9. Ontologies, Meta-models, and the Model-Driven Paradigm

Uwe Aßmann and Steffen Zschaler

Technische Universität Dresden, Institut für Software- und Multimedi-atechnik, Lehrstuhl für Softwaretechnologie, 01062 Dresden
uwe.assmann@tu-dresden.de, steffen.zschaler@tu-dresden.de

Gerd Wagner

Brandenburgisch-Technische Universität Cottbus, Institute of Informatics, Lehrstuhl Internet-Technologie, Postfach 101344, 03013 Cottbus
g.wagner@tu-cottbus.de

In memory of Emma Larsdotter-Nilsson who, while working on a thesis in biological modelling [26], died unexpectedly in October 2005

9.1 Introduction

Refinement-based software development centres around the production of several models, going from abstract to concrete (Fig. 9.1), cumulating in the implementation as the most refined model [45]. Step by step, constructs in abstract models are *refined* to more concrete model elements. Roughly speaking, development can be divided into two phases. The *analysis phase* constructs a requirement specification describing all features the user would like to have, building on a domain model, a business model, and a context model. Later on, the *design phase* produces an architectural design specification and a detailed design specification. In a last phase, the *implementation phase*, the design specifications are filled out to an implementation of the software system.

Fig. 9.1. Models in a typical object-oriented software development process

Model-driven engineering (MDE) is a variant of this refinement-based software development in which models are no longer loosely coupled, but connected in a systematic way [9, 10]. On the one hand, MDE improves on the software refinement method of the 1970s in the sense that more concrete phases are distinguished. On the other hand, every phase derives a more concrete model not only by manual refinement, but also by semi-automatic or automatic transformation. To this end, models must be connected; that is, model elements must be traceable from a more abstract model to a more concrete model and vice versa. This is achieved through meta-modelling: *meta-models* define sets of valid models, facilitating their transformation, serialization, and exchange.

In recent years, model-driven engineering has been popularized by a specific incarnation, *model-driven architecture (MDA)*. In this process, one specific type of model information, the *platform information*, plays an important role. In MDA, models differ in how much platform information they contain (Fig. 9.2). For instance, one platform can be the programming language of the system, another can be the employed libraries or frameworks, a third can be the binary component model. The designer begins with a high-level model that abstracts from all kinds of platform issues, and iteratively transforms the model to more concrete models, introducing more and more platform-specific information. Hence, all information that

relates to programming language, frameworks, or component model are added to the platform-independent model by platform-specific extensions.
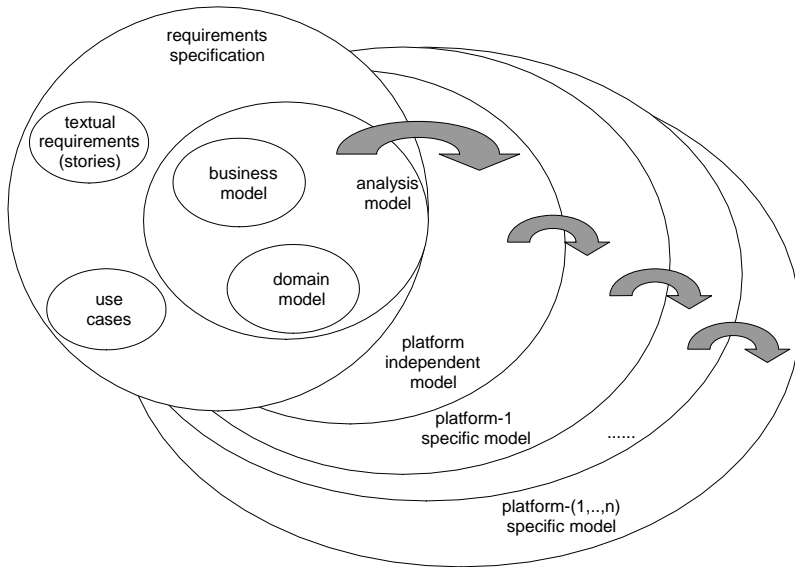

Fig. 9.2. Models in model-driven architecture (MDA)

Essentially, in MDA three types of viewpoints on models are distinguished [30]. The computationally independent (CI) viewpoint sees the system from the customer's point of view, and manifests it in a computation-independent model (CIM). This model is a typical analysis model, since it is expressed in terms of the problem domain:

> *The computation-independent viewpoint focuses on the environment of the system, and the requirements for the system; the details of the structure and processing of the system are hidden or as yet undetermined.* [30]

The CIM contains a *domain model*, describing the concepts of a domain and their interrelations, a *business model*, describing a company's rules of business, and, finally, the requirements. The platform-independent (PI) viewpoint sees the system from the designer's point of view, abstracts from all platforms a system may run on, and results in a *platform-independent model (PIM).* Roughly speaking, a PIM contains an architectural model, adorned with sufficient detail of platform-generic implemen-

tation issues. Finally, the platform-specific viewpoint adds platform-specific extensions and results in a *platform-specific model (PSM).* Either this model can be executed directly, or it is used to generate code.

To arrive at a PSM, the PIM must be extended with platform-specific information, for which it is merged with several platform-specific extensions (Fig. 9.3).
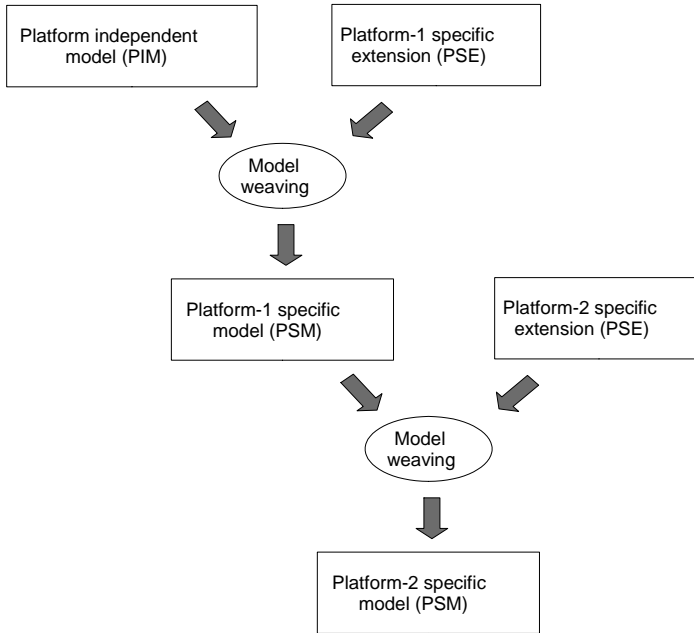

Fig. 9.3. The MDA pattern: weaving a platform-specific extension as an aspect into a PIM as a base

Because the platform-specific extension (PSE) can be regarded as an aspect that cross-cuts the platform-independent information [24], one can speak of *model weaving*. This *MDA pattern*, weaving PSMs from PIMs and PSE, can be repeated over several levels. Often, different kinds of platforms are involved and one would like to vary the system over all combinations of these platform instantiations; for example, by having a system with C# and Java, both on the web and GUI-client platforms. The idea of multi-level MDA is to repeat the model weaving process over several levels (Fig. 9.3), so that on every level, a PSM is reinterpreted as a new PIM for the next platform.

