brings us to the threshold of the next step in the evolution of the World Wide Web: the *Semantic Web*. Like the Internet before it, the Semantic Web promises to introduce a significant and qualitatively new phenomenon into our lives. This is because it endows the network of disparate information that is currently accessible on the Internet with *meaning*. Because this meaning can be gleaned and processed *automatically* by software, the Semantic Web opens up the exciting and awe-inducing possibility of a unified global intelligence accessible to all.

In the first half of the book, the authors navigate deftly through the prolific and highly confusing *gemischt* of technologies, tools, and standards such as XML, RDF, OWL, MDA, or UML and explain how they relate to each other in the context of the idea of the Semantic Web. They introduce the notion of *modeling spaces*, which provides a conceptually simple yet comprehensive framework for understanding and addressing issues within the domain considered. Using that framework, the second half of the book describes a practical strategy for realizing key elements of the Semantic Web based on existing industry standards.

The book is equally suited to those who merely want to be informed of the relevant technological landscape, to practitioners dealing with concrete problems, and to researchers seeking pointers to potentially fruitful areas of research. The writing is technical yet clear and accessible, and is illustrated throughout with useful and easily digestible examples.

I would also highly recommend this book to sociologists studying the interplay between society and technology. It clearly demonstrates that the core technologies required for constructing the Semantic Web are available and moving forward inexorably. Society must be prepared to deal with something so ripe with potential. We must understand not only how the Semantic Web can be useful but also what dangers lurk within it.

Ottawa, Canada

Bran Selic

December 2005

Preface

The idea of ontologies emerged in applied artificial intelligence some time ago as a means for sharing knowledge [Gruber, 1993]. Following the development of ontologies and related Web technologies (e.g., HTML and XML), Tim Berners-Lee, Jim Hendler, and Ora Lassila envisioned the next generation of the Web, called the Semantic Web [Berners-Lee et al., 2001]. Being based on ontologies, the Semantic Web has the potential for semantically richer representations of things (e.g., Web pages, applications, and persons) and their relations on the Web, and thus should provide us with more intelligent services. That idea might have initially sounded very futuristic and too enthusiastic, but it has recruited a lot of important players from both academia and industry into very extensive and well-funded research efforts. Today, we have quite impressive results, manifested by standards that have been adopted (RDF and OWL), development frameworks (Jena), best-practice and deployment recommendations, and many applications (e.g. PiggyBank).

Of course, researchers are still facing many challenges in their efforts to accomplish the full vision of the Semantic Web. Probably the first and most important goal is to persuade many industrial developers and software engineers to use and develop ontologies in their everyday practice. However, ontologies rely on well-defined and semantically powerful concepts in artificial intelligence such as description logics, reasoning, and rule-based systems. Since software engineers are largely unfamiliar with these concepts, ontologies have a price that must be paid for the benefits that they provide.

Trying to address the above problems, researchers have started exploring the potential of some widely adopted software engineering tools and methodologies for ontology development. Stephen Cranefield did the pioneering research by proposing that UML, a well-known software modeling language, should be used for ontology development [Cranefield, 2001a]. After him, several researchers have explored further the similarities, differences, and equivalences between UML and ontology languages, as well as the potential of the most recent software engineering initiative called the Model Driven Architecture (MDA), and its accompanying standards (the Meta-Object Facility (MOF) and XML

Metadata Interchange (XMI)) for ontology development [Baclawski et al., 2002a; Djurić et al., 2005a; Falkovych et al., 2003]. This resulted in the initiation of a process for adopting an MDA-based ontology standard by the Object Management Group (OMG), a software engineering standardization consortium [OMG ODM RFP, 2003]. The standard is intended to define the Ontology Definition Metamodel (ODM) using the MOF (used for specifying UML as well), a UML extension (the Ontology UML Profile, or OUP) to allow UML tools to be used to fully develop ontologies, and a set of transformations between the ODM, the OUP, UML, and Semantic Web ontology languages (e.g., RDF(S) and OWL). When completed, the ODM specification is expected to be in the form of an OMG language, like UML and CWM.

