

Timo Asikainen and Tomi Männistö. 2009. Nivel: a metamodelling language with a formal semantics. *Software and Systems Modeling*, volume 8, number 4, pages 521-549.

© 2008 by authors and © 2008 Springer Science+Business Media

Preprinted with permission from Springer Science and Business Media.

The final publication is available at <http://www.springerlink.com/>.

<http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s10270-008-0103-2>

Nivel: A metamodelling language with a formal semantics

Timo Asikainen, Tomi Männistö *

Helsinki University of Technology, Department of Computer Science and Engineering, P.O. Box 9210, FI-02015 TKK, Finland

Abstract Much work has been done to clarify the notion of metamodelling and new ideas, such as strict metamodelling, distinction between ontological and linguistic instantiation, unified modelling elements and deep instantiation, have been introduced. However, many of these ideas have not yet been fully developed and integrated into modelling languages with (concrete) syntax, rigorous semantics and tool support. Consequently, applying these ideas in practice and reasoning about their meaning is difficult, if not impossible. In this paper, we strive to add semantic rigour and conceptual clarity to metamodelling through the introduction of NIVEL, a novel metamodelling language capable of expressing models spanning an arbitrary number of levels. NIVEL is based on a core set of conceptual modelling concepts: class, generalisation, instantiation, attribute, value and association. NIVEL adheres to a form of strict metamodelling and supports deep instantiation of classes, associations and attributes. A formal semantics is given for NIVEL by translation to weight constraint rule language (WCRL), which enables decidable, automated reasoning about NIVEL. The modelling facilities of NIVEL and the utility of the formalisation are demonstrated in a case study on feature modelling.

Key words conceptual modelling – metamodelling – NIVEL – weight constraint rules – formal semantics

Send offprint requests to: Timo Asikainen

* We gratefully acknowledge the financial support from the Finnish Funding Agency for Technology and Innovation (Tekes) and from the Technology Industries of Finland Centennial Foundation.

Correspondence to: Timo Asikainen, tel. +358 9 4515364, fax +358 9 4514958, e-mail: timo.asikainen@tkk.fi

1 Introduction

Modelling is one of the fundamental paradigms underlying any mature field of engineering. This is also the case in software engineering, where initiatives such as the model-driven engineering one highlight the importance of models as engineering artefacts [24]. Metamodelling is an active area of research within the software modelling community with a long history, see, e.g., [15, 18].

Recently, much work has been done to clarify the notion of metamodelling [22] and interesting ideas, such as strict metamodelling [4, 5], distinction between ontological and linguistic instantiation [6, 22], unified modelling elements [5] and deep instantiation [3, 7] have been introduced and argued for; see Section 2 for an overview. Although these ideas are intuitively clear and described by some formal rules [3, 4, 23], they have not been developed into modelling languages with either formal semantics or supporting tools, which essentially makes them practically inapplicable.

A number of reasons speak in favour of giving a modelling language a formal semantics. Without a semantics, a language only amounts to a collection of notations. Such a collection may well be useful for communication purposes, but may lead into disputes over the proper usage and interpretation of the notations [14]. Although the semantics of a modelling language would be intuitively clear, it may still not be precise enough to enable formal (automated) reasoning and implementing model transformations required, e.g., in a model-driven development approach. The importance of a formal semantics is also emphasised by the large number of papers formalising parts of UML, e.g., [9, 14, 25]. Finally, it is particularly useful to give a formal semantics to a metamodelling language: a metamodel, when viewed as a modelling language, expressed in the metamodelling language is given at least a partial semantics, as demonstrated in Section 5.

On the other hand, it is not the case that a language with a formal semantics would in general be unaccept-

able to users [16]. Instead, such a language can be made accessible through a suitable concrete syntax or supporting tools; this is an approach sometimes referred to as “logic through the backdoor” in artificial intelligence research [33].

In this paper, we introduce NIVEL, a novel meta-modelling language capable of expressing models spanning an arbitrary number of levels. NIVEL is intended to cover a large variety of different modelling purposes and is therefore not committed to any single modelling paradigm, such as object-orientation. NIVEL is based on a core set of modelling concepts—class, generalisation, attribute, value and association—found in many existing conceptual modelling languages. In addition, emphasis is put on defining the instantiation relationship between classes, which enables NIVEL to be credited as a meta-modelling language. Finally, NIVEL incorporates a number of more recent ideas including strict meta-modelling [4, 5], distinction between ontological and linguistic instantiation [6, 22], unified modelling elements [5] and deep instantiation [3, 7]

NIVEL elaborates on its modelling constructs in various ways. The notion of deep instantiation previously defined for classes and attributes is extended to cover associations. A role in an association may be played by more than one class. Generalisations between associations are considered in detail, including cases in which roles are redefined in association subclasses. Both multiple classification and restricting to single classification are supported. Classes on a higher level exercise control over their instances whether and how they may participate in generalisations and define attributes. Although most of these ideas are not new to NIVEL, we believe that integrating them into a single yet relatively concise language with a formal semantics is a novel and significant contribution from a theoretical point of view and also a step towards practical applications.

A formal semantics is given for NIVEL by a translation to Weight Constraint Rule Language (WCRL) [32], a general purpose knowledge representation language for which efficient, decidable reasoning procedures are available. This enables decidable, automated reasoning about NIVEL. The modelling facilities of NIVEL are demonstrated using a running example and a case study on feature modelling, a modelling initiative popular in the software product family domain; the case study also highlights the utility of the formalisation.

The remainder of this paper is organised as follows. Next, in Section 2 we review a number of ideas related to meta-modelling that will be incorporated in NIVEL. Section 3 provides an overview of WCRL. Thereafter, NIVEL is introduced and formalised by translation to WCRL in Section 4. Section 5 presents a case study in which NIVEL is applied to feature modelling. Discussion and comparison to previous work follow in Section 6. Finally, conclusions and an outline for further work round up the paper in Section 7.

2 Modelling and meta-modelling

In this section, we provide an overview of terminology and ideas related to meta-modelling that will be incorporated in NIVEL in Section 4.

2.1 Model

In this paper, we adopt the view that being a model is a relationship between a *model* and an *original* [34, 22]. That is, being a model or an original is not an intrinsic property of an entity but roles played by them in a binary relationship.

A model is expressed using a *modelling language*. The term *language element* is used to refer to entities constituting a modelling language; examples include *proposition* and *connective* in propositional logic and *entity* and *relationship* in ER modelling. Entities constituting models are referred to as *model elements*.

2.2 Metaness

The notion of *metaness* often occurs when modelling is discussed, especially in the software engineering context. In this paper, we adopt a characterisation of metaness based on a binary relation R and the relation R^n defined for $n \geq 1$ as [22]:

$$e_1 R^n e_2 = \begin{cases} e_1 R e_2 & \text{if } n = 1 \\ \exists e : e_1 R^{n-1} e \wedge e R e_2 & \text{if } n \geq 2 \end{cases}$$

A *relation suitable for building metalevels* is required to satisfy three properties: the *acyclic*, *anti-transitive* and *level-respecting* properties [22]. The level-respecting property implies the anti-transitive property [22] and the acyclic property, see Appendix B. Hence, the level-respective property defined as [22]

$$\forall n, m : (\exists e_1, e_2 : e_1 R^n e_2 \wedge e_1 R^m e_2) \rightarrow n = m$$

suffices to characterise a relation suitable for building metalevels.

Given a relation R suitable for building metalevels, element e_1 can be characterised as *meta* with respect to e_2 if (and only if) $e_1 R^n e_2$ for $n \geq 2$. For example, if R is the binary relation $modelOf(m, o)$, m is termed a *model of original* o and $modelOf^2(m, o)$ implies that m is a *metamodel* of o . Further, each element m corresponds to a *level*, or *metalevel* or *layer*, alternative terms used. The level corresponding to m is said to be *above* the level corresponding to o whenever $modelOf^p(m, o)$ for $p \geq 1$; conversely, the level corresponding to o is said to be *below* the one corresponding to m . Finally, the terms *top* and *bottom* are used to refer to levels above and below all other levels, respectively.

If R is the *instanceOf*(i, t) relation, i is termed an *instance of* t and t a *type* of i . For $instanceOf^p(i, t)$, i is termed an *instance of order* p *of* t . Instances of order 2 or higher are collectively termed *higher-order instances*.

2.3 Strict metamodeling

The *strict metamodeling rule* has been defined as [4,5]:

In an n -level modeling architecture M_0, M_1, \dots, M_{n-1} , every element of an M_m -level model must be an instance-of exactly one element of an M_{m+1} -level model, for all $m < n-1$, and any relationship other than the instance-of relationship between two elements X and Y implies that $level(X) = level(Y)$.

The rule guarantees that the models constituting a metamodeling framework have well-defined boundaries; in the absence of such boundaries, the notion of metaness discussed in Section 2.2 could only be applied to individual model elements, not models.

Effectively, the rule requires that an element may not be an instance of another element two or more levels above. Also, the rule enforces single classification: an element may not be an instance of two or elements.

2.4 Ontological and linguistic instantiation

A typical metamodeling framework is based on a single form of instantiation. It has been argued that resorting to a single form of instantiation leads to a number of problems, most importantly *dual classification* [5]. As an example of dual classification, a model element, say *Collie* in Figure 1, is both an instance of *Class* and *Breed*; although both of these classifications are relevant, they are different in nature and, ideally, a metamodeling framework should acknowledge the difference.

As a remedy, it has been suggested that metamodeling frameworks should distinguish between two forms of classification (instantiation), termed the *logical* and *physical* dimension [5], or alternatively, *ontological* and *linguistic instantiation* [6,22]. In this paper, the latter set of terms is adopted.

As the term *linguistic* suggests, linguistic instantiation pertains to the (abstract syntax of the) language used to *represent* model elements or, in other words, the form model elements take [3,5]. In terms of model and language elements as defined and discussed in Section 2.1, linguistic instantiation is a relation between model and language elements where the former play the role of instance and the latter that of class or type. An ontological instantiation relationship relates two models (or model elements) that are in the same domain of discourse, e.g., biology or retail sales, but on different levels [22].

Figure 1 illustrates the distinction between ontological (dashed vertical arrows) and linguistic instantiation (dashed horizontal arrows). In the figure, *Lassie* is an ontological instance of *Collie* and a linguistic instance of *Object*. *Collie*, on the other hand, is an ontological instance of *Breed* and a linguistic instance of *Class*. *Breed* is a linguistic instance of *Metaclass*.

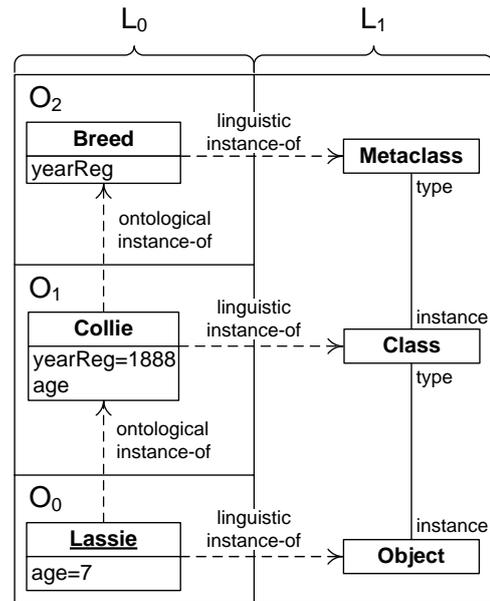


Fig. 1 Ontological and linguistic instantiation. Adapted from [22]

2.5 Unified modelling elements

In a metamodeling framework, model elements on a level are typically represented using the same language element: for example, in Figure 1, level O_0 is associated with *Object*, level O_1 with *Class* and level O_2 with *Metaclass*. Should another ontological level be added, a new language element would need to be introduced, probably termed *Metametaclass* in this case. If it is not possible to add language elements, the maximum number of available ontological levels is determined by the set of language elements at hand [5].

The notion of *unified modelling elements* has been suggested as a means to avoid undue restrictions of the above kind on the number of ontological levels. The notion is based on the observation that model elements on intermediate levels, i.e., other than the top and bottom level, are uniform in the sense that they exhibit both characteristics of classes and objects. Consequently, the model elements can be represented by a single language element, termed the *unified modelling element*. Further, model elements on the top and bottom levels can be viewed as special cases, top-level elements with a degenerate object and bottom-level elements with a degenerate type facet. In order to assign model elements to levels, the unified modelling element includes an integer attribute for level. [5]

Figure 1 has been redrawn as Figure 2, employing unified modelling elements. In Figure 2, the unified modelling element *ModelElement* unifies *Metaclass*, *Class* and *Object* in Figure 1. *Breed*, *Collie* and *Lassie* are all linguistic instances of *ModelElement*.

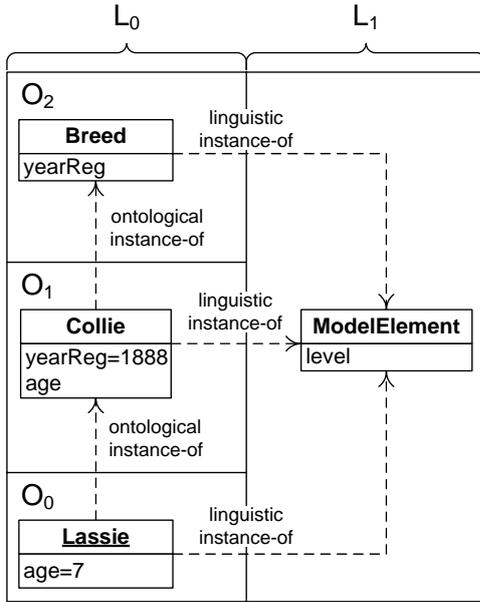


Fig. 2 Unified modelling elements

2.6 Deep instantiation

The case in which a model element may have higher-order instances is referred to as *deep instantiation* as opposed to the case of *shallow instantiation* where only first-order instances can be represented [3, 5].

The maximum order of instances of a class is specified by its *potency* [3, 5]. Potency is a non-negative integer assigned independently for each class that is not an instance of another class; the potency of a class is denoted by a superscript after its name, cf. Figure 3. An instance of a class has the potency value of its type decremented by one. However, the potency value may not be less than zero. Hence, the potency of a class gives the maximum order of its instances. A class under the notion of shallow instantiation is of potency 1.

Under deep instantiation, it is useful to enable a model element to characterise its higher-order instances. As an example, in Figure 3 the attribute denoted age^2 implies that each second order instance of *Breed*, e.g., *Lassie*, has a value for *age*; the intuition is that every dog must have an age. More generally, each attribute of a class is assigned an integer value termed *potency* that resembles the potency assigned for classes in that an instance of the class has an attribute by the same name as its type but with the potency decremented by one; an attribute of potency one turns into a value in an instance. Just as for classes, the potency of an attribute is denoted using a superscript after its name. [3, 5]

In Figure 3, the attribute *age* of *Breed* is of potency 2. Further, by virtue of being an instance of *Breed*, *Collie* has an attribute by the same name but of potency 1. Finally, a value (7) is assigned for *age* in *Lassie*.

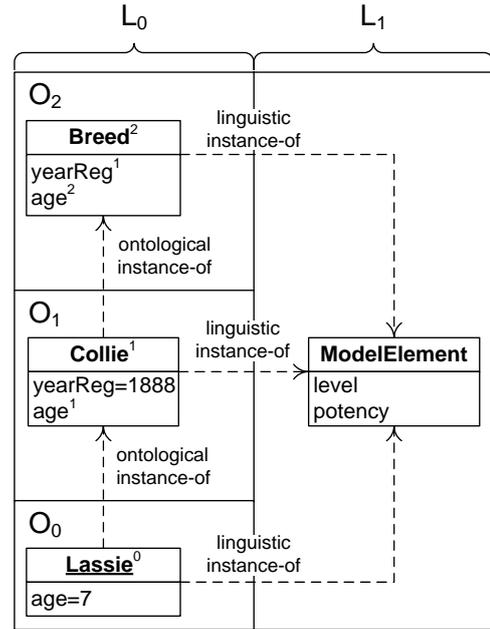


Fig. 3 Deep instantiation

3 Weight Constraint Rule Language

In this section, we provide a brief overview of Weight Constraint Rule Language (WCRL) based on [32], where the interested reader should refer for exact definitions, proofs and examples. WCRL will be used to give a formal semantics for NIVEL in Section 4.

WCRL is a general-purpose knowledge representation language similar to logic programs with a declarative formal semantics based on the stable model semantics of logic programs: the rules of the program are interpreted as constraints on a solution set for the program, whereas the usual logic programming paradigm is based on a goal-directed backward chaining query evaluation [29].