A heretical spectator could remark that MDA (and hence MDE) is not a new technology, but just refinement-based software development. How-

ever, since MDA discerns platform-specific information as the main criterion for refinement, the entire process is much more structured than the "free-style" refinement of the 1970s. Also, in MDA, all models are graph-based, while standard refinement worked mainly for syntax trees.

Recently, the Semantic Web has popularized another notion of model: *ontologies*. Ontologies are *formal explicit specifications of a shared conceptualization* [18]. They describe the concepts of a domain, similar to the domain model of a CIM. While they are currently used mainly in the Semantic Web, they could be useful also in general software development [1, 8]. But then, the question arises how ontologies should be integrated into MDE, and more specifically, into the process architecture of MDA. And this is what the rest of the chapter is about. In Sect. 9.2, ontologies are compared to general models, resulting in the insight that ontologies describe reality while models specify artifacts. Section 9.3 investigates these relationships in more detail and explains how the specification relationship instance-of can be used to build up a stack of models, the so-called IRDS meta-pyramid. Section 9.4 extends the meta-pyramid with ontologies, distinguishing a descriptive dimension. A comparison to related work concludes the chapter.

9.2 Models and Ontologies

In this section, we discuss the fundamental terms 'model' and 'ontology' and investigate their primary commonalities and differences. We begin by looking at definitions of 'model' and 'ontology', go on to discuss a fundamental property of models—namely, whether they are descriptive or prescriptive—and finish by showing how this distinction can be applied to distinguish between ontologies and other software models.

9.2.1 What's in a Model?

Models are representations, descriptions, and specifications of things. Pidd defines:

> *A model is a representation of reality intended for some definite purpose.* [34]

Hence, models represent reality (in the following denoted by the `is-represented-by` relation).

Models have a *causal connection* to the modelled part of reality: they must form *true* or *faithful* representations so that queries of the model give reliable statements about reality, or manipulations of the model result in reliable adaptations of reality. Pidd characterizes this as follows:

> *A model is an external and explicit representation of a part of reality as seen by the people who wish to use that model to understand, change, manage, and control that part of reality.* [34]

Secondly, while models represent reality faithfully, they may *abstract* from irrelevant details. For instance, while models are finite descriptions, they may well describe an infinite language—that is, an infinite set of things or systems. Usually, then, abstractions are involved—for example, about the number of elements in the language.

A model can represent many different kinds of realities, e.g. domains, languages, or, in particular, systems. Hence, we can distinguish *domain models* from *system models*, models that describe or control a set of systems:

> *A model of a system is a description or specification of that system and its environment for some certain purpose.* [31]

where the environment of a system is described by a domain model.

Models can describe structure or behaviour. In the former case, models describe the concepts of a reality and their interrelation, the *static semantics* of a domain, its *context-free* or *context-sensitive structure*. Well-formedness rules (*integrity constraints*) describe valid configurations of reality.

Example 1. UML class diagrams are frequently used together with an *Object Constraint Language* [31]. The OCL integrity constraints describe valid configurations and interrelationships of classes and objects in a UML class model.

Secondly, while a structural model contains abstractions and their interrelationships, a behavioural model also specifies their behaviour, their *dynamic semantics*. In this case, a model may state assertions on the behaviour of things in a domain or of some systems. Models can express such assertions either in a conceptual or in a transitional way. In the former case, dynamic features of a system are expressed as concepts and their interrelationships are explained by constraints. In the latter case, dynamic features and their relationships are expressed in terms of transitions on state spaces [23] or as modifications of a denotational semantics [42].

Sometimes, such transitions or modifications can in turn be expressed in logic. However, as the following example shows, this need not be appropriate. If the state space of the dynamic semantics is continuous, the semantics is better expressed by numerical means—for example, through differential equations.

Example 2. Modelica is a multi-domain modelling language for simulation, visualization, and controlling technical systems. Hence, it is a prescriptive modelling language for the dynamic semantics of technical systems [13].

9.2.2 What's in an Ontology?

Recently, the Semantic Web has popularized another notion of model— *ontologies*. One of the most-cited definitions is:

> Ontologies are *"formal explicit specifications of a shared conceptualization".* [18]

Since concepts are abstractions and play an important role in models, an ontology is certainly a special kind of model. But what is the exact difference? To answer this question, we have to introduce some other qualities of models.

Following the above definition, an ontology is a model shared by a group of people in a certain domain. This includes ontologies that have been *standardized* by international organizations (such as the Dublin Core ontology [27]), ontologies that are shared by large user groups (such as the gene ontology [3]), and ontologies that are shared between companies and their customers (such as the wine ontology [28]). In general, models need not be shared. For instance, the design model of a product, if it is shared only between the few developers of a small company, should not be regarded as an ontology, but rather as a plain artifact model. Of course, *sharedness* is a relative notion: it is often a matter of taste to consider a user group of a model large enough so that the model can be called an ontology of the user group.

An important property of ontologies is the so-called *open-world assumption* [20]. It states, intuitively, that anything not explicitly expressed by an ontology is unknown. Hence, ontologies use a form of partial description or under-specification as an important means of abstraction. In contrast, most system models underlie the assumption that what has not been specified is either implicitly disallowed or implicitly allowed

(*closed-world assumption*), to restrict arbitrary extensions of the system, which could introduce inconsistencies.

It is important to distinguish whether models describe or control reality. If they describe, they monitor reality and form *true*, or *faithful*, abstractions. If they control, they prescribe reality; that is, they specify well-formedness conditions what reality should be like, once it has been constructed. It can also be said that such models are templates or schemas of reality. Hence, a most fundamental feature of a model is that it can be *descriptive* or *prescriptive* [38]. In the former case, the model describes reality, but reality is not constructed from it. In the latter case, the model prescribes the structure or behaviour of reality and reality is constructed according to the model; that is, the model is a *specification* of reality. Favre [11] observes that in a descriptive model truth lies in reality, whereas in a prescriptive model, truth lies in the model itself. Descriptive models are, of course, used in analysis and re-engineering, specifications in design and forward engineering. Since most specifications model systems, a prescriptive system model is also called a *system specification*.