In this book we try to fill the gap in the literature covering the subject of applications of the MDA for ontology development on the Semantic Web. Other books cover either the MDA initiative [Kleppe et al., 2003; Mellor et al., 2003b] or the Semantic Web (i.e., ontology development) [Fensel, 2004; Stuckenschmidt & van Harmelen, 2005; Zhong et al., 2003] only. This book gives a comprehensive overview of both themes, with the main emphasis on how we can employ MDA-related standards to develop Semantic Web ontologies. The book is closely related to the recent OMG initiative for the ODM. The book is the first description of that new language.

The book is based on our experience obtained from a series of tutorials entitled “MDA Standards for Ontology Development” that we have given at several international conferences on the Semantic Web (the International Semantic Web Conference and the European Semantic Web Conference) and on software engineering (the International UML Conference and the International Conference of Web Engineering).

Organization and Structure

The book is divided into three parts. Part I covers the basics of both the main topics – ontologies and the MDA. First, Chap. 1 gives a brief overview of the field of knowledge representation in artificial intelligence. Chapters 2 and 3 introduce the main concepts of ontologies, the Semantic Web, standards, applications, tools, and some open research questions. Next, Chap. 4 explains the Model Driven Architecture, and its main standards (the MOF and XMI) and mechanisms (UML profiles). Part I is concluded by Chap. 5 with modeling spaces, a conceptual framework

defined to provide an easier understanding of approaches to modeling such as ontologies and MOF-defined modeling languages (UML and the ODM).

Part II is the central part of the book. It starts with Chap. 6, which presents a comprehensive review of several approaches and tools that aim to bridge the gap between ontology development and software engineering methodologies. This chapter also lists the relations between UML and ontology languages. Chapter 7 explains the motivation for the forthcoming OMG ontology development standard for the ODM, and the requirements the standard has to fulfill. Next, Chaps. 8 and 9 describe the current specifications of the OMD and the Ontology UML Profile, respectively. Finally, Chap. 10 analyzes the mappings between MDA-based languages (the ODM and the OUP) and Semantic Web ontology languages.

Part III is dedicated to applications that will support the practical use of languages that conform to the OMG ontology development standard, and to some practical aspects of how to develop ontologies using those MDA-based languages. First, Chap. 11 is a short tutorial showing how to develop ontologies using the OUP in two state-of-the-art UML tools (MagicDraw and Poseidon for UML). Chapter 12 describes an implementation of an ontology-building platform called AIR, developed entirely following MDA principles. Finally, Chap. 13 discusses two examples of ontologies developed using the OUP and MDA standards.

Throughout the book, we use many ontologies, UML and other MDA-based models, and transformations between them. In order to allow you to try them out and use them in practice, we have created a Web page containing supplementary resources. You can reach this Web page at <http://www.modelingspaces.org>. Besides the resources referred to in the book, this Web page contains the slide handouts of the tutorials that we have given at many international conferences.

Acknowledgments

Some of the content of this book has been previously published by organizations that are not part of the Springer group. We thank these organizations for their kind permission to use the material for this book, in particular

- IEEE Press
- Rinton Press
- Journal of Educational Technology & Society, and
- Journal of Object Technology.

We are very grateful to Jelena Jovanović, Nenad Krdžavac, and Colin Knight for a huge amount of help in preparing this book and for giving us many useful comments and suggestions. We would also like to thank all members of the GOOD OLD AI research group (<http://goodoldai.org.yu>) at the University of Belgrade, Serbia and Montenegro for providing us with a great environment and many useful ideas that enormously improved the quality of the book. Dragan Gašević is very grateful to Marek Hatala, Tom Calvert, Griff Richards, and other colleagues from the LORNET research network at Simon Fraser University in Vancouver, Canada, for their great support and understanding.