3.1 Syntax of weight constraint rules

A *weight constraint rule* is an expression of the form

$$C_0 \leftarrow C_1, \dots, C_n$$

where each C_i is a *weight constraint*. A weight constraint is of the form:

$$L \leq \{a_1 = w_{a_1}, \dots, a_n = w_{a_n}, \text{not } b_1 = w_{b_1}, \dots, \text{not } b_m = w_{b_m}\} \leq U$$

where each a_i and b_j is an *atom*. Atom a and not-atom $\text{not } b$ are called *literals*. Each literal is associated with a weight: e.g., w_{a_1} is the weight of a_1 . The numbers L and U are the *lower* and *upper bound* of the constraint, respectively.

A number of notational shorthands are used. Constraints where every weight has value 1, i.e., of the form

$$L \leq \{a_1 = 1, \dots, a_n = 1, \\ \text{not } b_1 = 1, \dots, \text{not } b_m = 1\} \leq U$$

are termed *cardinality constraints* and written as

$$L \{a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m\} U$$

A missing lower bound is interpreted as $-\infty$ and a missing upper bound as ∞ .

A rule of the form

$$\{a_1, \dots, a_n\} \leftarrow C_1, \dots, C_n$$

is termed a *choice rule*.

Constraints of the form $1 \leq \{l = 1\}$, where l is a literal, are written simply as l .

The shorthand

$$\leftarrow C_1, \dots, C_n$$

is used for rules where the head C_0 is an unsatisfiable constraint, such as $1 \leq \{\}$. This kind of rules are termed *integrity constraints*. Finally, rules with an empty body, that is, of the form

$$C_0 \leftarrow$$

are termed *facts*.

3.2 Stable model semantics

A set of literals S is defined to *satisfy* a weight constraint if (and only if):

$$L \leq \sum_{a_i \in S} w_{a_i} + \sum_{b_i \notin S} w_{b_i} \leq U$$

That is, a weight constraint is satisfied if the sum of weights of the literals satisfied by S is between the lower and upper bounds. A weight constraint rule is *satisfied* by a set of atoms S if and only if the head C_0 is satisfied whenever each constraint C_i in its body is satisfied. A *program* P is defined as a set of weight constraints. P is satisfied by S if and only if every weight constraint rule in P is satisfied by S .

A set of atoms S is a *stable model* of program P if (i) S satisfies P and (ii) S is *justified* by P . Intuitively, S is justified or *grounded* by P if every atom a in S has a non-circular justification. An atom is justified by a program if it appears in the head of a satisfied rule. Non-circular justification requires in addition that the body of the justifying rule is satisfied without assuming a .

Example Consider the program:

$$0 \leq \{a = 1, b = 1\} \leq 2 \leftarrow a$$

The set of atoms $\{a\}$ satisfies the only rule in program, and thus the program itself. However, the atom a only has a circular justification: the body is obviously not satisfied without assuming a . Hence, $\{a\}$ is not a stable model of the program. The only stable model is the empty set.

3.3 Rules with variables

For practical purposes, WCRL has been extended with a form of *variables* and *function symbols*. In more detail, *domain-restricted rules* with variables are allowed. Intuitively, a domain-restricted program P is divided into two parts: P_{Do} containing the definition of *domain predicates* and P_{Ot} containing all the other rules. The form of rules in P_{Do} is restricted in such a way that P_{Do} has a unique, finite stable model. All the rules in P_{Ot} must be domain-restricted in the sense that every variable occurring in a rule must appear in a domain predicate which occurs positively in the body of the rule. The domain predicates in P_{Do} are defined using stratified rules allowing a form of recursion [36].

A rule with variables is treated as a shorthand for all its ground instantiations with respect to the *Herbrand universe* of the program, i.e., the set of atoms that can be composed by applying functional composition from the basic symbols. The ground instantiation contains all the rules that can be obtained by substituting each variable in the rule with one of its possible values, restricted by one or more domain predicates. The atoms occurring in ground rules thus formed are termed the *Herbrand base* of the program and a model consisting of such atoms a *Herbrand model*.

Example Let us demonstrate different kinds of rules using a generalisation relation as an example, represented as a binary predicate *subclassOf*(a, b) in WCRL. Intuitively, the predicate gives that a is a *subclass of* b , and b is a *superclass of* a . In addition, we define the predicate *subclassOf_d*(a, b) with the semantics that a is a *direct subclass* of b . The unary domain predicate *class_p*(c) gives that c is a class.

To begin with, facts of the form

$$\text{class}_p(a) \leftarrow$$

can be used to define the set of classes of interest.

To specify that a direct subclass is also a non-direct subclass, we can write:

$$\text{subclassOf}(A, B) \leftarrow \text{subclassOf}_d(A, B), \\ \text{class}_p(A), \text{class}_p(B)$$

In the above rule, the literals *class_p*(A) and *class_p*(B) are needed to make it domain-restricted.

A choice rule can be used to enable direct generalisations between pairs of classes:

$$\{\text{subclassOf}_d(A, B)\} \leftarrow \text{class}_p(A), \text{class}_p(B)$$

That a class may not be a direct subclass of itself can be represented using the integrity constraint

$$\leftarrow \text{subclassOf}_d(A, A), \text{class}_p(A)$$

The possibility that a class (A) has two or more superclasses (B) can be ruled out using an integrity constraint:

$$\leftarrow 2 \{subclassOf(A, B) : class_p(B)\}, class_p(A)$$

The above rule also serves as an example of a rule including a *conditional literal*. Conditional literals are of the form $l : d$, where l is any predicate and d is a domain predicate. When instantiated, a conditional literal corresponds to the sequence of literals l' obtained by substituting the variables in l by all the combinations allowed by the domain predicate d .

3.4 Computational complexity and implementation

Given a ground WCRL program P , i.e., one not containing variables, and a set of atoms S , it can be decided in polynomial time whether S is a stable model of P . Deciding whether program P has a stable model is NP-complete.

An inference system, *smodels*¹, implements WCRL and has been shown to be competitive in performance compared with other solvers, especially in the case of problems including structure [32]. The *smodels* system includes built-in function symbols for integer arithmetic over domains restricted by domain predicates. This allows rules such as

$$weightOn(B, P, G \times M) \leftarrow mass(B, M), gravity(P, G)$$

where the predicate $mass(b, m)$ gives mass m of body b , the predicate $gravity(p, g)$ gravitational acceleration g on planet p and the predicate $weightOn(b, p, w)$ weight w of body b on planet p . The built-in function denoted by \times is the integer product.

The *smodels* system can be asked to compute a desired number of stable models of a program that agrees with a set of literals given as a *compute statement*.

4 Nivel—a metamodelling language

In this section, we introduce NIVEL, a metamodelling language capable of expressing models spanning an arbitrary number of levels. NIVEL is based on established conceptual modelling concepts: class, generalisation, attribute, value and association. Special emphasis is put on defining the instantiation relationship between classes, which enables NIVEL to be credited as a metamodelling language. In addition, NIVEL incorporates more recent metamodelling ideas discussed in Section 2 above.

The modelling facilities of NIVEL are demonstrated using a running example throughout the section and further in a case study on feature modelling in Section 5.

A formal semantics is given for NIVEL by translation to WCRL. The semantics enables automated, decidable reasoning on NIVEL which is essential in solving different computational tasks related to NIVEL. WCRL is particularly well suited as a knowledge representation language for NIVEL, as will be argued in Section 6.2. In short, the number of complexity of rules required is moderate and the translation is intuitive and modular.

This section begins with a description of the principles applied in the formalisation. Thereafter, the language elements of NIVEL are discussed. The section is concluded by an overview of computational problems related to NIVEL.

4.1 Formalisation principles

A formal semantics is given for NIVEL by translation to WCRL. In the terminology of [26], the approach followed is *translational semantics*: the syntactic constructs of a language are mapped onto constructs of another language, one that has a semantics defined. The symbol t will be used to denote this mapping. Further, the mapping denoted by t can be termed the *semantic mapping* and WCRL the *semantic domain* [17].

In symbols, we write: $t : \mathcal{N} \mapsto \mathcal{W}$, where \mathcal{N} denotes the set of syntactically well-formed NIVEL models and \mathcal{W} the set of WCRL programs. A model M entering the mapping t is termed an *input model*.

4.1.1 Input and valid models, possible and actual elements The formal semantics capture the notion of a *valid model*. Intuitively, the notion of validity pertains to the interrelations that may and must exist between model elements: for example, a class must not be an instance of itself. The mapping t is constructed in such a way that each stable model of $t(M)$, where $M \in \mathcal{N}$, corresponds to a valid NIVEL model.

An input model M need not be valid. In general, only a subset of the elements of M is found in each valid spun by M : an input model contains elements that may or may not appear in valid models. Intuitively, an input model spans a search space for valid models. We will use the terms *possible* and *actual* to refer to elements in input models and valid models, respectively.

The distinction between possible and actual entities is illustrated in Figure 4. Consider the model illustrated in Figure 4 (a) as an input model. It is assumed that each professor may possibly lecture any course. The input model contains all these possibilities, e.g., all the possible instances of the *lectures* association. There is also a constraint requiring that each course is lectured by a single professor and a professor must lecture at least one and at most two courses. A valid model must satisfy these constraints and may thus only include a subset of the possible elements. A valid instantiation of the *lectures* association is illustrated in Figure 4 (b).

¹ See <http://www.tcs.tkk.fi/Software/smodels/>

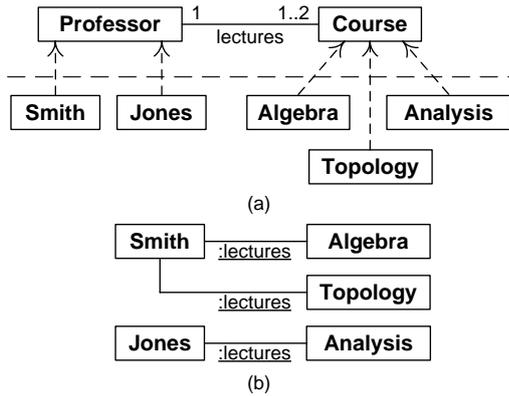


Fig. 4 Input and valid models, possible and actual model elements. (a) A NIVEL model spanning a number of valid models. It is assumed that each professor may possibly lecture any course. (b) A sample instantiation of the *lectures* association, assuming the model of part (a) as an input model.

4.1.2 Variants of predicates In the formalisation, it is necessary to distinguish between knowledge on input models and valid models, or in other words, between possible and actual elements. We draw this distinction by using different sets of predicates to represent these two classes of knowledge.

We adopt the convention that predicates referring to possible entities, i.e., those appearing in input models, are subscripted with p (for “possible”). These predicates will be defined in such a way that they are domain predicates; recall the definition from Section 3.3. As a special case of possible entities, predicates subscripted with D (for “declared”) are used to represent knowledge that is explicitly declared in an input model.

Knowledge on actual entities, i.e., those appearing in valid models, is represented using predicates without a subscript. For some language elements we will further distinguish between *direct* and other, *non-direct*, cases. The subscript d (for “direct”) is used to denote the direct case, also in conjunction with possible predicates. The capital subscript D implies both possibility and directness, the latter only if applicable.

4.1.3 A syntactic shorthand While input predicates are domain predicates, actual predicates are not. As can be recalled from Section 3.3, each variable occurring in a rule must appear in a domain predicate which occurs positively in the body of the rule. Consequently, in many rules concerning actual predicates, the same predicate will appear positively in the body applied to same arguments both as an actual variant and a possible variant needed to make the rule domain-restricted. The rules thus resulting may be unnecessarily difficult to read.

To improve readability, we adopt a syntactic shorthand: such a pair appearing in the body of a rule may be replaced by a single predicate with the subscripts joined using a comma (,); the dash (–) symbol is used when a

predicate without a subscript is joined. As an example,

$$pred_{d,D}(\dots)$$

can be used in place of

$$pred_d(\dots), pred_D(\dots)$$

4.1.4 Axiomatic and model/problem-specific rules The program $t(M)$ can be partitioned into two parts: an *axiomatic part* t_a and a part related to a particular model and a computational problem $t_{\mathcal{P}}(M)$.

The axiomatic part t_a contains rules shared by all models $M \in \mathcal{N}$. The axiomatic rules cover the interrelations between both possible and actual elements, and how possible elements may be actualised in valid models. The majority of the description of the formalisation given in Sections 4.2 through 4.9 is devoted to the axiomatic part. The mapping for possible elements will be illustrated in parallel with the axiomatic rules using a running example. Issues related to computational problems will be discussed in Section 4.10 and as a part of a case study in Section 5.2.

We now proceed to discuss the language elements of NIVEL. The abstract syntax of NIVEL is illustrated in Figure 5 and will be discussed in detail in the subsections on individual language elements to follow.

4.2 Class

Class is the most important language element in that all other language elements are directly related to class.

A *model* consists of a set of classes, some of which are *top-level* classes. The top-level classes are special in a number of ways to be detailed in the following subsections. A class is identified by a *name*.

Semantics A formalisation of NIVEL by translation to WCRL will be given under headings such as this. In describing the formalisation, we will concentrate on rules we believe to be most interesting to the reader; other rules are deferred to Appendix A. Table 1 contains a summary of the predicates used in the formalisation.

A class is represented using an object constant; an object constant may represent at most one class. The predicate $class(c)$ gives that the class represented by c is a class in a valid model. In the remainder of this paper we simply say “class c ” instead of “the class represented by c ”, and similarly for other model elements represented by object constants.

The unary domain predicate $topLevel(c)$ gives that class c is on the top level. A class declared to be on the top level is also a possible class:

$$class_p(C) \leftarrow topLevel_D(C)$$

and an actual class:

$$class(C) \leftarrow topLevel(C) \quad (1)$$

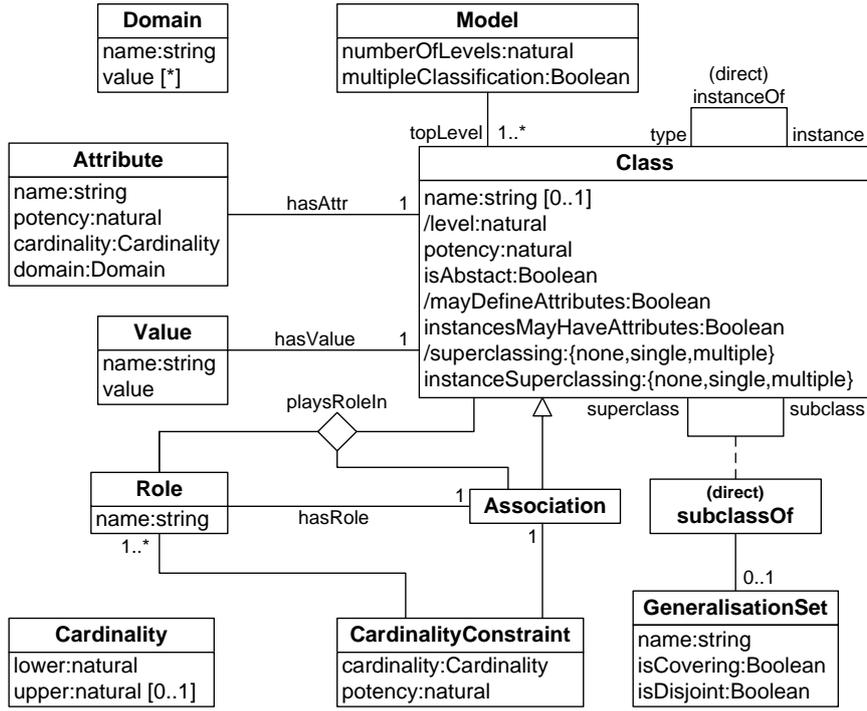


Fig. 5 The abstract syntax (language elements) of NIVEL

Hence, the top level is *fixed* in the sense that there is no distinction between entities possibly and actually on top level. The top level is also fixed in other respects: in each valid model corresponding to an input model, the top level is exactly as described in the input model. The rules capturing this idea are relatively simple and straightforward and are therefore deferred to Appendix A throughout the formalisation.

Example We will use a running example to demonstrate the concepts of NIVEL and their translation to WCRL. The example is illustrated in Figure 6, also containing a legend of the notation used. The corresponding WCRL rules are given as Figure 7

The running example includes a number of classes, *ProductType*, *Book*, *Novel*, *Video*, etc. In addition, as will be discussed in Section 4.8, an association is a special case of class and hence all the associations in the model, e.g., *adapts* and *dramatisation*, are classes as well.

Class *ProductType* and association *adapts* are declared to be on the top level as follows:

$$\begin{aligned} \text{topLevel}_D(\text{productType}) &\leftarrow \\ \text{topLevel}_D(\text{adapts}) &\leftarrow \end{aligned}$$

4.3 Generalisation

A generalisation is a binary relationship between classes, represented by the *subclassOf* relation. More specifically, if *subclassOf*(*a*, *b*) holds, we say that *a* is a *subclass* of *b* and that *b* is a *superclass* of *a*. We also introduce the notion of *direct subclass*, represented using the *subclassOf_d*

relation, and use the terms *direct subclass* and *direct superclass* in their obvious meanings.