Models are abstractions from reality for some purpose [34]. Ontologies are special models. Most of the models used in software development and design are of a prescriptive nature in that they form the templates from which the system is later implemented. In contrast, because of their open-world assumption, ontologies should be regarded as descriptive models. This is so, because the open-world assumption does not allow for a complete and final description: Anything that has not been said explicitly is unknown. Two very different systems may satisfy an ontology, if they differ in areas not explicitly mentioned in the ontology.

On the other hand, we concede that ontologies can also be—and often are—used in a prescriptive manner. We argue, however, that then they should better not be called ontologies, but specification models. When a model is used as a prescription for systems, it should confine their legal structure, for which the closed-world assumption is required. At least, at a certain point in development, the world must be closed; that is, the additional assumption has to be introduced that everything that has not yet been specified or cannot be derived is wrong. Such a *world closure* is not only hard to comprehend because it changes the semantics of the underlying logic, but may also require the insertion of additional facts in the database or a change in the logic reasoner.

Taking this discussion into account, we define the following:

> An ontology is a *shared, descriptive, structural model, representing reality by a set of concepts, their interrelations, and constraints under the open-world assumption.*

> A specification model is a *prescriptive model, representing a set of artifacts by a set of concepts, their interrelations, and constraints under the closed-world assumption.*

These definitions deserve some elucidating remarks. When comparing hallmark papers, such as [18] and [38], specification models and ontologies look very similar. Both provide vocabulary for a language and define validity rules for the elements of the language. Both specification models and ontologies use integrity constraints to limit the valid instances of the domain.[32]

However, there are also differences. Ontologies are *shared* knowledge; that is, they must be standardized in a certain group of people. Ontologies are not specification models, but descriptive models in Seidewitz's sense. Ontologies do not describe systems, only domains. Hence, in a software engineering process, they should play the role of an analysis model, not of a design or implementation model. With this view we contradict Devedzic, "Generally, an ontology is a meta-model describing how to build models" [8], and Gruber, because he maintains that ontologies are specifications [18]. However, this conceptual distinction creates a natural place for ontologies in model-driven engineering, as will be seen in Sect. 9.4.

To summarize, we will assume the following: Specification models focus on the *specification*, *control*, and *generation* of *systems*, ontologies on *description* and *conceptualization (structural modelling)* of *things*. Both kinds of models have in common the qualities of *abstraction* and *causal connection.* So, under these circumstances, how can ontologies and specification models cohabit in model-driven engineering?

9.3 Similarity Relations and Meta-modelling

The previous arguments make it possible to distinguish two basic notions of the is-represented-by relation between a model and the corresponding part of reality (Fig. 9.4). In a descriptive model—for example, an ontology—the model describes the world; that is, the world's objects are in relation is-described-by with concepts of the descriptive model. In a specification model, the system's objects are created from the model; that is, an object is an instance-of a model element. Both relationships are representation relations: one is descriptive, the other is prescriptive. Their generalization is-represented-by is a similarity relation, in which a causal connec-

---

[32] Both are structural models in the sense that while they can contain concepts that model behaviour, they usually do not model dynamic semantics.

tion—delivering true and faithful statements—is defined between the represented things and the representing model. Beyond that, more similarity relations can be defined; for example, two things may share features (often expressed as is-a, i.e., structural or behavioural inheritance), or they may be included in a hierarchy of sets (set inclusion, subset-of). In Fig. 9.4, is-a is defined as a sub-relationship of subset-of, because inheritance usually has a set-based semantics; namely, that all objects in a subclass are also members of the superclass. Additionally, is-a is a sub-relationship of is-described-by, because a superclass also describes all objects in a subclass. In contrast, is-a cannot be a sub-relationship of instance-of, because a superclass cannot necessarily be regarded as a template, schema, or specification for a subclass.
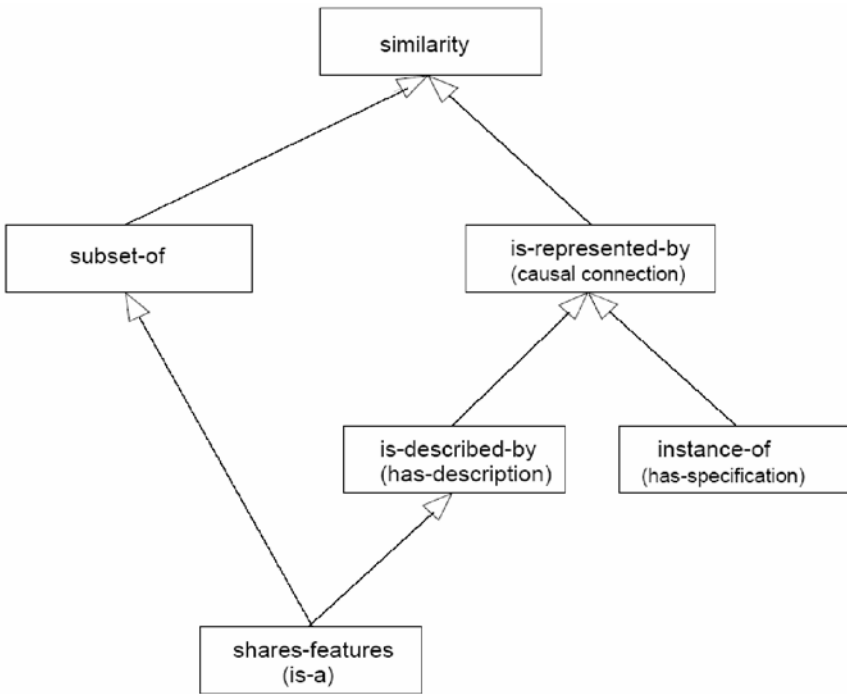

Fig. 9.4. A classification of similarity relations

9.3.1 Meta-models

In MDE, the specification relationship instance-of plays a special role. When the specification principle is applied repeatedly, models are re-

garded as the reality or system under study, so that models specifying models can be defined: namely, meta-models. Meta-models represent and specify models; that is, they describe about what are the valid ingredients of a model. More precisely:

> *A meta-model makes statements about what can be expressed*
> *in the valid models of a certain modelling language.* [38]

Hence, a meta-model is a prescriptive model of a modelling language [38]. In general, meta-models are language specifications, not only of modelling, but also of arbitrary languages. In the current stage of MDE, they are mainly concerned with the static semantics—that is, with context-sensitive syntax of models, integrity, and well-formedness constraints. However, modelling languages for dynamic semantics could also be applied to construct meta-models [42].