We thank Bran Selic of IBM Rational Software for agreeing to write the book's foreword. We are also very grateful to Jean Bézivin of the University of Nantes whose research in the field of model-driven engineering inspired us to start exploring relations between MDA and ontologies. A special mention should be given to Ralf Gerstner of Springer-Verlag, who invited us to write this book and had a lot of patience in working with us. We are very thankful to Evan Wallace, Elisa Kendall, Bob Colomb, Anna Gerber, Dan Chang, Lewis Hart, and other colleagues from the OMG's Ontology PSIG for giving us up-to-date information about the OMG's ODM standardization efforts in the last three years.

Vancouver, Canada
Belgrade, Serbia and Montenegro
Belgrade, Serbia and Montenegro

Dragan Gašević
Dragan Djurić
Vladan Devedžić

March 2006

Contents

Part I Basics

1. Knowledge Representation.....	3
1.1 Basic Concepts	4
1.2 Cognitive Science	7
1.3 Types of Human Knowledge	11
1.4 Knowledge Representation Techniques.....	14
1.4.1 Object–Attribute–Value Triplets	15
1.4.2 Uncertain Facts	15
1.4.3 Fuzzy Facts	16
1.4.4 Rules	17
1.4.5 Semantic networks.....	18
1.4.6 Frames	19
1.5 Knowledge Representation Languages.....	19
1.5.1 Logic-Based Representation Languages	20
1.5.2 Frame-Based Representation Languages	27
1.5.3 Rule-Based Representation Languages.....	29
1.5.4 Visual Languages for Knowledge Representation	32
1.5.5 Natural Languages and Knowledge Representation	35
1.6 Knowledge Engineering	36
1.7 Open Knowledge Base Connectivity (OKBC)	39
1.8 The Knowledge Level	41
2. Ontologies	45
2.1 Basic Concepts	46
2.1.1 Definitions	46
2.1.2 What Do Ontologies Look Like?	48
2.1.3 Why Ontologies?	50
2.1.4 Key Application Areas.....	55
2.1.5 Examples.....	57
2.2 Ontological Engineering.....	58
2.2.1 Ontology Development Tools.....	58
2.2.2 Ontology Development Methodologies	65
2.3 Applications.....	69
2.3.1 Magpie	69
2.3.2 Briefing Associate.....	70

2.3.3	Quickstep and Foxtrot.....	71
2.4	Advanced Topics.....	72
2.4.1	Metadata, Metamodeling, and Ontologies.....	72
2.4.2	Standard Upper Ontology.....	74
2.4.3	Ontological Level.....	76
3.	The Semantic Web.....	79
3.1	Rationale.....	80
3.2	Semantic Web Languages.....	81
3.2.1	XML and XML Schema.....	81
3.2.2	RDF and RDF Schema.....	84
3.2.3	DAML+OIL.....	87
3.2.4	OWL.....	90
3.2.5	SPARQL.....	92
3.3	The Role of Ontologies.....	95
3.4	Semantic Markup.....	96
3.5	Semantic Web Services.....	100
3.6	Open Issues.....	104
3.7	Quotations.....	107
4.	The Model Driven Architecture (MDA).....	109
4.1	Models and Metamodels.....	109
4.2	Platform-Independent Models.....	110
4.3	Four-Layer Architecture.....	112
4.4	The Meta-Object Facility.....	114
4.5	Specific MDA Metamodels.....	117
4.5.1	Unified Modeling Language.....	117
4.5.2	Common Warehouse Metamodel (CWM).....	118
4.5.3	Ontology Definition Metamodel.....	119
4.6	UML Profiles.....	120
4.6.1	Examples of UML Profiles.....	121
4.7	An XML for Sharing MDA Artifacts.....	123
4.8	The Need for Modeling Spaces.....	126
5.	Modeling Spaces.....	127
5.1	Modeling the Real World.....	128
5.2	The Real World, Models, and Metamodels.....	129
5.3	The Essentials of Modeling Spaces.....	131
5.4	Modeling Spaces Illuminated.....	134
5.5	A Touch of RDF(S) and MOF Modeling Spaces.....	137
5.6	A Touch of the Semantic Web and MDA Technical Spaces.....	139
5.7	Instead of Conclusions.....	141