The *subclassOf* relation is required to be antisymmetric. That is, it cannot be the case that two classes, *a* and *b*, are mutually subclasses of each other, unless *a* equals *b*.

We term a class that is on the top-level and does not have any superclasses a *root class*.

Semantics The predicates *subclassOf* and *subclassOf_d* are used in the meaning described above.

The predicate *subclassOf* is the reflexive closure of *subclassOf_d*, which is represented as follows, for the possible case:

$$\text{subclassOf}_p(C, C) \leftarrow \text{class}_p(C) \quad (2)$$

and actual elements:

$$\text{subclassOf}(C, C) \leftarrow \text{class}_{-,p}(C) \quad (3)$$

Similarly for the transitive closure, we write for the possible case:

$$\begin{aligned} \text{subclassOf}_p(C_{sub}, C_{super}) &\leftarrow \text{class}_p(C_{super}), \\ \text{subclassOf}_D(C_{sub}, C), \text{subclassOf}_p(C, C_{super}) & \quad (4) \end{aligned}$$

and for the actual case:

$$\begin{aligned} \text{subclassOf}(C_{sub}, C_{super}) &\leftarrow \\ \text{subclassOf}_{d,D}(C_{sub}, C), \text{subclassOf}_{-,p}(C, C_{super}) & \quad (5) \end{aligned}$$

Observe the close resemblance between Rules 2 and 3, and Rules 4 and 5. Rules for the actual variant (Rules 3

Table 1 Predicates used in the formalisation in the order they appear in text

Predicate	Variants	Semantics
<i>class</i> (<i>c</i>)	– <i>p</i>	Object constant <i>c</i> represents a class
<i>topLevel</i> (<i>c</i>)	– <i>D</i>	Class <i>c</i> is on top level
<i>subclassOf</i> (<i>a</i> , <i>b</i>)	– <i>p d D</i>	Class <i>a</i> is a subclass of class <i>b</i>
<i>superclassingSingle</i> (<i>c</i>)		At most one superclass may be defined for class <i>c</i>
<i>superclassingMultiple</i> (<i>c</i>)		Any number of superclasses may be defined for class <i>c</i>
<i>instanceOf</i> (<i>i</i> , <i>t</i>)	– <i>p d pd D</i>	Class <i>i</i> is an instance of class <i>t</i>
<i>instanceOf</i> (<i>i</i> , <i>t</i> , <i>o</i>)	<i>tp tpd</i>	Class <i>i</i> is an instance of class <i>t</i> of order <i>o</i>
<i>singleClassification</i>		No multiple classification is allowed
<i>abstract</i> (<i>c</i>)		Class <i>c</i> is abstract
<i>onLevel</i> (<i>c</i> , <i>l</i>)		Class <i>c</i> is on level <i>l</i>
<i>level</i> (<i>l</i>)		Natural number <i>l</i> is a possible level number
<i>hasPotency</i> (<i>c</i> , <i>p</i>)	– <i>D</i>	Class <i>c</i> is of potency <i>p</i>
<i>instanceSuperclassingSingle</i> (<i>c</i>)		At most one superclass may be defined for instances of type <i>t</i>
<i>instanceSuperclassingMultiple</i> (<i>c</i>)		Any number of superclasses may be defined for instances of type <i>t</i>
<i>inGSet</i> (<i>sub</i> , <i>super</i> , <i>g</i>)	– <i>D</i>	Class <i>sub</i> is a subclass of class <i>super</i> in generalisation set <i>g</i>
<i>inGSet</i> (<i>sub</i> , <i>g</i>)	– <i>D</i>	Class <i>sub</i> is a subclass in generalisation set <i>g</i>
<i>gSetOf</i> (<i>g</i> , <i>super</i>)	– <i>D</i>	Generalisation set <i>g</i> is a generalisation set of class <i>super</i>
<i>gSet</i> (<i>g</i>)		Constant symbol <i>g</i> represents a generalisation set
<i>covering</i> (<i>g</i>)		Attribute <i>isCovering</i> is <i>true</i> for generalisation set <i>g</i>
<i>disjoint</i> (<i>g</i>)		Attribute <i>isDisjoint</i> is <i>true</i> for generalisation set <i>g</i>
<i>contains</i> (<i>d</i> , <i>v</i>)		Domain <i>d</i> contains value <i>v</i>
<i>hasAttr</i> (<i>c</i> , <i>n</i> , <i>p</i> , <i>d</i> , <i>l</i> , <i>u</i>)	– <i>p D</i>	Class <i>c</i> has attribute named <i>n</i> with potency <i>p</i> , domain <i>d</i> , cardinality lower bound <i>l</i> and upper bound <i>u</i>
<i>mayDefineAttributes</i> (<i>c</i>)		Class <i>c</i> may be defined attributes
<i>instancesMayDefineAttributes</i> (<i>c</i>)		Attributes may be defined for instances of class <i>c</i>
<i>hasValue</i> (<i>c</i> , <i>n</i> , <i>v</i>)	– <i>p D</i>	Class <i>c</i> has value <i>v</i> under name <i>n</i>
<i>association</i> (<i>a</i>)		Class <i>a</i> is an association
<i>hasRole</i> (<i>a</i> , <i>r</i>)	– <i>p D</i>	Association <i>a</i> has role <i>r</i>
<i>playsRoleIn</i> (<i>c</i> , <i>r</i> , <i>a</i>)	– <i>d p D</i>	Class <i>c</i> plays role <i>r</i> in association <i>a</i> ; (<i>c</i> , <i>r</i> , <i>a</i>) is a roleplay
<i>rolePlayed</i> (<i>r</i> , <i>a</i>)		Role <i>r</i> is played in association <i>a</i>
<i>validRolePlay</i> (<i>c</i> , <i>r</i> , <i>a_i</i> , <i>a_t</i>)		The roleplay (<i>c</i> , <i>r</i> , <i>a_i</i>) is valid with respect to type <i>a_t</i>

Note: Symbols in the *Variants* column: ‘*p*’ – possible, ‘*d*’ – direct, ‘*D*’ – declared, ‘*t*’ – transitive, ‘–’ – actual. In case more than one variant is defined, the semantics are described for the actual variant.

and 5) can be roughly obtained by changing the possible predicates to actual ones and appending an appropriate possible predicate to the actual one in the body: a possible predicate (denoted by subscript *p*) is typically appended to a non-direct actual one (missing subscript) and a declared predicate (*D*) to a direct one (*d*).

A resemblance of the above-described kind is a recurrent pattern in the formalisation. Therefore, we simplify the discussion by deferring most definitions for possible variants to Appendix A.

The predicate *subclassOf* is antisymmetric:

$$\leftarrow \text{subclassOf}_{-p}(A, B), \text{subclassOf}_{-p}(B, A), A \neq B$$

The notion of root class coincides with the predicate *topLevel_D* defined above.

Example The running example includes one case of subclassing: the class *Book* is a superclass of two classes, *Novel* and *Anthology*. The latter case is represented as:

$$\text{subclassOf}_D(\text{anthology}, \text{book}) \leftarrow \square$$

There are restrictions on whether and how many superclasses a class may have. These restrictions are imposed by the *superClassing* attribute of *Class*. The attribute takes one of the values *none*, *single* and *multiple* with the intuitive semantics that a class may have no, at most one or any number of superclasses, respectively.

A value for *superClassing* is not independently assigned for classes but derived in a way to be detailed in Section 4.4.5. Top-level classes are a special case: any number of superclasses may be defined for them.

Semantics The predicate *superclassingSingle*(*c*) gives that the attribute *superclassing* of *c* has value *single*. Similarly, the value *multiple* is implied by the predicate *superclassingMultiple*(*c*). The value *none* is the default in the sense that it is implied by neither of the above-mentioned predicates holding for a class.

Direct generalisations are justified by the choice rule:

$$\begin{aligned} & \{ \text{subclassOf}_d(C_{sub}, C_{super}) \} \leftarrow \\ & \text{subclassOf}_D(C_{sub}, C_{super}), \\ & \text{superclassingSingle}(C_{sub}), \text{class}(C_{super}) \end{aligned} \quad (6)$$

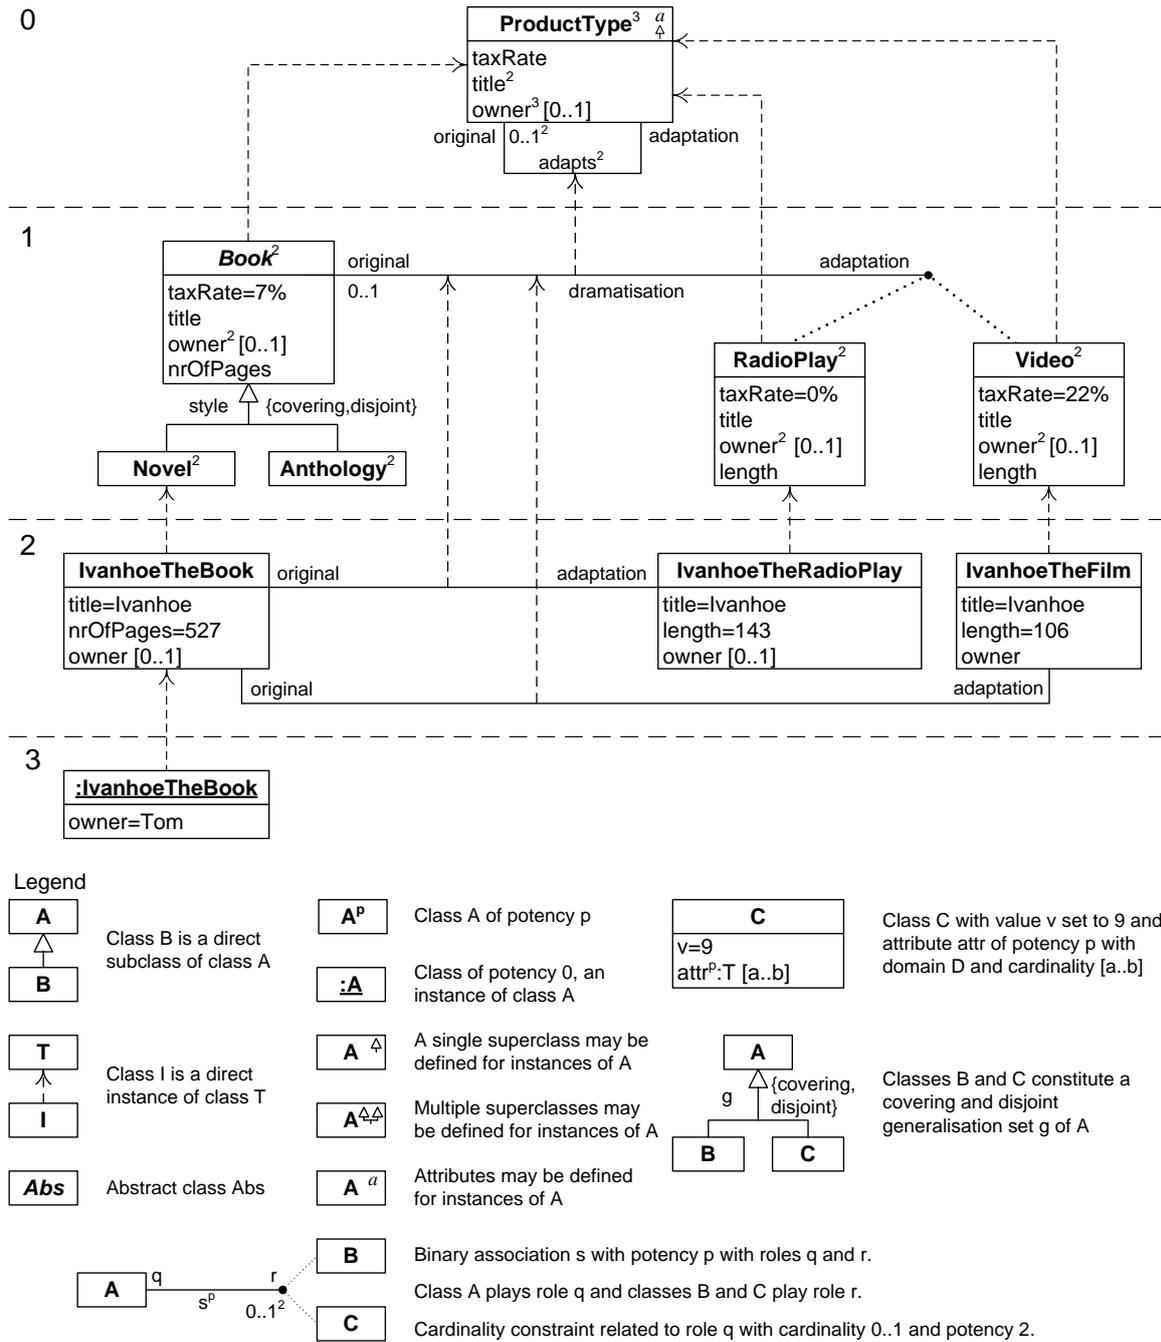


Fig. 6 A sample NIVEL model

Only classes with multiple allowed superclasses may have two or more of them:

$$\leftarrow 2 \{ subclassOf_d(C_{sub}, C_{super}) : subclassOf_D(C_{sub}, C_{super}) \}, not\ superclassingMultiple(C_{sub}), class_p(C_{sub})$$

4.4 Instantiation

Instantiation is a binary relationship between classes. Whenever $instanceOf(i, t)$ holds, we say that i is an instance of t and t is a type of i . We also introduce the

notion of *direct instantiation* and use the terms *direct instance* and *direct type*.

4.4.1 Single and multiple classification A NIVEL class may have multiple direct types; in other words, multiple classification is allowed. However, a model may be confined to single classification by setting the attribute *multipleClassification* of *Model* to *false*.

Semantics The binary predicate $instanceOf(i, t)$ has the semantics that class i is an instance of class t . Similarly, $instanceOf_d(i, t)$ gives that i is a direct instance

```

const maxLevel = 3.
singleClassification ←

topLevelD(productType) ←
hasPotencyD(productType, 3) ←
instancesMayDefineAttributes(productType) ←
instanceSuperclassingSingle(productType) ←

hasAttrD(productType, taxRate, 1, trDomain, 1, 1) ←
hasAttrD(productType, title, 2, titleDomain, 1, 1) ←
hasAttrD(productType, owner, 3, ownerDomain, 1, 1) ←

contains(trDomain, 0) ←
contains(trDomain, 7) ←
contains(trDomain, 22) ←

contains(titleDomain, ivanhoe) ←
contains(titleDomain, senseAndSensibility) ←
contains(titleDomain, lassieReturns) ←

contains(ownerDomain, tom) ←
contains(ownerDomain, dick) ←
contains(ownerDomain, harry) ←

topLevelD(adapts) ←
association(adapts) ←
hasPotencyD(adapts, 2) ←
hasRoleD(adapts, original) ←
hasRoleD(adapts, adaptation) ←
playsRoleInD(productType, original, adapts) ←
playsRoleInD(productType, adaptation, adapts) ←

instanceOfD(book, productType) ←
hasValueD(book, taxRate, 7) ←
hasAttrD(book, nrOfPages, 1, pageDomain, 1, 1) ←
contains(pageDomain, 520..530) ←
abstract(book) ←

inGSetD(novel, book, style) ←
inGSetD(anthology, book, style) ←
covering(style) ←
disjoint(style) ←

instanceOfD(video, productType) ←
hasValueD(video, taxRate, 22) ←
hasAttrD(video, length, 1, lengthDomain, 1, 1) ←
contains(lengthDomain, 105..110) ←

instanceOfD(radioPlay, productType) ←
hasValueD(radioPlay, taxRate, 0) ←
hasAttrD(radioPlay, length, 1, lengthDomain, 1, 1) ←
contains(lengthDomain, 140..145) ←

instanceOfD(dramatisation, adapts) ←
playsRoleInD(book, original, dramatisation) ←
playsRoleInD(video, adaptation, dramatisation) ←
playsRoleInD(radioPlay, adaptation, dramatisation) ←

instanceOfD(ivBook, novel) ←
hasValueD(ivBook, title, ivanhoe) ←
hasValueD(ivBook, nrOfPages, 527) ←

instanceOfD(ivFilm, video) ←
hasValueD(ivFilm, title, ivanhoe) ←
hasValueD(ivFilm, length, 106) ←

instanceOfD(ivRadioPlay, radioPlay) ←
hasValueD(ivRadioPlay, length, 143) ←
hasValueD(ivRadioPlay, title, ivanhoe) ←

instanceOfD(ivFilm, dramatisation) ←
playsRoleInD(ivBook, original, ivFilm) ←
playsRoleInD(ivFilm, adaptation, ivFilm) ←

instanceOfD(ivRadioPlayAdpt, dramatisation) ←
playsRoleInD(ivBook, original, ivRadioPlayAdpt) ←
playsRoleInD(ivRadioPlay, adaptation, ivRadioPlayAdpt) ←

instanceOfD(tomsBook, ivBook) ←
hasValueD(tomsBook, owner, tom) ←

adapts(C, O) ← instanceOftp(A, adapts, 2),
playsRoleIn-p(C, adaptation, A),
playsRoleIn-p(O, original, A),

adaptsp(C, O) ← instanceOftp(A, adapts, 2),
playsRoleInp(C, adaptation, A),
playsRoleInp(O, original, A)

← 2 { adapts(Ci, O) : adaptsp(Ci, O) },
playsRoleIn-p(Ct, adaptation, A),
instanceOftp(A, adapts, 1),
instanceOf-p(Ci, Ct)

```

Fig. 7 The model-specific rules resulting from applying the mapping t to the model in Figure 6

of t . The atom *singleClassification* gives that multiple classification is not allowed.