A language concept or construct in a meta-model is captured by a *meta-class.* While its structure and embedding describe the static semantics of the language constructs, its methods describe the dynamic behaviour of the language construct. Usually, meta-classes are assembled in a behavioural meta-model, the *meta-object protocol (MOP)* [25], a reflective meta-model that describes an interpreter for the language.

A big incentive for meta-modelling has been the need of CASE (Computer-Aided Software Engineering) tool vendors to exchange models [32]. Since a meta-model describes, rather specifies, valid instances of a modelling language—models—it enables control over the structure and validity of models. If two CASE tools agree on the same meta-model, they impose the same structure on their models, so that they can easily exchange them.

A language, described by a meta-model, can have a specific purpose or domain in which it is applied. Such purposes or modelling domains are called the *subject areas* of meta-models [12].

Example 3.   For instance, the common warehouse meta-model (CWM) [29] defines a data specification language, a meta-model for data and information system applications. Work-flow systems are another special subject area whose data, functions, and tasks can also be described with meta-models [36]. Software processes, being specific work flows, can be meta-modelled [14] and used to construct software environments [5].

Subject areas can be organized in hierarchies or partial orders. Then, meta-models in a certain subject area can build on others from lower-level subject areas, so that complex languages can reuse simpler languages [12].

Example 4. The CASE Data Interchange Format (CDIF) has structured its meta-model into several subject areas (Fig. 9.5). The Foundation

module contains information about names and relations; the `Common` module defines name aliasing for objects; and the `Data` module describes access paths to data and roles of objects. Based on these, data flow can be defined (`Data Flow` module). Another module specifies facilities for the presentation of objects. Finally, the full integrated meta-model uses all other modules and provides their concepts in an integrated way to the users.
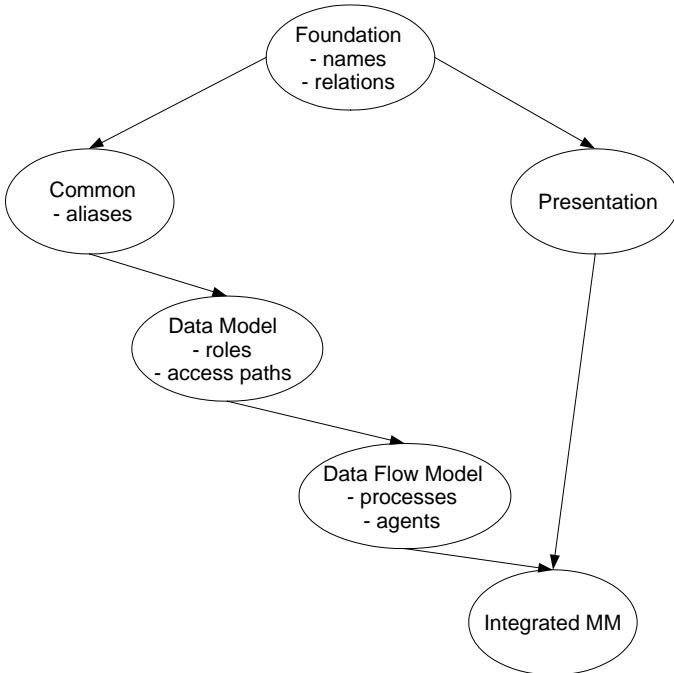

Fig. 9.5. The subject areas of CDIF and their meta-models in a use relationship

9.3.2 Metameta-models

The specification principle can be applied repeatedly. Metameta-models *represent* and *specify* meta-models; that is, they describe what are the valid ingredients of a meta-model. They specify languages, and are thus a form of *language specification languages (meta-languages).*

In order to model anything useful, such a minimal meta-language should contain the following concepts [12]:

- classes (concepts);
- attributes (or properties) of classes, contained in the classes;
- binary relations between classes.

Thus, the Entity–Relationship Diagram (ERD) language [6] can be used as a very simple meta-language. It defines modelling concepts, their attributes, and their relationships. Other meta-languages exist that describe other forms of languages, or describe specific aspects:

1. Grammar specification languages—for example, EBNF—specify the concrete or abstract syntax of a text-based language [17].
2. Attribute grammars describe context-sensitive syntax in the form of attribution rules of syntax trees [7].
3. Natural semantics can be employed for type systems, but are also able to specify dynamic semantics of systems [23].
4. In SGML [16], mark up languages can be defined. XML [44] is a variant of SGML, allowing for defining context-free mark up languages.
5. EXPRESS [37], a modelling language in the spirit of UML, is frequently used in mechanical engineering.

9.3.3 The Meta-pyramid, the Modelling Architecture of MDE

Based on the meta-principle, a so-called meta-pyramid can be defined, which displays systematically the mentioned stack of models and meta-models [22]. In essence, a meta-pyramid is a specification hierarchy linked by the instance-of relation, in which upper-level meta-models in some way specify other sets of lower-level models. Since sets of models can be regarded as languages, the meta-pyramid is a hierarchy of language specifications.

In this chapter, we focus on the standard meta-pyramid of OMG, originally presented in the ISO Information Resource Dictionary System (IRDS) standard [22] (Fig. 9.6), which contains four levels: M0 level (objects), M1 level (models), M2 level (meta-model or language level), M3 level (metameta-model or language description level). There are alternatives and a debate is going on whether the IRDS meta-pyramid is precise enough, because it is one-dimensional, while multi-dimensional model pyramids exist [2]. However, at the moment, this is the mainstream meta-pyramid of MDE.