Part II The Model Driven Architecture and Ontologies

6. Software Engineering Approaches to Ontology Development.....	145
6.1 A Brief History of Ontology Modeling	145
6.1.1 Networked Knowledge Representation and Exchange Using UML and RDF.....	145
6.1.2 Extending the Unified Modeling Language for Ontology Development	150
6.1.3 The Unified Ontology Language	155
6.1.4 UML for the Semantic Web: Transformation-Based Approach.....	156
6.1.5 The AIFB OWL DL Metamodel.....	159
6.1.6 The GOOD OLD AI ODM Proposal	160
6.2 Ontology Development Tools Based on Software Engineering Techniques	160
6.2.1 Protégé.....	161
6.2.2 DUET (DAML UML Enhanced Tool)	164
6.2.3 An Ontology Tool for IBM Rational Rose UML Models.....	165
6.2.4 Visual Ontology Modeler (VOM)	167
6.3 Summary of Relations Between UML and Ontologies	168
6.3.1 Summary of Approaches and Tools for Software Engineering-Based Ontology Development	169
6.3.2 Summary of Differences Between UML and Ontology Languages	169
6.3.3 Future Development	172
7. The MDA-Based Ontology Infrastructure.....	173
7.1 Motivation	173
7.2 Overview	174
7.3 Bridging RDF(S) and MOF.....	176
7.4 Design Rationale for the Ontology UML Profile.....	178
8. The Ontology Definition Metamodel (ODM).....	181
8.1 ODM Metamodels	181
8.2 A Few Issues Regarding the Revised Joint Submission	183
8.3 The Resource Description Framework Schema (RDFS) metamodel.....	184
8.4 The Web Ontology Language (OWL) Metamodel.....	190
9. The Ontology UML Profile	201
9.1 Classes and Individuals in Ontologies	201
9.2 Properties of Ontologies	204
9.3 Statements.....	206
9.4 Different Versions of the Ontology UML Profile.....	207

10. Mappings of MDA-Based Languages and Ontologies	211
10.1 Relations Between Modeling Spaces.....	211
10.2 Transformations Between Modeling Spaces	214
10.3 Example of an Implementation: an XSLT-Based Approach	217
10.3.1 Implementation Details.....	218
10.3.2 Transformation Example	219
10.3.3 Practical Experience	222
10.3.4 Discussion.....	225

Part III Applications

11. Using UML Tools for Ontology Modeling	229
11.1 MagicDraw	230
11.1.1 Starting with MagicDraw	230
11.1.2 Things You Should Know when Working with UML Profiles	232
11.1.3 Creating a New Ontology	234
11.1.4 Working with Ontology Classes	237
11.1.5 Working with Ontology Properties.....	240
11.1.6 Working with Individuals	244
11.1.7 Working with Statements	246
11.2 Poseidon for UML.....	247
11.2.1 Modeling Ontology Classes in Poseidon	249
11.2.2 Modeling Ontology Individuals and Statements in Poseidon	250
11.3 Sharing UML Models Between UML tools and Protégé Using the UML Back End.....	251
12. An MDA Based Ontology Platform: AIR	255
12.1 Motivation	255
12.2 The Basic Idea	256
12.3 Metamodel – the Conceptual Building Block of AIR	258
12.4 The AIR Metadata Repository.....	259
12.5 The AIR Workbench	262
12.6 The Role of XML Technologies.....	264
12.7 Possibilities.....	265
13. Examples of Ontology	267
13.1 Petri Net Ontology.....	267
13.1.1 Organization of the Petri Net Ontology	269
13.1.2 The Core Petri Net Ontology in the Ontology UML Profile	272
13.1.3 Example of an Extension: Upgraded Petri Nets.....	275