Possible, direct instantiations may become actual, direct instantiations:

$$\{instanceOf_d(I, T)\} \leftarrow instanceOf_{pd}(I, T), \quad (7)$$

$$class_{-p}(T)$$

A direct instance is a class:

$$class(I) \leftarrow instanceOf_{d,pd}(I, T) \quad (8)$$

Multiple classification, if *singleClassification* holds, is ruled out as follows:

$$\leftarrow 2 \{ instanceOf_d(I, T) : instanceOf_{pd}(I, T) \},$$

$$class_{-p}(I), singleClassification$$

Example The running example is restricted to single classification:

$$singleClassification \leftarrow$$

The example contains a large number of declared instantiations. Only instantiations shown using the dashed line notation in Figure 6 need be explicitly represented in WCRL. As an example, that *Book* is a direct instance of *ProductType* is represented as follows:

$$\leftarrow \text{instanceOf}_D(\text{book}, \text{productType})$$

4.4.2 Abstractness A class may be *abstract*, represented by the attribute *isAbstract* set to *true*, with the semantics that an abstract class may not have direct instances.

Semantics The predicate *abstract*(*c*) gives that class *c* is abstract. The semantics of abstractness is captured by:

$$\leftarrow \text{abstract}(T), \text{instanceOf}_{d,pd}(I, T)$$

Example Class *Book* is abstract:

$$\text{abstract}(\text{book})$$

4.4.3 Level Each class is on exactly one *level* identified by a natural number. The level of a class is represented using the *level* attribute. The top level corresponds to level 0. The level of other than top-level classes is determined based on the instantiation relation: if class *i* is an instance of class *t* on level *l*, class *i* is on level *l* + 1.

Semantics Each level is represented by its number. The binary predicate *onLevel*(*c*, *l*) gives that class *c* is on level *l*. The top level is assigned number 0 by the rule:

$$\text{onLevel}(C, 0) \leftarrow \text{topLevel}(C)$$

If class *i* is an instance of class *t* on level *l*, class *i* is on level *l* + 1:

$$\begin{aligned} \text{onLevel}(I, L + 1) &\leftarrow \text{instanceOf}_{-p}(I, T), \\ &\text{onLevel}(T, L), \text{level}(L) \end{aligned} \quad (9)$$

Above, the domain predicate *level*(*l*) gives that *l* is a level.

The requirement that a class must not be on two or more levels is represented using an integrity constraint:

$$\leftarrow 2 \{ \text{onLevel}(C, L) : \text{level}(L) \}, \text{class}_{-p}(C) \quad (10)$$

4.4.4 Potency Each class must have exactly one natural number as its *potency*. A potency is independently assigned for root classes. Other top-level classes have the same potency as their superclasses. If class *i* is an instance of class *t* of potency *p*, class *i* is of potency *p* - 1. That is, the potency value is decremented by one when a class is instantiated. A class of potency 0 cannot be instantiated. Hence, the potency value gives the maximum order of instances a class may have.

Semantics The predicate *hasPotency*(*c*, *p*) gives that class *c* is of potency *p*. Potency is decremented by one when a class is instantiated:

$$\begin{aligned} \text{hasPotency}(I, P - 1) &\leftarrow \text{instanceOf}_{d,pd}(I, T), \\ &\text{hasPotency}(T, P), \text{level}(P), P > 0 \end{aligned}$$

A class must have a single potency value. That is, it may not be the case that class has no potency value:

$$\leftarrow \{ \text{hasPotency}(C, P) : \text{level}(P) \} 0, \text{class}_{-p}(C)$$

or two or more potency values:

$$\leftarrow 2 \{ \text{hasPotency}(C, P) : \text{level}(P) \}, \text{class}_{-p}(C)$$

Example The top-level classes *ProductType* and *adapts* are declared potencies as follows:

$$\begin{aligned} \text{hasPotency}_D(\text{productType}, 3) &\leftarrow \\ \text{hasPotency}_D(\text{adapts}, 2) &\leftarrow \end{aligned}$$

4.4.5 Instantiations implied by generalisations A generalisation may imply instantiations, both direct and non-direct ones. The direct case is illustrated in Figure 8 (a): if *I_{super}* is a direct instance of *T* and *I_{sub}* is a subclass of *I_{super}*, then *I_{sub}* is a direct instance of *T*.

However, there is an exception to this rule, illustrated in Figure 8 (b): if *I_{sub}* is a direct instance of *T_{sub}* that is a subclass of *T_{super}*, a direct instantiation between *I_{sub}* and *T_{super}* is *not* implied; however, a non-direct instantiation between the pair (*I_{sub}*, *T_{super}*) is implied. More generally, we require that class *i* may not be a direct instance of class *t* and a superclass *t'* of *t*, even if multiple classification is allowed.

Non-direct instantiations are implied analogously but without a similar exception, see Figure 8 (c). If class *I_{super}* is a (direct) instance of class *T_{sub}*, then a subclass *I_{sub}* of *I_{super}* is an instance of each superclass *T_{super}* of *T_{sub}*, including *T_{sub}* itself.

Semantics Direct generalisations may be implied:

$$\begin{aligned} \{ \text{instanceOf}_d(I_{sub}, T) \} &\leftarrow \\ \text{instanceOf}_{d,pd}(I_{super}, T), & \\ \text{subclassOf}_{-p}(I_{sub}, I_{super}) & \end{aligned} \quad (11)$$

False cases of implied direct instantiation are ruled out by the integrity rule

$$\begin{aligned} \leftarrow \text{instanceOf}_{d,pd}(I, T_{sub}), \text{instanceOf}_{d,pd}(I, T_{super}), \\ \text{subclassOf}_{-p}(T_{sub}, T_{super}), T_{sub} \neq T_{super} \end{aligned}$$

Non-direct instantiations are always implied:

$$\begin{aligned} \text{instanceOf}(I_{sub}, T_{super}) &\leftarrow \\ \text{instanceOf}_{d,pd}(I_{super}, T_{sub}), & \\ \text{subclassOf}_{-p}(T_{sub}, T_{super}), & \\ \text{subclassOf}_{-p}(I_{sub}, I_{super}) & \end{aligned} \quad (12)$$

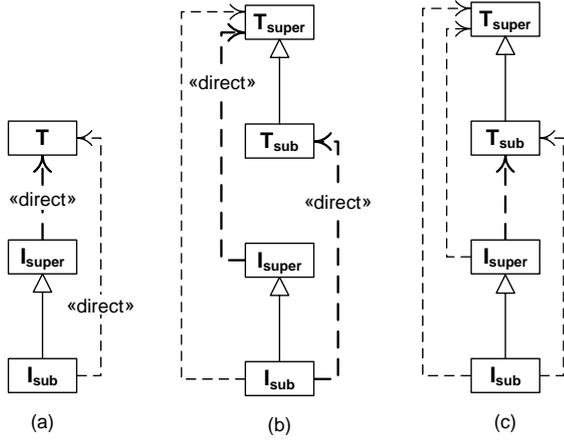


Fig. 8 The implications of subclassing on instantiation. Instantiations depicted using a thick line are assumed, others are implied. (a) Direct instantiations implied. (b) Only non-direct instantiation implied between I_{sub} and T_{super} . (c) Non-direct instantiations implied.

Example A number of instantiation relationships are implied in the running example: e.g, *IvanhoeTheBook* is, in addition to being a direct instance of *Novel*, a non-direct instance of *Book*. *Novel* is a direct instance of *ProductType*. \square

The types of a class exercise control whether and how many superclasses may be defined for the class. The relevant language construct is the *instanceSuperclassing* attribute of *Class* taking one of the values *none*, *single* and *multiple*. Superclass(es) may be defined for a class if at least one of its direct types enables this.

Semantics That a single superclass may be defined for the direct instances of class c is given by the predicate *instanceSuperclassingSingle*(c). Similarly, the predicate *instanceSuperclassingMultiple*(c) gives that any number of superclasses may be defined for the instances. The former is formalised as follows:

$$\begin{aligned} \text{superclassingSingle}(I) &\leftarrow \text{instanceOf}_{-pd}(I, T), \\ \text{instanceSuperclassingSingle}(T) \end{aligned} \quad (13)$$

The definition of the latter is similar and deferred to Appendix A.

Example A single superclasses may be defined for instances of *ProductType*:

$$\text{instanceSuperclassingSingle}(\text{productType}) \leftarrow$$

4.5 Generalisation set

A generalisation may belong to a *generalisation set*. A generalisation set is identified by a *name*. All the generalisations in a generalisation set must share the same superclass. Given a generalisation set g containing generalisation $(sub, super)$, we say that g is a *generalisation set of super* and that sub is *in generalisation set g*.

Semantics A generalisation set is represented by an object constant. The predicate *inGSet*($sub, super, g$) gives that generalisation $(sub, super)$ belongs to generalisation set g .

A generalisation declared in the context of a generalisation set is a declared generalisation:

$$\text{subclassOf}_D(C_{sub}, C_{super}) \leftarrow \text{inGSet}_D(C_{sub}, C_{super}, G)$$

A generalisation possibly in a generalisation set may actually be in it, provided that the generalisation actually holds:

$$\begin{aligned} \{ \text{inGSet}(C_{sub}, C_{super}, G) \} &\leftarrow \\ \text{inGSet}_D(C_{sub}, C_{super}, G), & \\ \text{subclassOf}_d(C_{sub}, C_{super}) \end{aligned} \quad (14)$$

Example The classes *Novel* and *Anthology* constitute a generalisation set *style* of *Book*. This is represented by the facts:

$$\begin{aligned} \text{inGSet}_D(\text{novel}, \text{book}, \text{style}) &\leftarrow \\ \text{inGSet}_D(\text{anthology}, \text{book}, \text{style}) &\leftarrow \quad \square \end{aligned}$$

Disjointness and *covering constraints* may apply to a generalisation set, imposed by attributes *isCovering* and *isDisjoint* of *GeneralisationSet*.

If *isCovering* is set to *true*, the generalisation set is said to be *covering* and an instance of the superclass must be an instance of (at least) one of the subclasses. Similarly, if *isDisjoint* is *true*, the generalisation set is said to be *disjoint* and an instance of the superclass may not be an instance of two or more of the subclasses.

Semantics The predicate *covering*(g) gives that generalisation set g is covering. The semantics is defined as an integrity constraint:

$$\begin{aligned} &\leftarrow \{ \text{instanceOf}(I, T_{sub}) : \text{inGSet}_D(T_{sub}, G) \} 0, \\ &\quad \text{covering}(G), \\ &\text{instanceOf}_{-p}(I, T_{super}), \text{gSetOf}(G, T_{super}) \end{aligned}$$

Similarly for the case of disjointness, the predicate *disjoint*(g) gives that generalisation set g is disjoint. The related constraint is:

$$\begin{aligned} &\leftarrow 2 \{ \text{instanceOf}(I, T_{sub}) : \text{inGSet}_D(T_{sub}, G) \}, \\ &\quad \text{disjoint}(G), \\ &\text{instanceOf}_{-p}(I, T_{super}), \text{gSetOf}(G, T_{super}) \end{aligned}$$

Example The generalisation set *style* of *Book* is both covering and disjoint:

$$\begin{aligned} \text{covering}(\text{style}) &\leftarrow \\ \text{disjoint}(\text{style}) &\leftarrow \end{aligned}$$

4.6 Attribute

A class may be characterised by *attributes*. An attribute is described by a *name*, a *potency*, a *cardinality* and a *domain*. The name is used to distinguish the attributes of a class from one other. Potency is a non-zero natural number. The cardinality consists of a *lower* and an optional *upper* bound, both of which are natural numbers. The domain of an attribute is a set of values identified by a name.

Given a class with an attribute, an instance of the class has an attribute by the same name and with the same domain as its type but with the potency decremented by one, for potencies greater than 1.

Attributes are *inherited* by subclasses: if a class has an attribute, its subclasses have the same attribute.

Semantics Attribute names, domains and values contained in domains are all represented using object constants. The predicate $contains(d, v)$ gives that domain d contains value v .

The predicate $hasAttr(c, n, p, d, l, u)$ gives that class c has attribute named n with potency p , domain d and cardinality with lower bound l and upper bound u . If the upper bound is omitted, a special symbol ∞ is used in the formalisation.

A direct instance has the attributes of its type, with potencies decremented by one:

$$\begin{aligned} hasAttr(I, N, P - 1, D, L, U) \leftarrow \\ hasAttr_{-p}(T, N, P, D, L, U), \\ instanceOf_{d, pd}(I, T), P > 1 \end{aligned} \quad (15)$$

A subclass inherits the attributes of its superclass:

$$\begin{aligned} hasAttr(C_{sub}, N, P, D, L, U) \leftarrow \\ hasAttr_{-p}(C_{super}, N, P, D, L, U), \\ subclassOf_{-p}(C_{sub}, C_{super}) \end{aligned} \quad (16)$$

Example The running example includes a number of attributes: e.g., the attribute *title* of *ProductType* is represented as follows:

$$hasAttr_D(productType, title, 2, titleDomain, 1, 1) \leftarrow$$

As an example of domains, *pageDomain* is represented as follows:

$$contains(pageDomain, 520..530) \leftarrow \quad \square$$

A class exercises control over whether attributes may be defined for its instances or not. A class may enable its direct instances to define attributes by setting its Boolean attribute *instancesMayDefineAttributes* to *true*. As a result, the attribute *mayDefineAttributes* will be *true* for its instances, with the semantics that attributes may be defined for them. Attributes may always be defined for top level classes.

Semantics A declared attribute of class c for which $mayDefineAttributes(c)$ holds may become an actual attribute:

$$\begin{aligned} \{hasAttr(C, N, P, D, L, U)\} \leftarrow \\ hasAttr_D(C, N, P, D, L, U), \\ mayDefineAttributes(C) \end{aligned} \quad (17)$$

Above, the predicate $mayDefineAttributes(c)$ gives that attributes may be defined for class c :

$$\begin{aligned} mayDefineAttributes(I) \leftarrow instanceOf_{d, pd}(I, T), \\ instancesMayDefineAttributes(T) \end{aligned}$$

Note that in the case of multiple classification, a class may be defined attributes if at least one of its direct types allows it.

Example Attributes may be defined for instances of *ProductType*:

$$instancesMayDefineAttributes(productType) \leftarrow$$

4.7 Value

When a class with an attribute of potency 1 is instantiated, the instances may and must have *values* corresponding to the attribute. The name of the value is the same as that of the attribute. The number of such values must agree with the cardinality of the attribute: there must be at least *lower* and at most *upper* values by the name of the attribute. Finally, each of the values must be in the domain of the attribute.

Semantics The predicate $hasValue(c, n, v)$ gives that class c has value v under name n .

Attributes of potency 1 result in values. For an attribute with and upper bound, i.e., $U \neq \infty$

$$\begin{aligned} L \{hasValue(I, N, V) : contains(D, V) : \\ hasValue_p(I, N, V)\} U \leftarrow \\ hasAttr_{-p}(T, N, 1, D, L, U), U \neq \infty, \\ instanceOf_{d, pd}(I, T) \end{aligned} \quad (18)$$

Similarly for an attribute without an upper bound ($U = \infty$):

$$\begin{aligned} L \{hasValue(I, N, V) : contains(D, V) : \\ hasValue_p(I, N, V)\} \leftarrow \\ hasAttr_{-p}(T, N, 1, D, L, U), U = \infty, \\ instanceOf_{d, pd}(I, T) \end{aligned}$$

Example The fact below gives that *Book* has value 7 declared for the attribute *taxRate*:

$$hasValue_D(book, taxRate, 7) \leftarrow$$

4.8 Association

An association is a relationship between a set of classes. An association defines a set of *roles* each of which is *played* by one or more classes. Specifically, a role can be played by multiple classes that are not subclasses of a single superclass. An association is also a class. Hence, all the language elements related to class discussed above apply to association as well: associations may participate in generalisations and generalisation sets, be instantiated and have attributes and values.