On level M3, the IRDS/OMG meta-pyramid employs the *meta-object facility (MOF)* as metameta-model. Essentially, its concepts are similar to those of the ERD. The stereotypical models of MDA, CIM, PIM, and PSM live on level M1. All of them are specified on level M2 by meta-models (CIM-MM, PIM-MM, PSM-MM), dialects of UML, enriching the UML core by *profiles* containing markup for model elements (stereotypes and tagged values).
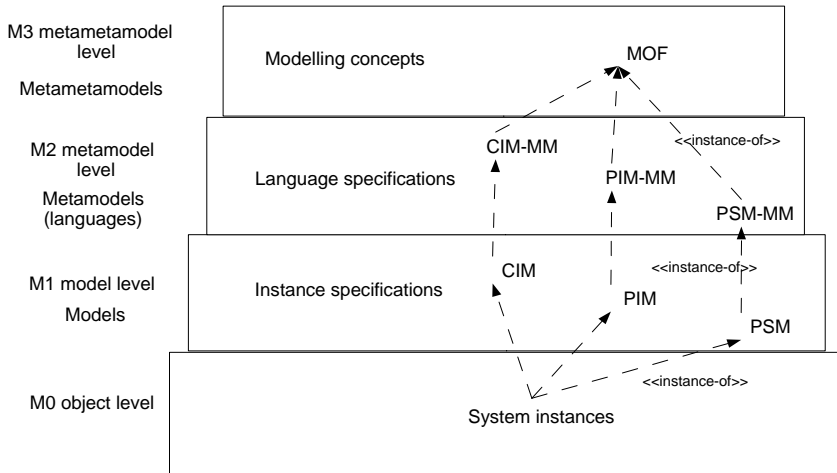

Fig. 9.6. The meta-pyramid with the MDA-related model types CIM, PIM, PSM

Each of these meta-models covers different subject areas of a PSM: the CIM-MM covers the requirements, the PIM-MM covers the platform-independent concepts, while the PSM-MM adds the platform issues. While all of these models are prescriptive—that is, using the `instance-of` relationship—the question remains how ontologies, being models relying on `described-by`, can be integrated into the meta-pyramid. This is the topic of the next section.

9.4 MDE and Ontologies

This section discusses the role of descriptive and structural models, in particular ontologies, in the model-driven process. First, the different role of domain and upper-level ontologies is discussed. We postulate that upper-level ontologies can also be used as language descriptions. Secondly, we propose an embedding of parts of the CIM as ontologies into the MDA meta-pyramid (ontology-aware meta-pyramid). In fact, this delivers a first

ontology-aware mega-model of MDE [10], and we discuss its conceptual advantages. On the one hand, the mega-model suggests an extended, ontology-aware software process. On the other hand, the technologies for tool construction in the MDA and MOF world can be transferred to the ontology world.

9.4.1 Domain and Upper-Level Ontologies

The basic idea of the ontology-aware meta-pyramid is that most models in MDE are specifications, but can integrate ontologies on different meta-levels as descriptive analysis models. Since ontologies differ from specifications due to their descriptive nature, the standard M0–M3 meta-pyramid can be refined from using pure specification models to also using ontologies.

Depending on the meta-level, an ontology may serve different purposes. In fact, there are different qualities of ontologies in the literature. First of all, the word *ontology* stems from philosophy, where it characterizes *Existence*:

> *Ontology is a systematic account of Existence.* [18]

We call such a systematic account of existence a *World ontology*, a conceptualization of the world, that is all existing concepts. Usually, a World ontology is split into an *upper-level ontology (concept ontology, frame ontology)*, providing basic concepts for classification and description, and several lower-level ontologies, *domain ontologies* describing domains of the world [19, 41]. Sowa characterizes domain ontologies as follows:

> *The subject of ontology is the study of the categories of things that exist or may exist in some domain. The product of such a study, called an ontology, is a catalogue of the types of things that are assumed to exist in a domain of interest D from the perspective of a person who uses a language L for the purpose of talking about D. The types in the ontology represent the predicates, word senses, or concept and relation types of the language L when used to discuss topics in the domain D.* [40]

In contrast, upper-level ontologies can be defined as follows:

> *An upper ontology is limited to concepts that are meta, ge-*
> *neric, abstract and philosophical, and therefore are general*
> *enough to address (at a high level) a broad range of domain*
> *areas. Concepts specific to given domains will not be in-*
> *cluded; however, this standard will provide a structure and a*
> *set of general concepts upon which domain ontologies (e.g.,*
> *medical, financial, engineering, etc.) could be constructed.*
> [21]

Usually, concepts of the domain ontology *inherit* from concepts in the upper-level ontology. For better interoperability and understanding, some researchers try to create a normalized upper-level ontology, from which all possible domain ontologies may inherit [33]. If a standardized upper-level ontology with modelling concepts existed, all domain ontologies could rely on a standardized concept vocabulary.

9.4.2 Relationship of Ontologies and System Models on Different Meta-levels

With this terminological distinction, we can relate the different forms of ontologies to meta-levels in the meta-pyramid. Domain ontologies live on level M1; they correspond to models. An upper-level ontology, also a standardized one, should live on level M2, because it provides a language for ontologies. Figure 9.7 summarizes this insight, showing both dimensions and descriptive and prescriptive models, on different meta-layers.

Interestingly, on the ontology side, inheritance is used as the connecting relation of M1 and M2, and not `instance-of`. We believe that this historic choice, which might have been made unconsciously, has a deep semantic reason in the difference between descriptiveness and prescriptiveness. A concept in a domain ontology on M1 needs to express its `similarity` to a modelling concept of an upper-level ontology (on M2). For this, the `is-a` relationship is sufficiently precise (cf. Fig. 9.4), and therefore it has been selected in the ontology world to connect the meta-levels. A concept in a specification model, however, has to express *that it has been made from* a specification model, which is clearly a more specific relationship than `is-a`. And this is the reason why in the IRDS world the `instance-of` relationship has been employed.

We argue that on level M3 of the descriptive side of the ontology-aware meta-pyramid, also a specification meta-language should be employed (Fig. 9.7). The language that describes or specifies an ontology language cannot be descriptive, because ontology languages are not something

given, but artificial languages. Hence, a model to represent them should be prescriptive. We argue that the same meta-language can be used on the ontology as well as on the system model side.
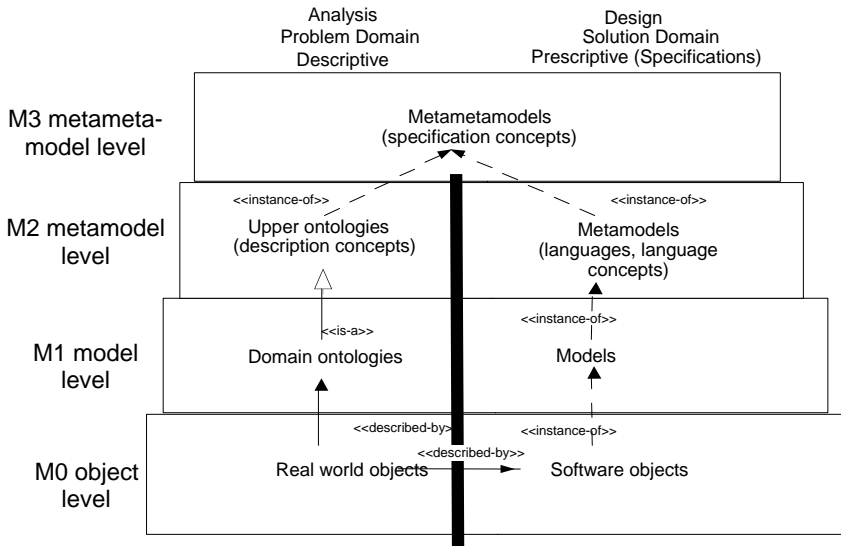