Semantics An association is also a class and is consequently represented by an object constant. The predicate $association(a)$ gives that a is an association.

Example The fact

$$association(adapts) \leftarrow$$

gives that $adapts$ is an association.

4.8.1 Role The distinguishing characteristic of association is that they may have roles. A role is identified by a *name* that must be unique in the context of an association. The set of roles of an association is fixed in the sense that all the subclasses and instances of an association must have the same set of roles as the association itself: roles may only be defined for root classes.

Each role of an association must be *played* by at least one class. If class c plays role r in association a , we say that the triple (c, r, a) is a *roleplay*. Further, an instance of c may play r in instances of a .

Semantics The predicate $hasRole(a, r)$ gives that association a has role r . An instance of an association has the roles of its type:

$$\begin{aligned} hasRole(I, R) &\leftarrow hasRole_{-p}(T, R), \\ instanceOf_{-p}(I, T) & \end{aligned} \quad (19)$$

The ternary predicate $playsRoleIn(c, r, a)$ gives that class c plays role r in association a . The following rule gives that a declared roleplay (c_i, r, a_i) may be an actual roleplay, provided that (c_t, r, a_t) is, where c_t is a type of c_i and a_t is a type of a_i :

$$\begin{aligned} \{playsRoleIn_d(C_i, R, A_i)\} &\leftarrow \\ &playsRoleIn_D(C_i, R, A_i), \\ &playsRoleIn_{-p}(C_t, R, A_t), \\ instanceOf_{-p}(C_i, C_t), &instanceOf_{d, pd}(A_i, A_t) \end{aligned} \quad (20)$$

The predicate $rolePlayed(r, a)$ gives that role r is played by at least one class in association a :

$$rolePlayed(R, A) \leftarrow playsRoleIn_{-p}(C, R, A)$$

All the roles of an association must be played:

$$\begin{aligned} \leftarrow not\ rolePlayed(R, A), &hasRole_{-p}(A, R), \\ &association(A) \end{aligned}$$

Example Roles *original* and *adaptation* are declared for the association *adapts*:

$$\begin{aligned} hasRole_D(adapts, original) &\leftarrow \\ hasRole_D(adapts, adaptation) &\leftarrow \end{aligned}$$

Both roles are played by *ProductType*:

$$\begin{aligned} playsRoleIn_D(productType, original, adapts) &\leftarrow \\ playsRoleIn_D(productType, adaptation, adapts) &\leftarrow \end{aligned}$$

Roleplays on lower levels are declared similarly, see Figure 7 for the facts.

4.8.2 Association generalisation Given that an association is also a class, there can be generalisations between associations and all the properties of generalisations discussed above apply to generalisations between associations as well. In this section, we will discuss how roles and the fact that roles are played by classes are affected by generalisations.

By default, a subclass of an association inherits the roleplays of its superclasses. However, a role can be *redefined* in an association subclass. Intuitively, a role is redefined in a subclass when something is intentionally said about the set of classes playing the role in the subclass. The fact that role r is redefined in an association subclass a_i implies that the roleplays (c, r, a_t) of the association superclass a_t are *not* necessarily roleplays of a_i ; only roleplays intentionally defined for a_i apply to it.

Semantics The notion of role redefinition is captured by the predicate $redefinedRole(r, a)$, with the semantics that role r of association a is redefined:

$$\begin{aligned} redefinedRole(R, A_{sub}) &\leftarrow playsRoleIn_{d,D}(C, R, A_{sub}), \\ subclassOf_{d,D}(A_{sub}, A_{super}) & \end{aligned}$$

A class c playing role r in association a_{super} plays the role in a subclass a_{sub} of a_{super} only if r has not been redefined in a_{sub} :

$$\begin{aligned} playsRoleIn(C, R, A_{sub}) &\leftarrow \\ not\ redefinedRole(R, A_{sub}), & \\ playsRoleIn_{-p}(C, R, A_{super}), & \\ subclassOf_{d,D}(A_{sub}, A_{super}) &\quad \square \end{aligned} \quad (21)$$

Intuitively, an instance of a subclass must be in all respects a valid instance of all its superclasses. For associations, matching this intuition requires some extra care. Towards this end, we define the notion of an association instance being valid with respect to its type:

Definition Association instance a_i is valid with respect to its type a_t , if every role r of a_i is valid with respect to a_t . Role r of a_i is valid with respect to a_t , if r is played by class c_i such that c_i is an instance of c_t such that c_t plays r in a_t .

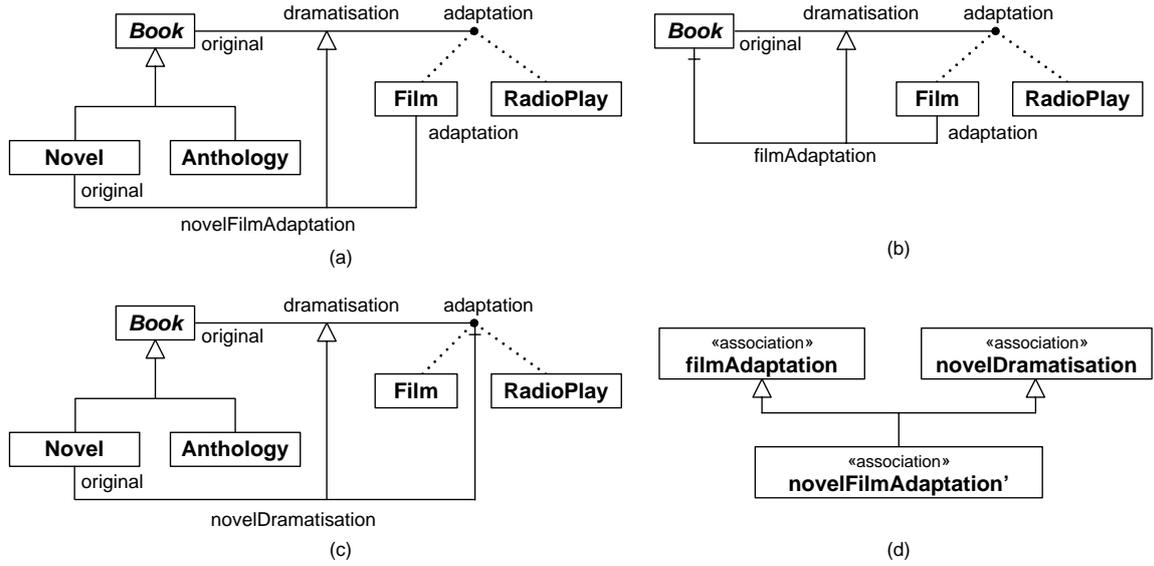


Fig. 9 Association generalisations. (a) Both roles redefined. (b) Only the role *adaptation* redefined. (c) Only the role *original* redefined. (d) An alternative definition of *novelFilmAdaptation*. An orthogonal line next to a role signals that the role is *not* redefined.

Role redefinitions may lead to situations in which an instance of an association is valid with respect to its direct type but invalid with respect to one or more of its non-direct types. An obvious way in which this may occur is when a role is redefined to be played by a class that is not a subclass of any of the classes playing the role in the association superclass.

However, there is another way in which this may occur, illustrated in Figures 9 (b) through (d). Parts (b) and (c) show how the associations *filmAdaptation* and *novelDramatisation*, respectively, both redefine one of the two roles, *adaptation* and *original*, of *dramatisation*, respectively. The instances of *novelFilmAdaptation'* defined in Figure 9 (d) should intuitively be valid instances of both *filmAdaptation* and *novelDramatisation*. However, neither role is redefined in *novelFilmAdaptation'* and therefore, without the requirement that an association must be valid with respect its both direct and non-direct types, *novelFilmAdaptation'* would unintendedly be equivalent to *dramatisation*.

Semantics We formalise the notion of a role being valid with respect to a type using the quaternary predicate $validRolePlay(c, r, a_i, a_t)$, with the semantics that the roleplay (c, r, a_i) is valid with respect to type a_t of a_i :

$$validRolePlay(C_i, R, A_i, A_t) \leftarrow \\ playsRoleIn_{-p}(C_i, R, A_i), playsRoleIn_{-p}(C_t, R, A_t), \\ instanceOf_{-p}(C_i, C_t), instanceOf_{-p}(A_i, A_t)$$

An association must be valid with respect to all its types:

$$\leftarrow not\ validRolePlay(C, R, A_i, A_t), \\ playsRoleIn_{-p}(C, R, A_i), \\ instanceOf_{-p}(A_i, A_t), not\ topLevel(C)$$

4.9 Cardinality constraints

An association may be subject to *cardinality constraints*. A cardinality constraint is said to *apply* to the instances of the association. Intuitively, a cardinality constraint restricts the possible combinations of classes participating in the instances. A cardinality constraint relates to one or more roles of the association. In addition, a cardinality constraint includes a *cardinality* and a *potency*. Cardinality is structurally equivalent to an attribute cardinality. Potency is the order of instances to which the constraint applies.

Example The formalisation of cardinality constraints is straightforward but lengthy and will not be discussed here for space reasons. However, the formalisation is demonstrated by showing how the cardinality constraint in the running example can be formalised.

The cardinality constraint applies to second-order instances of the *adapts* association. The constrained role is *adaptation*; notice, however, that the UML convention is followed and the constraint is located next to the *original* role. The intuitive semantics is that a second-order instance of *ProductType*, e.g., *IvanhoeTheFilm* or *IvanhoeTheRadioPlay*, may adapt at most one original.

To formalise cardinality constraints applying to an association, a predicate representing the instances of the association as tuples must be defined. Hence, we define the predicate $adapts(c, o)$, with the semantics that (adaptation) c adapts (original) o , both second-order instances of *ProductType*:

$$adapts(C, O) \leftarrow instanceOf_{tp}(A, adapts, 2), \\ playsRoleIn_{-p}(C, adaptation, A), \\ playsRoleIn_{-p}(O, original, A)$$

Above, the auxiliary predicate $instanceOf_{tp}(i, t, o)$ gives that i is a possible instance of t of order o ; see Appendix A, Rule 24 for the definition.

The cardinality constraint can now be written as:

$$\begin{aligned} \leftarrow & 2 \{ adapts(C_i, O) : adapts_p(C_i, O) \}, \\ & instanceOf_p(A, adapts), \\ & playsRoleIn_{-p}(C_t, adaptation, A), \\ & instanceOf_{-p}(C_i, C_t) \end{aligned}$$

4.10 Computational problems

In this section, we discuss how different computational problems related to NIVEL can be defined by introducing WCRL rules relating the possible and actual elements, or in other words, input and valid models. We will consider the NIVEL CHECK problem in detail and outline how problems involving search can be represented. A complete search problem is specified in Section 5.2 as a part of a case study in feature modelling.

In general, a computational problem specifies which entities declared in an input model *must* or *may* appear in valid models. We will use the generic term *check* to refer to the former alternative and *search* to the latter. Additional rules are needed to check that possible elements are actualised: such rules serve to restrict the search space spun by the input model.

The relationship between declared and actual elements may be specified for all instances of a language element at once or for a specific subset.

The NIVEL CHECK problem is defined as follows: Given an input model m , check if m is valid. Intuitively, for an input model to be valid under NIVEL CHECK, all the declared entities must be actualised in a valid model.

Semantics The NIVEL CHECK is formalised as follows. All the declared instantiations must be actual, direct instantiations:

$$\leftarrow instanceOf_D(I, T), not instanceOf_d(I, T) \quad (22)$$

all the declared generalisations must be actual, direct generalisations:

$$\begin{aligned} \leftarrow & subclassOf_D(C_{sub}, C_{super}), \\ & not subclassOf_d(C_{sub}, C_{super}) \end{aligned}$$

all the generalisations declared to be in generalisation sets must actually be in them:

$$\begin{aligned} \leftarrow & inGSet_D(C_{sub}, C_{super}, G), \\ & not inGSet(C_{sub}, C_{super}, G) \end{aligned}$$

all the declared attributes must correspond to actual attributes:

$$\begin{aligned} \leftarrow & hasAttr_D(C, N, P, D, L, U), \\ & not hasAttr(C, N, P, D, L, U) \end{aligned}$$

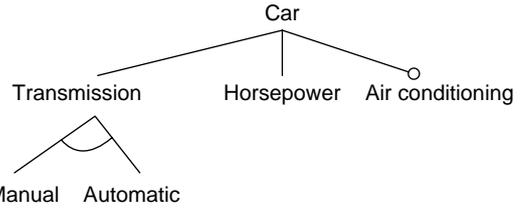


Fig. 10 A sample FODA model. Adapted from [21]

all the declared values must correspond to actual values:

$$\leftarrow hasValue_D(C, N, V), not hasValue(C, N, V)$$

all the roles declared to be played by classes must actually be played by them:

$$\leftarrow playsRoleIn_D(C, R, A), not playsRoleIn(C, R, A)$$

Note the role Rules 6, 7, 14, 17 and 20 play in setting up the search space: the rules are choice rules that allow but do not force entities to appear in valid models. Should the braces in the rule heads be dropped, the axiomatic part would in essence encode NIVEL CHECK.

5 Case study in feature modelling

In this section, we demonstrate the modelling capabilities of NIVEL by applying it to *feature modelling*, a modelling initiative gaining increasing popularity in the software product family domain. The section begins with a brief introduction to feature modelling. Thereafter, we show how NIVEL can be used to represent feature modelling concepts using three levels.

5.1 Feature modelling

This section provides an overview of feature modelling. We begin by discussing Feature Oriented Domain Analysis (FODA) [21], the first feature modelling language reported, and thereafter discuss a number of extensions.

In FODA, feature is defined as an attribute that directly affects end users. Features are organised into *feature models*. Figure 10 illustrates a sample feature model in FODA. A feature model is a tree where the root (in Figure 10, *Car*) is a feature, sometimes referred to as the *concept*. The root feature may have other features as its *subfeatures* which may in turn have other features as their subfeatures etc. A feature model is a description of a system family, e.g., a software product family.

There are a number of subfeature kinds: a *mandatory* subfeature (*Horsepower*) must be selected, i.e., included in feature configuration, whenever its parent is selected; an *optional* feature (*Air conditioning*) may be selected whenever its parent is selected, but needs not be selected; an *alternative* subfeature (*Transmission*) consists of a set of features (*Manual*, *Automatic*) of which exactly one must be selected whenever the parent feature is selected.

An individual product of the system family is described by a set of features that obeys the rules of the feature model. Such a set is termed a *feature configuration*, or simply *configuration*. The task of finding a valid configuration matching a specific set of requirements is termed *configuration task*.

A feature (other than the root feature) in a valid configuration must be *justified* by a subfeature in the model. That is, no other features than those required and enabled by the subfeatures in the model may be included in a configuration.

A number of feature modelling languages extending FODA have been suggested, see, e.g., [12, 13]. Most popular extensions include feature cardinalities, attributes and or-features, or more generally, group cardinalities.

A *feature cardinality* [13] specifies how many times a subfeature must be selected into a configuration. A mandatory feature corresponds to cardinality 1, and an optional subfeature to cardinality 0..1. Other cardinalities can be defined, such as 2..4 (two to four subfeatures must be selected) and 1..* (at least one subfeature must be selected). In some approaches, features may have *attributes* [13].

An *or-feature* [12] is a subfeature kind similar to an alternative feature, with the difference that at least one of the alternatives must be selected.

5.2 Using NIVEL to define a feature modelling language

In this subsection, we show how NIVEL can be used to define a feature modelling language. The discussion is based on a feature modelling language called Forfamel [1], although some aspects of the language have been left out to simplify the discussion.

In Forfamel, a distinction is made between *feature types* occurring in feature models, and their instances, termed *features*, appearing in (feature) configurations. Similarly for the notion of subfeature, Forfamel distinguishes between *subfeature definitions* involving feature types, and the subfeature relation between features. Forfamel supports four cardinality values for subfeature definitions: *mandatory*, *optional*, *or*, and *any*; each possible subfeature type may be instantiated at most once. Feature types may have attributes that are reflected as values in their instances. A single supertype may be defined for a feature type.

5.2.1 Forfamel metamodel in NIVEL Figure 11 illustrates the Forfamel metamodel in NIVEL. As a point of reference, Figure 12 contains a corresponding metamodel represented as a UML class diagram.

The metamodel contains two classes, *FeatureModel* and *FeatureType*, and the association *subfeature*, all of potency 2. An instance of *FeatureModel* represents a feature model and a second-order instance a feature configuration. Similarly, the instances of *FeatureType* repre-

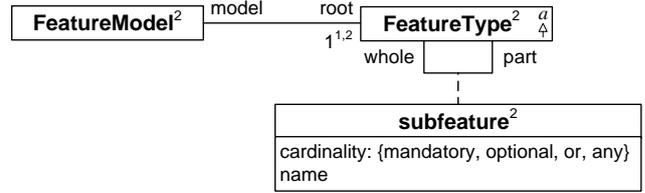


Fig. 11 A NIVEL metamodel for Forfamel

sent feature types and the second-order instances features. Finally, the instances of *subfeature* represent the subfeature definitions involving feature types and the second-order instances the subfeature relations between features.

5.2.2 Additional rules The major part of the intended semantics of feature models is captured by the semantics of NIVEL. However, a small number of additional rules are needed. These will be discussed in the following.

First, it is required for simplicity that there is at most one feature model in each NIVEL model:

$$\leftarrow 2 \{ \text{instanceOf}_d(M, \text{featureModel}) : \text{instanceOf}_{pd}(M, \text{featureModel}) \}$$

The *subfeature* association must be given additional semantics. Towards this end, we define the ternary predicate $\text{subf}^2(w, p, n)$ with the semantics that feature w has feature p as its subfeature under name n ; the superscript gives the order of instances of the association, in this case 2:

$$\begin{aligned} \text{subf}^2(W, P, N) \leftarrow & \text{playsRoleIn}_{-p}(W, \text{whole}, S_i), \\ & \text{playsRoleIn}_{-p}(P, \text{part}, S_i), \text{hasValue}_{-p}(S_t, \text{name}, N), \\ & \text{instanceOf}_{d,pd}(S_i, S_t), \text{instanceOf}_{pd}(S_t, \text{subfeature}) \end{aligned}$$

Note that the name N is a value of the subfeature definition S_t relating feature types, although the related elements are features, located on the level next below. Hence, the rule demonstrates that values and other properties of types are visible to their instances in WCRL.

As an example of applying subf^2 , the cardinality value *optional* is given a semantics as follows:

$$\begin{aligned} \leftarrow 2 \{ \text{subf}^2(W, P, N) : \text{subf}_p^2(W, P, N) \}, \\ \text{hasValue}_{-p}(S_t, \text{cardinality}, \text{optional}), \\ \text{hasValue}_{-p}(S_t, \text{name}, N), \\ \text{instanceOf}_{-p}(W_i, W_t), \text{playsRoleIn}_{-p}(W_t, \text{whole}, S_t) \end{aligned}$$

The notion of justification is represented by the predicate $\text{justified}(f)$ with the semantics that feature f is justified. A feature may be justified either by being a root feature

$$\begin{aligned} \text{justified}(F) \leftarrow & \text{root}^2(F, C), \text{root}_p^1(T, M), \\ & \text{instanceOf}_{pd}(F, T), \text{instanceOf}_{pd}(C, M) \end{aligned}$$

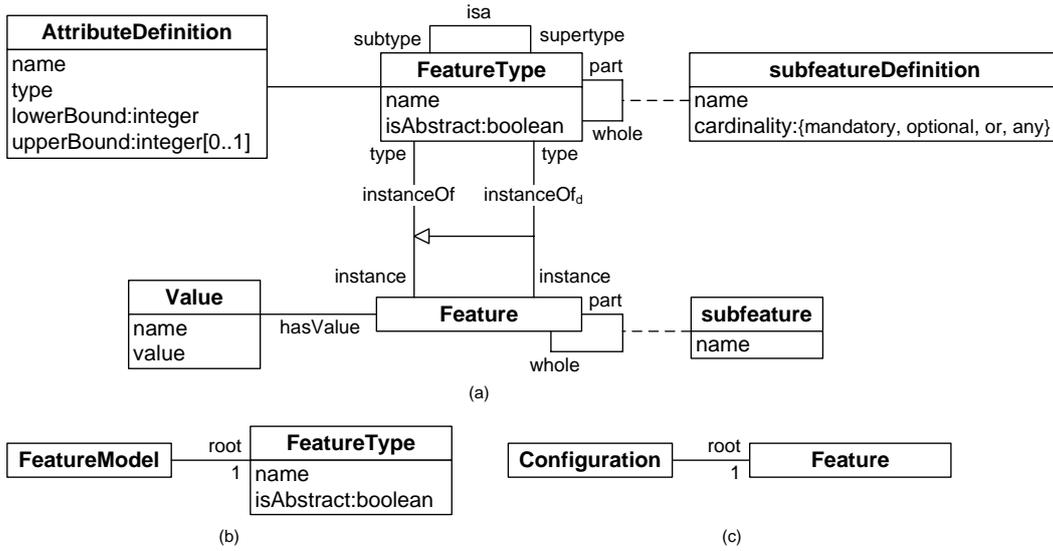


Fig. 12 The Forfamel metamodel as a UML class diagram. (a) Classes related to feature types and features, (b) feature model and (c) configuration. Adapted from [1]

where the predicate $root^2(f, c)$ is a tuple representation of the *rootFeature* association giving that feature f is a root feature in configuration c ; or by being a subfeature of another feature:

$$justified(F) \leftarrow subf_{-p}^2(W, F, N)$$

No feature without a justification may appear in a valid configuration:

$$\leftarrow instanceOf_{tp}(F, featureType, 2), class(F), \\ not\ justified(F)$$

For purposes to be explained in the next subsection, we restrict the allowed subfeature definitions in such a way that any valid configuration is guaranteed to be finite. One way for achieving this is to first consider the binary predicate $subf_p^1(w, p)$ with the semantics that a feature of type w possibly has a subfeature of type p :

$$subf_p^1(W, P_{sub}) \leftarrow playsRoleIn_p(W, whole, S), \\ playsRoleIn_p(P, part, S), \\ instanceOf_p(S, subfeature), subclassOf_p(P_{sub}, P)$$

The transitive closure $subf_{tp}^1$ of $subf_p^1$ is defined as:

$$subf_{tp}^1(W, P) \leftarrow subf_{tp}^1(W, P_0), \\ subf_p^1(P_0, P), instanceOf_p(W, featureType)$$

The desired effect can now be achieved using a simple integrity constraint:

$$\leftarrow subf_{tp}^1(C, C), instanceOf_p(C, featureType) \quad (23)$$

Intuitively, the rule forbids any feature type an instance of which could possibly have a transitive subfeature of the same type.

5.2.3 Configuration task The CONFIGURATION TASK was informally defined in Section 5.1 as: Given a feature model and a specific set of requirements for an individual product, find a valid configuration of the model. In this section, we will define the CONFIGURATION TASK for Forfamel defined in NIVEL.

A feature model consists of first-order instances of the top level classes and hence constitutes level 1 of the input model. That the feature model is taken as given implies that the problem is to check it. Hence, rules similar to those given Section 4.10 NIVEL CHECK applying exclusively to first-order instances of the top-level classes can be used. As an example of such rules,

$$\leftarrow instanceOf_{pd}(I, T), not\ instanceOf_d(I, T), \\ topLevel(T)$$

gives that the feature types, subfeature definitions and feature models are actualised as declared, cf. Rule 22.

On the other hand, the level 2 model elements constituting a valid configuration are subject to no additional constraints, with the exception that if a configuration is declared, it must be actualised:

$$\leftarrow instanceOf_{tp}(C, featureModel, 2), not\ class(C)$$

On the contrary, the attributes of feature types may take any value in their respective domains when instantiated in features:

$$hasValue_D(I, N, V) \leftarrow instanceOf_{pd}(I, T), \\ instanceOf_{pd}(T, featureType), \\ hasAttr_p(T, N, 1, D, L, U), contains(D, V)$$

Hence, finding a valid configuration can be characterised as a search problem. The search space \mathcal{S} , consisting of elements that may possibly appear in a configuration, must be explicitly given in the input model. \mathcal{S}

Input: feature type f_t to be instantiated
Output: set of rules representing f_i , an instance of f_t
 let f_i be a new object constant
 add fact $instanceOf_D(f_i, f_t) \leftarrow$
foreach s_t such that $s_t:subfeature \wedge f_t \in s_t \rightarrow whole$ **do**
 let s_i be a new object constant
 add fact $instanceOf_D(s_i, s_t) \leftarrow$
 add fact $playsRoleIn(f_i, whole, s_i) \leftarrow$
foreach p_t such that $p_t:featureType \wedge p_t \in s_t \rightarrow part$ **do**
foreach t such that $subtypeOf(t, p_t)$ **do**
if $t.isAbstract$ **then** continue
 $i := instantiate(t)$
 add fact $playsRoleIn_D(i, part, s_i) \leftarrow$
return f_i

Fig. 13 Algorithm $instantiate(t)$ for generating the search space \mathcal{S} for a feature model, initially called with the root feature type as the input parameter. For simplicity, same symbols are used for model elements and the object constants representing them. Shorthands used: $i:t \leftrightarrow instanceOf(i, t)$, $a \rightarrow r = \{c : playsRoleIn(c, r, a)\}$.

must include a number large enough of features and subfeature relationships to cover any valid configuration. A finite \mathcal{S} meeting this requirement can be generated using the algorithm given as Figure 13, provided that Rule 23 is satisfied.

5.2.4 Example feature model Figure 14 (a) illustrates a sample feature model describing a family of text editors. A text editor can manage different file formats (doc, rtf, pdf), has a clipboard of capacity between 1 to 10 items and at least one equation editor (*primary*) and optionally a secondary one, and supports spell checking in one to three languages in the set English, French and Dutch. A valid configuration of the model is illustrated in Figure 14 (b).

The sample model can be represented in WCRL in a similar manner as was shown for the running example of Figure 6 in Figure 7. The only part of the model requiring special care is the constraint that the primary and secondary equation editors must not be of the same type; this can be represented in WCRL as follows:

$$\leftarrow subf_{-p}^2(W, E_1, primary), subf_{-p}^2(W, E_2, secondary), \\ instanceOf_{pd}(E_1, T), instanceOf_{pd}(E_2, T)$$

5.3 Discussion

The number of model elements needed to represent the Forfamél concepts is significantly reduced in the NIVEL metamodel, shown in Figure 11, compared with the UML variant shown in Figures 12 (a) through (c): the classes *Feature*, *Configuration* and *subfeature* need not be explicitly represented in the NIVEL variant, as they are represented by the second-order instances of *FeatureType*, *FeatureModel* and *subfeature*, respectively.

In addition, the instantiation between feature types and features, and abstractness and attributes of and subclassing relation between feature types could be represented the language elements of NIVEL. In the UML variant these must be explicitly specified through associations *instanceOf*, *instanceOf_d*, and *isa*, and classes *Attribute* and *AttributeDefinition*. In conclusion, NIVEL helped in reducing the *accidental complexity* in the Forfamél metamodel in the sense defined and envisaged by Atkinson and Kühne [7].

In addition to abstract syntax, the NIVEL metamodel for Forfamél captures the major part of the semantics of Forfamél; a number of additional rules were required to complete the semantics. However, the number and complexity of the additional rules is modest compared with what would have been required to define a semantics for Forfamél from scratch. The same applies to the UML metamodel for Forfamél: the metamodel specifies a syntax for feature models and configurations but says nothing about their interrelations. Hence, constraints supplementing the UML metamodel, e.g., in OCL (Object Constraint Language) [30], would have to cover the entire Forfamél semantics.

The implementation of the case study in WCRL and results from test runs are discussed in Section 6.1

6 Discussion and comparison to previous work

In this section, we first discuss the validation currently existing for NIVEL. Thereafter, we motivate and discuss the choice of WCRL as the knowledge representation language for NIVEL. Section 6.3 shows how ideas discussed in Section 2 are reflected in NIVEL, followed by a comparison between NIVEL and previous metamodeling frameworks. Finally, individual language elements of NIVEL are discussed in Section 6.5.

6.1 Validation

In this section, we discuss how NIVEL has been validated. The purpose of the validation is *not* to evaluate the individual modelling concepts underlying NIVEL: as mentioned in the introduction, the concepts as such do not represent a novel contribution of NIVEL. The core concepts of NIVEL are all established in the modelling community and the more recent ideas incorporated in NIVEL have been argued for in length [3, 4, 5, 6, 7, 35].

Instead, the primary purpose of the validation is to verify that the translation t from NIVEL to WCRL is faithful in the sense that there is a bijective mapping between the stable models of $t(M)$ and the valid models of M for each well-formed M . In other words, we wish to ascertain that the formalisation matches the notion of valid model and overall intuition of the language elements as described in Section 4.

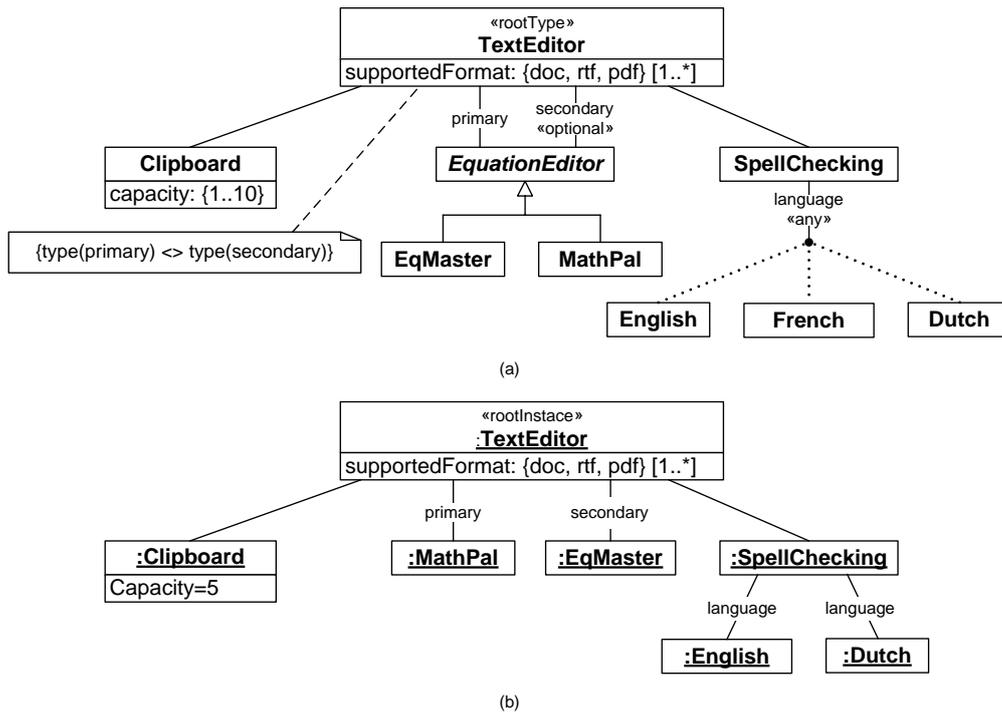


Fig. 14 (a) A sample feature model. (b) A sample feature configuration that is valid with respect to the above model.

A secondary goal of the validation is to check whether the stable models can be found efficiently enough for practical purposes, e.g., in an interactive modelling tool, to check whether a model is valid after each change.

We have implemented the axiomatic rules presented in Section 4, including those defining the NIVEL CHECK problem in Section 4.10, using the WCRL syntax accepted by the *smodels* system; see Section 3.4 and the references there for details on *smodels*. Similarly, we have implemented the feature modelling case study, including both the rules discussed in Section 5.2.2 and 5.2.3 and the example model illustrated in Figure 14 (a). Finally, we have created a number of other test models, including the running example of Figure 6.

It is easy to check by hand that the feature model in Figure 14 (a) has 1,960 valid configurations. This is also the number of stable models found by the *smodels* system given the corresponding WCRL program as input; a number of configurations were checked manually to verify that they are intuitively valid. Finding the first stable model took approximately 0.2 seconds (average over 10 runs) and all stable models were found in approximately 0.5 seconds on a Pentium D CPU 3.00 GHz desktop computer.

Other test models gave similar results: the results from test runs were intuitively correct. The processing time needed has been very acceptable: e.g., verifying that the running example represents a valid model requires approximately 0.02 seconds.

To conclude, the results from the test runs support the hypothesis that the formalisation captures the in-

tended semantics of NIVEL and that the implementation in WCRL could be used in tools supporting NIVEL or metamodels expressed in it. In particular, the semantics given for Forfamel, predominantly based on the semantics of NIVEL, seems to provide a basis for a sound and efficient tool supporting the CONFIGURATION TASK similar to what has been described in, e.g., [2].

The files related to the current implementation and instructions for downloading and running the *smodels* system can be found on a dedicated webpage².

6.2 Choice of knowledge representation language

In this section, we motivate the choice of WCRL as a knowledge representation language for NIVEL and contrast it with other formal languages.

WCRL possesses a number of characteristics that make it appropriate for representing knowledge in general and NIVEL in particular. A declarative formal semantics has been given for WCRL. Although the language includes a form of variables, predicates and function symbols, it is still decidable and of reasonable computational complexity. An efficient implementation for the language, *smodels*, is available under the GNU Public Licence.