Fig. 9.7. The ontology-aware meta-pyramid

In fact, inheritance is not required in Fig. 9.7. While, usually, concepts in a domain ontology *inherit* from a concept in an upper-level frame ontology, we suggest that to distinguish them better from concepts in specification models, ontology modelling should causally connect ontological concepts by the `described-by` relationship. This would introduce a parallelism to using `instance-of` on the specification side and retain the basic ontological modelling principle of descriptiveness. Because of the parallel structure to the specification dimension, the advantage of such a meta-pyramid is that connections from ontologies to specifications can easily be made. In particular, this holds for the application of the meta-pyramid in the MDE.

9.4.3 Employing Domain Ontologies in the MDA

This version of an ontology-aware meta-pyramid permits us to group the MDA-based models around ontologies. In particular, the CIM plays a special role.

A CIM contains information about the system from the perspective of the system user. It is an *analysis model*. As such, it may contain a domain model, a business model, and requirements (Fig. 9.1) [30]. The gap between descriptive and prescriptive models concerns the CIM in particular. The domain model of a CIM can be selected to be a domain ontology (CIM-DO in Fig. 9.8). A business model, capturing business rules for a company that should prevail in all software products, can also be regarded as a domain ontology, namely that of the rules of the company (i.e. a domain ontology for a company, CIM-BO in Fig. 9.8). However, the parts of the CIM that deal with requirements cannot be grasped by ontologies, because they *specify* requirements of a system to-be-made. Hence, this specification is grouped in CIM-RM in Fig. 9.8 as a specification model. This difference is also the reason why only for CIM-RM, the specification part of the CIM, is a meta-model needed. Concepts of CIM-DO or CIM-BO *describe* existing things, and may inherit from concepts on the language or concept ontological level. Concepts in CIM-RM, on the other hand, are instances of a CIM meta-model, because they specify parts of functions of a system.
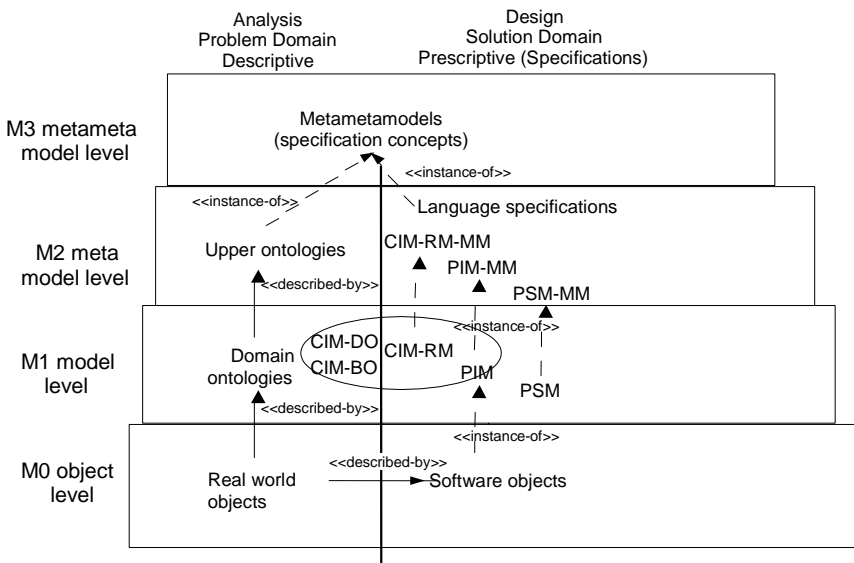

Fig. 9.8. A proposal for the role of ontologies in meta-pyramid of MDE and the MDA

Usually, a CIM is extended towards a PIM by hand, by enriching it with operational model elements. Hence, at least CIM-DO and CIM-BO play

the role of standardized analysis models, whose elements can be traced back from the PIM [1]:

> *In an MDA specification of a system, CIM requirements should be traceable to the PIM and PSM constructs that implement them, and vice versa.* [30]

Hence, surprisingly, the MDA can benefit from ontologies, because via the standardized domain and business ontologies, once parts of a CIM, connection to PIM specifications can be made in a clear and systematic way.

9.4.4 Conceptual Benefits of an Ontology-Aware Meta-pyramid

The ontology-aware meta-pyramid offers several other benefits. First of all, it suggests a more concrete model-driven software development process. The designer starts from standardized analysis models, ontologies, which may have been defined long before project start. These domain and business models are refined towards design models. First, the requirements are added to yield a complete CIM. This is refined to a PIM and, then, conventionally, via several PSMs towards an implementation. Employing ontologies as analysis models should increase the reliability of software products, because these models are well engineered, often used, and hence trustworthy. This avoids the risks of a self-made domain analysis.

Secondly, ontologies as analysis models offer a more common vocabulary for the software architect, customer, and domain expert. This should improve the understanding of the parties that order and construct software. Then, the standardization of the ontologies improves the interoperability of applications, because applications that use the ontology contain a common core of common vocabulary. Finally, domain and business ontologies can be reused in many software products. In particular, they may form the core of a software product line [1], around which many products are grouped, and from which they reuse domain terminology. Overall, this improves reuse in the software process.

It is also beneficial to make an explicit distinction between descriptive and prescriptive models in the MDA. Modelling becomes easier, because designers and domain experts can always answer the question: where does the truth lie? In the model or in reality? Specification models have to confine themselves to the modelling of *artificial things*, things that are made, while ontologies can focus on the description of *real things*, things that ex-

ist. (In particular, this can be seen from the example of the CIM, which in fact contains descriptive and prescriptive models.)

Finally, the ontology-aware meta-pyramid distinguishes conceptual from behavioural models. It seems to be convenient to centre software modelling around the concepts of a domain, or structure of a domain, while adding behaviour to it step by step. In essence, this supports one of the central ideas of MDA: refinement.