For the case of NIVEL in particular, cardinality constraints allow the compact representation of a number of NIVEL language elements, most noteworthy cardinality bounds in attributes and cardinality constraints. Further, the mapping given in Section 4 is *modular* in the

² See <http://www.soberit.tkk.fi/nivel/>

sense that each model element can be translated independently of other elements. The translated program is *polynomial in size* with respect to the input model. We also claim that the translation is relatively *easy to understand*: most model elements are mapped into simple facts, the most notable exception being cardinality constraints that result in slightly more complex but standard rules; the axiomatic rules and rules encoding computational problems are likewise relatively simple. Thus, the translation can also be easily automated. Finally, as discussed in the previous subsection, the translated model is *detailed enough to enable automated reasoning*.

The *groundedness principle* is a semantic property of WCRL instrumental in achieving the above-discussed desirable characteristics. As can be recalled from Section 3.2, each atom in a stable model must be grounded. The principle enables deducing, e.g., that $class(c)$ can be in a stable model based on Rules 1 and 8 only. The groundedness principle helps to avoid so-called frame axioms and thus leads to a more compact formulation of NIVEL than could otherwise have been achieved.

Besides WCRL, other knowledge representation languages include a form of groundedness principle. An important class are typical logic programming systems, such as Prolog implementations. However, such systems are different from WCRL in a fundamental respect: their task is to compute a yes/no answer, or more generally, an answer substitution, for a given query. WCRL, on the other hand, is based on the *answer set programming* paradigm where no explicit query is needed. Instead, each stable model of a program represents a solution (in NIVEL, valid model) to the problem (an input model) encoded by the program. The answer set programming paradigm seems to suite NIVEL better than the query-oriented one: querying for a valid model would require the definition of a predicate capturing the notion of valid model, hence resulting in a larger number of axiomatic rules.

Another difference between WCRL and Prolog-like systems is that WCRL rules may include weight constraints whereas Prolog-like languages must be extended to allow forms of constraints in rules; such integrated systems are known under the *constraint logic programming* paradigm [20]. However, constraints in WCRL are an integral part of the language and its declarative semantics, whereas constraint logic programming is typically based on at least a partially procedural semantics. Further, weight constraint rules can themselves be understood as constraints on the solutions [29].

On the other hand, Berardi et al. [9] have given UML class diagrams with a limited form of constraints and ignoring implementation issues a “natural” formal semantics by translation to first-order predicate logic and subsequently to **EXPTIME**-decidable description logics. The encoding in first-order predicate logic serves as a point of reference in the sense that it is possible to argue

for the correctness of other formalisations by showing that it is equivalent to the first-order encoding.

In the case of NIVEL, following a similar approach would bring a number of benefits. A WCRL encoding would probably be readily understandable to a wider audience than a WCRL one. Similarly, a first-order encoding would likely provide an easier point of comparison for formalisations of NIVEL in alternative knowledge representation languages.

On the other hand, unlike in the case of first-order predicate logic and description logics, it is questionable whether an encoding in first-order predicate logic is more natural than one in WCRL: specifically, the lack of a groundedness property in first-order predicate logic implies that some form of frame axioms would be required to prevent elements without justification appearing in a valid model.

First-order predicate logic could be used as a basis for inferences. Unlike when using WCRL as described in Section 5.2.3 and shown in Figures 13, unrestricted first-order predicate logic does not require an a priori instantiation of model elements, thus resulting in a more simple translation. Of course, this would come at the cost of losing decidability. A third approach is to resort to decidable subsets of first-order predicate logic, e.g., the description logics \mathcal{DLR}_{ifd} and \mathcal{ALCQI} used by Berardi et al. [9]. However, such subsets imply syntactic restrictions that many, especially practitioners, may find both unintuitive and hard to accept.

6.3 Recent metamodelling ideas

In this section, we discuss how NIVEL relates to the notions of strict metamodelling, distinction between ontological and linguistic instantiation, unified modelling elements and deep instantiation.

NIVEL adheres to the strict metamodelling rule [4]: each class is an instance of a class on a level next above it and other relationships between classes (generalisation and association) may not cross levels. However, if multiple classification is adopted, a class may be a direct instance of more than one class. See Appendix B for a detailed discussion and proofs.

NIVEL commits to the distinction between linguistic and ontological instantiation. NIVEL itself is taken as a language. This implies that language elements illustrated in Figure 5, are not model elements and are thus not included in models, such as the one illustrated in Figure 6; in particular, language elements do not form the top level of NIVEL models. Also, we do not suggest that language elements are instances of themselves, at least not in the sense described by the *instanceOf* relation discussed in Section 4.

The relationship between model and language elements can be termed linguistic instantiation. As an example, in the sample model of Figure 6, *ProductType* is a

linguistic instance of *Class*. Conversely, the *instanceOf* relation in NIVEL models, can be termed ontological.

However, the rules given in Section 4 cannot be used to check whether in instantiation in a NIVEL model is ontological in the sense defined in [6, 22]: ontological instantiation is based on the meaning of model elements in terms of a system external to the model, i.e., its original. Hence, one cannot in general perform such a check based on the information available in the model alone.

The NIVEL concept of class, and association as a kind of class, is an implementation of the idea of unified modelling elements: a class has characteristics typical to both classes and objects in other modelling languages: a class can assume both the role of a type and an instance, have both attributes and values, and participate in associations of different potencies.

Finally, the notion of deep instantiation is embodied in NIVEL in classes, including associations, attributes and cardinality constraints. The key construct in implementing deep instantiation is potency.

6.4 Metamodeling languages and frameworks

In this section, we discuss how NIVEL compares to a number of previous metamodeling frameworks.

MOF (Meta Object Facility) promises to support any number of levels through its “ability to navigate from an instance to its metaobject, i.e., its classifier” [27]. However, a number of issues differentiate MOF from NIVEL. To begin with, the original purpose of MOF was to provide a standard way of accessing run-time meta-information about objects [5]. Consequently, MOF is committed to the object-orientated paradigm through the use of terms such as *navigation* and *reference*. Further, MOF makes reference to specific technologies such as XML.

Unlike NIVEL, MOF neither makes a distinction between linguistic and ontological instantiation nor commits to the notion of unified modelling elements. MOF includes no notion of deep instantiation and seems to make no guarantees about the strictness of models expressed in it. Further, MOF is allegedly self-defining and constitutes the top level of model hierarchies defined in it, such as the UML four-layer architecture [38].

Another example of a metamodeling framework that does not subscribe to the strict metamodeling rule is the LOOPS scheme for metalevels, known as the *Golden Braid*. In LOOPS *Object*, *Class* and *MetaClass* pairwise instantiate each other (in that order) and inherit from each other (in reverse order); in addition, *MetaClass* is an instance of itself [10]. Clearly, the concepts of instantiation and inheritance (generalisation) are incompatible in NIVEL and LOOPS.

UML [39] can be used to define a metamodel through a *profile* consisting of a number of *stereotypes*. Each stereotype defines an extension that can be applied to one or more metaclasses, such as *Class*. A stereotype

may define additional properties and constraints that become properties and constraints of the instances of a metaclass to which the stereotype is applied.

However, the profile mechanism provides only limited support for metamodeling. First, the number of levels in a profile-based approach is restricted to three. Further, stereotypes are technically extensions to UML language and not intended for representing domain (meta)-types [7]. The relation between stereotypes and classes is different from the one between classes and their instances, either objects on M0 or *InstanceSpecifications* on M1. Moreover, the profile mechanism has itself been criticised for ambiguity, e.g., with respect to whether only classes or all metaclasses may be extended, and problems in its semantics [19]. Also, from the pragmatic point of view, there has been confusion about whether stereotypes apply to classes, objects or both [8].

Telos [28] resembles NIVEL in that an individual, roughly corresponding to a class in NIVEL, may be an instance of one or more individuals. Also, individuals may have attributes and there may be generalisations between them in both Telos and NIVEL. However, there are important differences between the languages. Telos is based on extensive use of attributes, whereas NIVEL includes the association concept as a language primitive. Unlike NIVEL, Telos includes no notion of strictness. Finally, in strong contrast with NIVEL, Telos seems to intentionally unify model elements, language elements and the formal entities representing them.

6.5 NIVEL language elements

6.5.1 Generalisation An important question related to generalisation is whether multiple superclasses are allowed or not. For example, in the Java programming language a class may extend at most one class, whereas UML makes no similar restriction.

Given the intended versatility of NIVEL, the language is equipped with support for multiple superclasses as well as a means for restricting the number of direct superclasses to one, or disallowing them altogether.

It has been argued that covering and disjointness constraints attached to generalisation sets are in practice the most commonly used constraint in UML class diagrams [9], which serves as a motivation for including the concept of generalisation set in NIVEL.

6.5.2 Instantiation The ability to explicitly represent instantiation relationships enables NIVEL to be credited as a metamodeling language.

A key issue related to instantiation is the question whether multiple classification is allowed or not. Single instantiation is commonly assumed in object-oriented modelling and programming, although UML seems to equivocate on the issue. On the other hand, multiple

classification is a reasonable means for modelling concepts such as roles and phases. Again, recalling its intended versatility, NIVEL supports both multiple classification and enables a model to be restricted to single classification.

6.5.3 Attribute and value In previous work on deep instantiation, value has been considered as a special case of attribute, i.e., as an attribute of potency 0 [3]; in NIVEL, value is defined as a language element of its own. A number of arguments support this decision. First, attributes and values are syntactically different, as shown in Figure 5 and reflected in predicates *hasAttr* and *hasValue*. Second, the semantics of instantiating an attribute of potency greater than 1, as captured by Rule 15, is drastically different from that of instantiating an attribute of potency 1 into values, cf. Rule 18.

6.5.4 Association Unlike in most previous conceptual modelling languages, a role in NIVEL can be played by one or more classes that do not share a superclass; a similar approach has been suggested by Steimann [35]. Arguably, assuming multiple superclasses can be defined for a class, the same effect can be achieved by defining a common superclass for the classes playing a certain role in an association. However, this may lead to a undesirably large number of classes that are in a sense artificial.

We also elaborate on association generalisation, especially on role redefinitions in association subclasses. Although the notion of association redefinition is included in UML since version 2.0, no detailed semantics has been given: “The interaction of association specialization with association end redefinition and subsetting is not defined” [39, p. 57]. On the other hand, Costal et al. have considered association redefinition as a mechanism for imposing additional constraints on associations [11] but similarly as UML and unlike NIVEL, have not discussed whether and how the notion of redefinition is related to generalisations between associations.

6.5.5 Constraint NIVEL defines no constraint language of its own, with the exception of cardinality constraints. However, as demonstrated in the case study on feature modelling in Section 5, it is possible to include arbitrary WCRL rules in the translated model $t(M)$. WCRL enables quantification over finite domains and it is straightforward to represent the standard Boolean connectives using weight constraint rules [29].

However, adopting WCRL as the constraint language for NIVEL would introduce a number of problems. First, arbitrary rules can easily be used to manipulate semantics given to predicates in axiomatic rules; however, this can be avoided with simple syntactic checks. More importantly, the potential users of NIVEL cannot be assumed to be familiar with WCRL and thus to be able write constraints in it. This problem seems not to pertain

to WCRL only, but also to other viable constraint languages, such as OCL or first-order logic: based on our experience, engineers and other professionals involved in software development typically seem not to be well familiar with the idea of expressing constraints in any knowledge representation language.

The cardinality constraints of NIVEL resemble those in entity-relationship modelling [37] and to a lesser extent those in UML. In UML, cardinality constraints are look-across constraints expressed using the multiplicities of individual ends and consequently and, in the case of associations with more than two ends, often too weak and awkward from the designer point of view [9].

7 Conclusions and further work

We have presented NIVEL, a novel conceptual modelling language capable of expressing models spanning an arbitrary number of levels. NIVEL synthesises and elaborates on a number of ideas introduced in the metamodelling community. The modelling facilities of NIVEL are demonstrated using a case study in feature modelling and an example based on product hierarchies.

NIVEL is given a formal semantics by translation to WCRL, a general-purpose, decidable knowledge representation language syntactically similar to logic programs. This enables both automated and other forms of reasoning about NIVEL.

There are a number of possible ways to extend the work presented in this paper: conceptual extensions, concrete modelling languages and tool support, and case studies. We will briefly elaborate on each of these.

Although NIVEL covers the most important conceptual modelling concepts, there are a number of additional concepts that could be integrated in the language. A constraint language is needed to capture the potentially complex dependencies that may occur between model elements. As a first step, a semantics could be specified for different forms of cardinality constraints [37]. Also, the notion of deriving model elements using constraints could be studied [31]. The strict metamodelling rules could be weakened to enable more flexible modelling style. The current requirement that all roles of an association must be specified on the top level may be unnecessarily strong and could be weakened. Including a notion of time in NIVEL would enable representing knowledge and reasoning about temporal properties of models, including notions such as dynamic vs. static typing. Finally, class-valued attributes would likely be a useful alternative for binary associations.

For practical purposes, concrete modelling languages and supporting tools are required. There are many conceivable forms of tool support. An elementary form of support would be a concrete syntax for NIVEL and an automated translation from this syntax to WCRL. A graphical modelling tool could be more usable than a

textual language. Likewise, a programming interface for NIVEL could be implemented; such interfaces could also be generated for individual NIVEL models.

Finally, case studies are needed to demonstrate the practical utility of NIVEL.

A Additional rules

This appendix contains the axiomatic rules that were skipped in Section 4.

A.1 Generalisation

A declared subclass is a possible class:

$$\text{class}_p(C_{\text{sub}}) \leftarrow \text{subclassOf}_D(C_{\text{sub}}, C_{\text{super}})$$

Multiple allowed superclasses for a class implies that a single superclass may be defined:

$$\begin{aligned} &\text{superclassingSingle}(I) \leftarrow \\ &\text{superclassingMultiple}(I), \text{class}_p(C) \end{aligned}$$

A possible generalisation between top-level classes always becomes an actual one:

$$\begin{aligned} &\text{subclassOf}_d(C_{\text{sub}}, C_{\text{super}}) \leftarrow \\ &\text{subclassOf}_D(C_{\text{sub}}, C_{\text{super}}), \text{topLevel}(C_{\text{super}}) \end{aligned}$$

A root class, as represented by the topLevel_D predicate, may not be a subclass:

$$\leftarrow \text{topLevel}_D(C_{\text{sub}}), \text{subclassOf}_D(C_{\text{sub}}, C_{\text{super}})$$

A.2 Instantiation

A declared instance is a possible class, cf. Rule 8:

$$\text{class}_p(I) \leftarrow \text{instanceOf}_D(I, T)$$

A class declared to be on the top level and all its possible subclasses are on the top level:

$$\begin{aligned} &\text{topLevel}(C_{\text{sub}}) \leftarrow \text{topLevel}_D(C_{\text{super}}), \\ &\text{subclassOf}_p(C_{\text{sub}}, C_{\text{super}}) \end{aligned}$$

The extension of the $\text{level}(l)$ predicate is defined using a model-specific constant maxLevel :

$$\text{level}(0..\text{maxLevel}) \leftarrow$$

The construct used is a shorthand way of stating that $\text{level}(l)$ holds for all values l between 0 and maxLevel , both included.

The predicate $\text{instanceOf}_{tp}(i, t, o)$ gives that class i is a possible instance of class t of order o . The subscript t stands for “transitive”. The rule for order 1 is:

$$\text{instanceOf}_{tp}(I, T, 1) \leftarrow \text{instanceOf}_p(I, T) \quad (24)$$

and for higher orders:

$$\begin{aligned} &\text{instanceOf}_{tp}(I, T, O + 1) \leftarrow \text{instanceOf}_p(I, T_0), \\ &\text{instanceOf}_{tp}(T_0, T, O), \text{class}_p(T), \text{level}(O) \end{aligned}$$

A corresponding predicate, instanceOf_{tpd} , is defined for the case of direct instantiation. For order 1, we write:

$$\text{instanceOf}_{tpd}(I, T, 1) \leftarrow \text{instanceOf}_{pd}(I, T)$$

and for higher orders:

$$\begin{aligned} &\text{instanceOf}_{tpd}(I, T, O + 1) \leftarrow \text{instanceOf}_{pd}(I, T_0), \\ &\text{instanceOf}_{tpd}(T_0, T, O), \text{class}_p(T), \text{level}(O) \end{aligned}$$

A non-root top-level class is assigned a potency as follows:

$$\begin{aligned} &\text{hasPotency}(C_{\text{sub}}, P) \leftarrow \text{hasPotency}_D(C_{\text{super}}, P), \\ &\text{topLevel}_D(C_{\text{super}}), \text{subclassOf}_p(C_{\text{sub}}, C_{\text{super}}) \end{aligned}$$

Direct implied instantiations for possible predicates are represented as follows, cf. Rule 11:

$$\begin{aligned} &\text{instanceOf}_{pd}(I_{\text{sub}}, T) \leftarrow \text{instanceOf}_D(I_{\text{super}}, T), \\ &\text{subclassOf}_p(I_{\text{sub}}, I_{\text{super}}) \end{aligned}$$

A similar rule can be written for the non-direct case, cf. Rule 12:

$$\begin{aligned} &\text{instanceOf}_p(I_{\text{sub}}, T_{\text{super}}) \leftarrow \text{instanceOf}_{pd}(I_{\text{super}}, T_{\text{sub}}), \\ &\text{subclassOf}_p(T_{\text{sub}}, T_{\text{super}}), \text{subclassOf}_p(I_{\text{sub}}, I_{\text{super}}) \end{aligned}$$

The definition of $\text{superclassingMultiple}$ is similar to that of $\text{superclassingSingle}$, cf. Rule 13:

$$\begin{aligned} &\text{superclassingMultiple}(I) \leftarrow \\ &\text{instanceSuperclassingMultiple}(T), \\ &\text{instanceOf}_{d,pd}(I, T) \end{aligned}$$

A.3 Generalisation set

We define the predicate $\text{inGSet}(c_{\text{sub}}, g)$ as a projection from $\text{inGSet}(c_{\text{sub}}, c_{\text{super}}, g)$ with the semantics that c_{sub} is a subclass in generalisation set g . The rule for the possible case is:

$$\text{inGSet}_D(C_{\text{sub}}, G) \leftarrow \text{inGSet}_D(C_{\text{sub}}, C_{\text{super}}, G)$$

and for the actual one:

$$\text{inGSet}(C_{\text{sub}}, G) \leftarrow \text{inGSet}_{-,D}(C_{\text{sub}}, C_{\text{super}}, G)$$

The binary predicate $\text{gSetOf}(g, c_{\text{super}})$ with the semantics that g is a generalisation set of c_{super} is defined in a similar manner. For the possible case, we write:

$$\text{gSetOf}_D(G, C_{\text{super}}) \leftarrow \text{inGSet}_D(C_{\text{sub}}, C_{\text{super}}, G)$$

and for the actual one:

$$\text{gSetOf}(G, C_{\text{super}}) \leftarrow \text{inGSet}_{-,D}(C_{\text{sub}}, C_{\text{super}}, G)$$

The unary domain predicate $gSet(g)$ is defined with the semantics that g is a generalisation set:

$$gSet(G) \leftarrow gSetOf_D(C_{super}, G)$$

The requirement that all the generalisations in a generalisation set must share the same superclass is represented as an integrity constraint:

$$\leftarrow 2 \{gSetOf(G, C_{super}) : gSetOf_D(G, C_{super})\}, gSet(G)$$

A.4 Attribute

A declared attribute of a class is a possible attribute for its instances, cf. Rule 15:

$$\begin{aligned} hasAttr_p(I, N, P - O, D, L, U) \leftarrow \\ hasAttr_D(T, N, P, D, L, U), \\ instanceOf_{tp}(I, T, O), P > O \end{aligned}$$

and for its subclasses, cf. Rule 16:

$$\begin{aligned} hasAttr_p(C_{sub}, N, P, D, L, U) \leftarrow \\ hasAttr_D(C_{super}, N, P, D, L, U), \\ subclassOf_p(C_{sub}, C_{super}) \end{aligned}$$

On the top level, a declared attribute always turns into an actual one:

$$\begin{aligned} hasAttr(C, N, P, D, L, U) \leftarrow \\ hasAttr_D(C, N, P, D, L, U), topLevel(C) \end{aligned}$$

A.5 Value

A value declared for a superclass is a possible value of its subclasses:

$$\begin{aligned} hasValue_p(C_{sub}, N, V) \leftarrow hasValue_D(C_{super}, N, V), \\ subclassOf_p(C_{sub}, C_{super}) \end{aligned}$$

A.6 Association

A direct instance of an association is an association:

$$association(I) \leftarrow association(T), instanceOf_{d,pd}(I, T)$$

Similarly for a subclass of an association:

$$\begin{aligned} association(A_{sub}) \leftarrow association(A_{super}), \\ subclassOf_{-p}(A_{sub}, A_{super}) \end{aligned}$$

Roles can be declared for root classes only:

$$\leftarrow hasRole_D(A, R), not topLevel_D(A)$$

and only for associations:

$$\leftarrow hasRole_D(A, R), not association(A)$$

A role that is declared for an association turns into an actual role of the association:

$$hasRole(A, R) \leftarrow hasRole_D(A, R)$$

and into a possible role:

$$hasRole_p(A, R) \leftarrow hasRole_D(A, R)$$

Instances of an association of all orders have the same possible roles, cf. Rule 19:

$$hasRole_p(I, R) \leftarrow hasRole_D(T, R), instanceOf_{tp}(I, T, O)$$

A role declared for a superclass is a possible role of its subclasses, cf. Rule 21:

$$\begin{aligned} playsRoleIn_p(C, R, A_{sub}) \leftarrow playsRoleIn_D(C, R, A_{super}), \\ subclassOf_p(A_{sub}, A_{super}) \end{aligned}$$

A direct roleplay implies an ordinary, non-direct one:

$$playsRoleIn(C, R, A) \leftarrow playsRoleIn_{d,D}(C, R, A)$$

The top level is a special case: a declared roleplay becomes an actual one, given that the association exists and has the relevant role:

$$\begin{aligned} playsRoleIn(C, R, A) \leftarrow playsRoleIn_D(C, R, A), \\ hasRole(A, R), topLevel(C), topLevel(A) \end{aligned} \quad (25)$$

B Level-respectiveness and strictness in Nivel

In this appendix, we show that the *instanceOf* relation in NIVEL meets the requirements for a relation suitable of building metalevels [22]. The acyclic, anti-transitive and level-respecting properties are defined as:

$$\text{acyclic: } \forall e_1, e_2, n : e_1 R^n e_2 \rightarrow \neg e_2 R e_1$$

$$\text{anti-transitive: } \forall n \geq 2 : R^n \cap R = \emptyset$$

level-respecting:

$$\forall n, m : (\exists e_1, e_2 : e_1 R^n e_2 \wedge e_1 R^m e_2) \rightarrow n = m$$

To see that *instanceOf* is level-respecting, note that level-respectiveness amounts to saying that each instantiation path from e_1 to e_2 is of the same length. In NIVEL, the length of the path equals the difference in the level numbers of the respective elements; each class is guaranteed to be on a single level by Rules 9 and 10.

Level-respectiveness implies anti-transitivity [22]. In addition, level-respectiveness implies acyclicity: acyclicity can alternatively be formulated as excluding paths from elements to themselves:

$$(\forall e_1, e_2, n : e_1 R^n e_2 \rightarrow \neg e_2 R e_1) \leftrightarrow \forall e, n : \neg e R^n e$$

Assume that R is not acyclic, i.e., $\exists e, n : e R^n e$ for $e = e_0$ and $n = m$. By the definition of R^n , it is also the case that $e_0 R^{2m} e_0$. But this is a contradiction with the definition of level-respecting, with $m \neq 2m$. Hence, a

relation that is not acyclic cannot be level-respecting, or in other words, a level-respecting relation is necessarily acyclic. We conclude that the *instanceOf* relation is a relation able to build metalevels.

To show strictness, we split the strict metamodeling rule [4] into three requirements:

1. Every element of an M_m -level model must be an instance-of *at least one* element of an M_{m+1} -level model, for all $m < n - 1$.
2. Every element of an M_m -level model must be an instance-of *at most one* element of an M_{m+1} -level model, for all $m < n - 1$.
3. Any relationship other than the instance-of relationship between two elements X and Y implies that $level(X) = level(Y)$.

The first requirement is reflected in NIVEL for classes, including associations, as follows. Class c can be in a valid model, represented by the predicate $class(c)$, for two reasons. First, based on Rule 8, a class can be in a model by virtue of being a direct instance of another class. Rule 9 guarantees that an instance is on the next level below its type. Second, according to Rule 1, a class can be in a model by being on the top level. This corresponds to the exception made for the M_{n-1} -level. Hence, NIVEL adheres to the first requirement.

NIVEL is *not* strict in the sense of the second requirement: a class is allowed to be an instance of multiple elements. However, single instantiation can be enforced by setting the *singleClassification* attribute for the model.

NIVEL conforms to the third requirement, with association and generalisation as the relationships other than instance-of. For associations, Rule 25 guarantees that this is the case for the top level: only top-level classes may play roles on the top level. For other levels, this is guaranteed by Rule 20: each class playing a role in an association must be an instance of a class playing the same role in the type of the association. Recalling that Rule 9 guarantees that each instantiation results in the level of the class being incremented by one, we have the desired result.

To show the same for generalisations, let us assume generalisation (A, B) where A is on a level above B ; this is illustrated in Figure 15 (a). Class B is on level $m + 1$ by virtue of being a direct instance of C on level $m + 2$ and A on level m by being an instance of D on level $m + 1$. Rule 12 implies that A is an instance of C . Hence, A is an instance of both C and D and must be on both levels m and $m + 1$. This is in conflict with Rule 10, a contradiction.

The case where a subclass is on a level below its superclass is shown in Figure 15 (b). The instantiation (B, D) is implied by Rule 12, resulting in a contradiction.

References

1. Asikainen, T., Männistö, T., Soininen, T.: A unified conceptual foundation for feature modelling. In: L. O'Brien

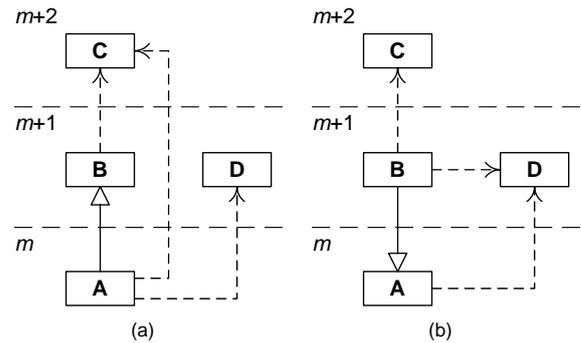


Fig. 15 Generalisations between classes on different levels are not possible. (a) Subclass on a level below its superclass. (b) Subclass on a level above its superclass.

- (ed.) Proceedings of the 10th International Software Product Line Conference (SPLC 2006), pp. 31–40 (2006)
2. Asikainen, T., Männistö, T., Soininen, T.: Kumbang: A domain ontology for modelling variability in software product families. *Advanced Engineering Informatics* **21**(1), 23–40 (2007)
3. Atkinson, C., Kühne, T.: The essence of multilevel meta-modeling. In: M. Gogolla, C. Kobryn (eds.) Proceedings of The Fourth International Conference on the Unified Modeling Language UML 2001, *Lecture Notes in Computer Science*, vol. 2185, pp. 19–33 (2001)
4. Atkinson, C., Kühne, T.: Profiles in a strict metamodeling framework. *Science of Computer Programming* **44**(1), 5–22 (2002)
5. Atkinson, C., Kühne, T.: Rearchitcting the UML infrastructure. *ACM Transactions on Modeling and Computer Simulation* **22**(4), 290–321 (2002)
6. Atkinson, C., Kühne, T.: Model-driven development: A metamodeling foundation. *IEEE Software* **20**(5), 36–41 (2003)
7. Atkinson, C., Kühne, T.: Reducing accidental complexity in domain models. *Software and Systems Modeling* (2007). DOI 10.1007/s10270-007-0061-0
8. Atkinson, C., Kühne, T., Henderson-Sellers, B.: Stereotypical encounters of the third kind. In: J.M. Jézéquel, H. Hussmann, S. Cook (eds.) Proceedings of the 5th International Conference on the Unified Modeling Language (UML 2002), *Lecture Notes in Computer Science*, vol. 2460 (2002)
9. Berardi, D., Calvanese, D., Giacomo, G.D.: Reasoning on UML class diagrams. *Artificial Intelligence* **168**(1-2), 70–118 (2005)
10. Bobrow, D.G., Stefik, M.: The LOOPS manual. Tech. rep., Xerox Corporation (1983)
11. Costal, D., Gómez, C.: On the use of association redefinition in UML class diagrams. In: D.W. Embley, A. Olivé, S. Ram (eds.) 25th International Conference on Conceptual Modeling (ER2006), *Lecture Notes in Computer Science*, vol. 4215, pp. 513–527 (2006)
12. Czarnecki, K., Eisenecker, U.W.: Generative Programming. Addison-Wesley, Boston (MA) (2000)
13. Czarnecki, K., Helsen, S., Eisenecker, U.W.: Staged configuration through specialization and multilevel configuration of feature models. *Software Process: Improvement and Practices* **10**(2), 143–169 (2005)

14. Evans, A., France, R.B., Lano, K., Rumpe, B.: The UML as a formal modeling notation. In: Selected papers from the First International Workshop on The Unified Modeling Language UML'98, *Lecture Notes in Computer Science*, vol. 1618, pp. 336–348 (1999)
15. van Gigch, J.P.: System design, modeling and metamodeling. Plenum Press (1991)
16. Hall, A.: Seven myths of formal methods. *IEEE Software* **7**(5), 11–19 (1990)
17. Harel, D., Rumpe, B.: Meaningful modeling: What's the semantics of "semantics"? *IEEE Computer* **37**(10), 64–72 (2004)
18. Henderson-Sellers, B., Bulthuis, A.: COMMA: Sample metamodels. *JOOP* **9**(7), 44–48 (1996)
19. Henderson-Sellers, B., Gonzalez-Perez, C.: Uses and abuses of the stereotype mechanism in UML 1.x and 2.0. In: 9th International Conference on Model Driven Engineering Languages and Systems (MoDELS 2006), pp. 16–26
20. Jaffar, J., Lassez, J.L.: Constraint logic programming. In: POPL '87: Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, pp. 111–119. ACM Press (1987)
21. Kang, K.C., Cohen, S.G., Hess, J.A., Novak, W.E., Peterson, S.A.: Feature-oriented domain analysis (FODA)—feasibility study. Tech. Rep. CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie Mellon University (1990)
22. Kühne, T.: Matters of (meta-) modeling. *Software and Systems Modeling* **5**(4), 369–385 (2006)
23. Kühne, T., Steimann, F.: Tiefe Charakterisierung. In: Modellierung 2004, pp. 121–133 (2004). In German.
24. Ludewig, J.: Models in software engineering—an introduction. *Software and Systems Modeling* **2**(1), 5–14 (2003)
25. McUmber, W.E., Cheng, B.H.C.: A general framework for formalizing UML with formal languages. In: Proceedings of the 23rd International Conference on Software Engineering (ICSE 2001), pp. 433–442 (2001)
26. Meyer, B.: Introduction to the Theory of Programming Languages. Prentice Hall, New York (1990)
27. Meta Object Facility (MOF) core specification, OMG available specification, version 2.0. Tech. Rep. formal/06-01-01 (2006)
28. Mylopoulos, J., Borgida, A., Jarke, M., Koubarakis, M.: Telos: Representing knowledge about information systems. *ACM Transactions on Information Systems* **8**(4), 325–362 (1990)
29. Niemelä, I.: Logic programs with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence* **25**(3-4), 241–273 (1999)
30. Object constraint language, OMG available specification, version 2.0. Tech. Rep. formal/06-05-01, Object Management Group (2006)
31. Olivé, A.: Derivation rules in object-oriented conceptual modeling languages. In: J. Eder, M. Missikoff (eds.) Proceedings of the 15th International Conference on Advanced Information Systems Engineering, *Lecture Notes in Computer Science*, vol. 2681, pp. 404–420 (2003)
32. Simons, P., Niemelä, I., Soinen, T.: Extending and implementing the stable model semantics. *Artificial Intelligence* **138**(1-2), 181–234 (2002)
33. Soinen, T.: An approach to knowledge representation and reasoning for product configuration tasks. Ph.D. thesis, Helsinki University of Technology (2000)
34. Stachowiak, H.: Allgemeine Modelltheorie. Springer (1973)
35. Steimann, F.: Role = interface: A merger of concepts. *Journal of Object-Oriented Programming* **14**(4), 23–32 (2001)
36. Syrjänen, T.: Omega-restricted logic programs. In: Proceedings of the 6th International Conference on Logic Programming and Nonmonotonic Reasoning, *Lecture Notes in Artificial Intelligence*, vol. 2713, pp. 267–280 (2001)
37. Thalheim, B.: Fundamentals of cardinality constraints. In: G. Pernul, A.M. Tjoa (eds.) Proceedings of the 11th International Conference on the Entity-Relationship Approach (ER'92), *Lecture Notes in Computer Science*, vol. 645, pp. 7–23 (1992)
38. Unified Modeling Language: Infrastructure, version 2.1.1. Tech. Rep. formal/2007-02-06, Object Management Group (OMG) (2007)
39. Unified Modeling Language: Superstructure, version 2.1.1. Tech. Rep. formal/2007-02-05, Object Management Group (OMG) (2007)