9.4.5 Tools Based on an Ontology-Aware Meta-pyramid

Ontology-aware meta-pyramids not only deliver a conceptual integration of the Semantic Web and MDE, but also enable us to compare engineering practices of both paradigms to derive common tools.

In MDE, type systems are mediated by an *interface definition language (IDL)* [39]. Based on the meta-models for two type systems (on level M2), automatic conversion code (on level M1) between objects typed in type system 1 and objects typed in type system 2 can be generated. This is the task of an IDL compiler and facilitates *interoperability* between components and services, because data can easily be serialized and de-serialized in appropriate forms. At the moment, interoperability between ontology-based applications is an unsolved problem, but it might be possible to transfer the IDL tools to ontology languages.

The division of M1 models into platform-aware subject areas (CIM, PIM, PSM) is a structuring principle that can be applied to the ontology world. Because the principle has been invented for the reuse of models in product families (CIM and PIM are reused in many PIMs and PSMs, respectively), it could enable reuse of abstract ontologies in *ontology families*. Domains are not always disjoint, but often overlap. This suggests that *abstract ontologies* should be developed that can be shared between domains and are refined towards *concrete ontologies* by adding the differences of domains. Whether the notion of *platform* is the right criterion for abstraction remains to be seen; however, MDE tools, such as MDE code generators, could easily be transferred to such ontology families.

The success of ontologies and ontology languages suggests the use of logic in specification models. This is often the reason why, in practice, ontologies are abused in a prescriptive way. However, it would be more beneficial to reflect the role of open- and closed-world assumptions in ontology and specification languages. For a given modelling language, when is it possible to change the assumption? And how far can tools be reused if the assumption is orthogonal to the modelling language?

In the MDE world, the exchange of meta-data has been simplified by the XMI standard [32]. Essentially, XMI defines meta-model mappings on level M2 between the UML meta-model, XML schema definitions, and a programming language—for example, Java. Based on these mappings, serialization of graph-like UML models to tree-shaped XML models can be automated. Also, Java class models, which use a restricted form of inheritance, can be generated automatically. XMI lays the foundation for meta-data repositories such as MDR [43] or Eclipse-MDR [15], which seem to be the basis for future CASE tools and integrated software–development environments. Based on the ontology-aware meta-pyramid, the XMI technology could be transferred to ontology repositories.

Figure 9.8 suggests a common meta-language for the ontology and specification world. It should be clear by now that such a meta-language should be based on an expressive logic. If this logic is decidable (as in the case of OWL-DL), decidable tool technology can be built. If the logic is undecidable, it is more expressive, which might be more useful. Perhaps it is possible to define a hierarchy of compatible logic languages that combines expressive power with flexibility of use. Such a language hierarchy would certainly be of great help to build tools in both the descriptive ontology as well as the prescriptive specification world.

9.4.6 The mega-Model of Ontology-Aware MDE

The above-presented ontology-aware meta-pyramid can be called a mega-model of ontology-aware MDE:

> *A mega-model is a model that describes a meta-pyramid.* [11]

A mega-model stands outside of the meta-pyramid and describes all its levels. It has a global influence on all levels of the meta-pyramid. As such, the presented mega-model sheds new light on the relation of ontologies and meta-models in MDE. Systematically, ontologies can be related to specification models and meta-models in the meta-pyramid. It is important to distinguish the representation relations `is-described-by` and `instance-of`, because then ontologies can be differentiated from specification models on all levels. As a whole, we propose that:

1. An ontology-aware MDA should employ domain and business ontologies as parts of the CIM.
2. An ontology-aware MDE should additionally incorporate a second dimension of ontologies as descriptive models in the meta-

pyramid, and maintain interrelations between the descriptive and prescriptive models on all levels.

9.5 Related Work

One of the works integrating meta-models and ontologies is [35], which extends software process and measurement ontologies to meta-models from which software can be built. The work demonstrates the usefulness of ontologies in a meta-modelling scenario.

The standard aforementioned meta-pyramid is not undebated in the literature. Other pyramids can be described, in particular if some design principles for meta-pyramids found in the literature are varied. A central role is played by the similarity relations: since different notions can be defined, different model hierarchies result.

Favre dissects the `instanceOf` relation into `representationOf` and `member-of` [9]. A model *represents* a language, and a system is an element of that language. This leads to a *relative* model hierarchy which is not restricted to four levels, but in which certain composite patterns denote more complex similarity relations, such as `instance-of` or `described-by`.

If every element on level $n+1$ is an instance of exactly *one* element on level n, a meta-pyramid is called *strict* [2]. With strict similarity, meta-pyramids must be lists or trees and are essentially one-dimensional. Based on this distinction, [2] defines a non-strict meta-pyramid consisting of two dimensions arranged in a matrix. One dimension of the matrix is characterized by physical (technical, *linguistic*) instantiation. The linguistic similarity describes the *specification language aspect* of modelling: which language construct is an instance of which language concept. Linguistic similarity is distinguished from logical *(ontological)* similarity, which spans the other dimension, the matrix-like meta-pyramid. Ontological similarity describes the similarity of real-world concepts, e.g. that a dog is a mammal, and Fido is a dog. Clearly, this dimension corresponds to our descriptive, ontological dimension. However, [2] does not distinguish prescriptive vs. descriptive models, nor further different forms of similarity relations. Future work will combine both approaches; at this time, it seems unclear whether a two-dimensional matrix-like approach or the presented approach of parallel descriptive and prescriptive dimensions will prevail.

9.6 Conclusions

Ontologies are no silver bullet. They can be employed in the software process as descriptive standardized domain models, domain-specific languages, and modelling (description) languages. However, they should not be mingled with specifications of software systems. In MDE, both forms of models are needed and complement each other. It is time to develop appropriate mega-models that clarify the role of ontologies in MDE. This chapter has presented one approach; however, this can be only an intermediate step, because we restricted ourselves to the standard IRDS metapyramid. Other, more sophisticated meta-pyramids exist and must be extended to be ontology-aware.

9.7 Acknowledgments

References

1. Aßmann, U., Reuse in semantic applications. In Norbert Eisinger and Jan Małuszynski, editors, Reasoning Web, First International Summer School 2005. Lecture Notes in Computer Science 3564, Springer, Berlin, July 2005.
2. Atkinson, C., and Kühne, T., Model-driven development: A metamodeling foundation. IEEE Software, 20(5):36–41, 2003.
3. Smith B., Williams, J., and Schulze-Kremer, S., The ontology of the Gene Ontology. In AMIA 2003 – Annual Symposium of the American Medical Informatics Association, 2003.
   http://www.gene-ontology.org
4. Bry, F. et al. Rules in a Semantic Web Environment (REWERSE). EU Project 6th framework. IST-2004-506779.
   http://www.rewerse.net
5. Canfora, G., García, F., Piattini, M., Ruiz, F., and Visaggio, C.A., Applying a framework for the improvement of software process maturity. Software - Practice and Experience, 36(3): 283–304, March 2005, Wiley, New York.
6. Chen, P.P.-S., The entity-relationship model - towards a unified view of data. Transactions on Database Systems, 1(1):9–36, 1976.
7. Deransart, P., Jourdan, M., and Lorho, B., Attribute grammars - definitions, systems and bibliography. Lecture Notes in Computer Science 323, Springer, Berlin, 1988.

8. Devedzic, V., Understanding ontological engineering. Communications of the ACM, 45(4):136–144, 2002.

9. Favre, J.-M., Foundations of model (driven) (reverse) engineering: Models. Technical Report, vol. 1–3ADELE Team, Laboratoire LSR-IMAG, Université Joseph Fourier, Grenoble, France, 2004.

10. Favre, J.-M., Megamodeling and etymology - a story of words: From MED to MDE via MODEL in five milleniums. In Dagstuhl Seminar on Transformation Techniques in Software Engineering, no. 05161 in DROPS 04101. IFBI, 2005.

11. Favre, J.-M., and Nguyen, T., Towards a megamodel to model software evolution through transformations. Electronic Notes in Theoretical. Computer Science 127(3):59–74, 2005.

12. Flatscher, R., Metamodeling in EIA/CDIF - meta-metamodel and metamodels. ACM Transactions on Modeling and Computer Simulation, 12(4):322–342, 2002.

13. Fritzson, P, and Engelson, V., Modelica—A unified object-oriented language for system modeling and simulation. In Eric Jul, editor, ECOOP '98 – Object-Oriented Programming, Lecture Notes in Computer Science 1445 pages 67–90. Springer, Berlin, 1998.

14. García, F., Ruiz, F., Piattini, M., and Polo, M., Conceptual architecture for the assessment and improvement of software maintenance. In Mario Piattini and Joaquim Filipe, editors, Enterprise Information Systems IV (ICEIS), pages 219–226. Kluwer Academic Publishers, Dordrecht, 2002.

15. Geer, D., Eclipse becomes the dominant Java IDE. IEEE Computer, 38(7):16–18, 2005.

16. Goldfarb, S.F., The SGML Handbook. OUP, Oxford, 1990.

17. Goos, G., and Waite, W.M., Compiler Construction. Springer, Berlin, 1984.

18. Gruber, T.R., A translation approach to portable ontology specifications. Knowledge Acquisition, 5(2):199–220, 1993.

19. Guizzardi, G., Herre, H., and Wagner, G., On the general ontological foundations of conceptual modeling. In S. Spaccapietra, S.T. March, and Y. Kambayashi, editors, 21st International Conference on Conceptual Modeling (ER 2002), Lecture Notes in Computer Science 2503, pages 65–78, Springer, Berlin, 2002.

20. Horrocks, I., Patel-Schneider, P., and van Harmelen, F., From SHIQ and RDF to OWL: The making of a web ontology language. Journal of Web Semantics, 1(1):7–26, 2003.

21. IEEE. Standard upper ontology knowledge interchange format. Technical Report, 2003.
http://suo.ieee.org/suo-kif.html

22. ISO and IEC. Information technology – information resource dictionary system (IRDS). International Standard ISO/IEC 10027, 1990.

23. Kahn, G., Natural semantics. Report no. 601, INRIA, February 1987.

24. Kiczales, G., Aspect-oriented programming. ACM Computing Surveys, 28(4), December 1996.

25. Kiczales, G., des Rivières, J., and Bobrow, D.G., The Art of the Metaobject Protocol. MIT Press, Cambridge, MA, 1991.
26. Larsdotter-Nilsson, E., and Fritzson, P.,. Using Modelica for modeling of discrete, continuous and hybrid biological and biochemical systems. In the 3rd Conference on Modeling and Simulation in Biology, Medicine and Biomedical Engineering. The University of Balamand, May 2003.
27. Needleman, M.H., Dublin core metadata element set. Serials Review, 24(3–4):131–135, Elsevier, 1998.
28. Fridman, N., and Musen, M.A., Ontology versioning in an ontology management framework. IEEE Intelligent Systems, 19(4):6–13, 2004.
29. Object Management Group (OMG). Common warehouse metamodel (CWM), February10, 2000.
30. OMG. MDA Guide, June 2003.
    http://www.omg.org/mda
31. OMG. UML 2.0 Object Constraint Language (OCL) specification, 2003.
    http://www.omg.org/docs/ptc/03-10-14.pdf
32. OMG. XML Metadata Interchange (XMI), January 2002.
    http://www.omg.org/technology/documents/format/xmi.htm
33. Pease, A., Niles, I., and Teknowledge Corporation. Towards a standard upper ontology. In FOIS, Ogunquit, Maine, ACM, October 2001.
34. Pidd., M., Tools for Thinking - Modeling in Management Science. Wiley, New York, 2000.
35. Ruiz, F., Vizcaíno, A., Piattini, M., and García, F., An ontology for the management of software maintenance projects. International Journal of Software Engineering and Knowledge Engineering, 14(3):323–349, 2004.
36. Scheer, A.-W., ARIS - Business Process Frameworks. Springer, Berlin, 1998.
37. Schenck, D., The express language reference manual. Technical Report ISO TC184/SC4/WG1 N466 Working Document, ISO, March 1990.
38. Seidewitz, E., What models mean. IEEE Software, 20:26–32, September 2003.
39. Siegel, J., OMG overview: CORBA and the OMA in enterprise computing. Communications of the ACM, 41(10):37–43, October 1998.
40. Sowa, J.F., Ontologies Website.
    http://www.jfsowa.com/ontology/index.htm
41. Sowa, J.F., Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks Cole Publishing, Belmont, 2000.
42. Stoy, J.E.,, Denotational Sematics: The Scott-Strachey Approach to Programming Language Theory. MIT Press, Cambridge, MA, 1977.
43. Surveyer, J., Sun adds to opensource Java IDE roster: A review of NetBeans Java IDE. Application Development Trends, 11(9):48–48, 2004.
44. W3C. Extensible markup language (XML) 1.0. Technical Report REC-xml-19980210, February 1998.
45. Wirth, N., Program development by stepwise refinement. Communications of the ACM, 14(4): 221–227, 1971.