

Deep Meta-Modelling with METADEPTH

Juan de Lara¹ and Esther Guerra²

¹ Universidad Autónoma de Madrid (Spain), Juan.deLara@uam.es

² Universidad Carlos III de Madrid (Spain), eguerra@inf.uc3m.es

Abstract. Meta-modelling is at the core of Model-Driven Engineering, where it is used for language engineering and domain modelling. The OMG's Meta-Object Facility is the standard framework for building and instantiating meta-models. However, in the last few years, several researchers have identified limitations and rigidities in such scheme, most notably concerning the consideration of only two meta-modelling levels *at the same time*.

In this paper we present METADEPTH, a novel framework that supports a dual linguistic/ontological instantiation and permits building systems with an arbitrary number of meta-levels through deep meta-modelling. The framework implements advanced modelling concepts allowing the specification and evaluation of derived attributes and constraints across multiple meta-levels, linguistic extensions of ontological instance models, transactions, and hosting different constraint and action languages.

1 Introduction

Model-Driven Engineering (MDE) is a software development paradigm aiming at speeding up development times, while increasing quality and maintainability. MDE pursues these goals by treating models as the key assets of the process, being no longer mere documentation but used actively to (re-)generate code, as well as for validation and verification. Therefore, these activities demand computer-processable models with precise syntax. In MDE, models' syntax is defined through a meta-model that describes the set of valid models. Hence, meta-modelling is one of the pillars of MDE, being used for language engineering and domain modelling, and it is also at the core of other related approaches like product lines, feature oriented development [6] and method engineering [10].

The OMG has proposed the Meta-Object Facility (MOF) [20] as the meta-modelling approach in the Model-Driven Architecture (MDA) [17], a particular incarnation of MDE. MOF has been adopted as a standard by many meta-modelling tools and frameworks, most notably by the Eclipse Modelling Framework (EMF) [22]. MDA proposes a four layer, linear meta-modelling architecture and a style of meta-modelling called *strict* in which an element of a meta-layer is the instance of exactly one element at the upper meta-level. Several authors have pointed out limitations of this approach [4, 5, 9, 10], in particular concerning the existence of only one kind of instantiation relation and the constraint of considering only two adjacent meta-levels at the same time. This is enough to cover

the *linguistic* case, where an object is instance of exactly one class, but cannot capture in addition *ontological* instantiation relations within a domain. Hence, often engineers are forced to squeeze into two meta-modelling layers concepts that would naturally span several layers, resulting in more complex and cluttered models [5]. Moreover, the lack of uniformity employed in the concepts at the different layers in most approaches (e.g., UML associations are structurally different from links) makes difficult to treat in a uniform way meta-models and models, as well as to link models in different meta-levels (since “meta-” is a relative term and meta-models are also models).

Several solutions have been proposed to these problems [5, 9, 10]. Their common idea is to increase the flexibility of the meta-modelling architecture by allowing an arbitrary number of meta-levels. In [5] a mechanism called *potency* was proposed, so that one model can control the properties of models that are indirect instances of it. In [1, 4] a dual ontological and linguistic instantiation is proposed, allowing an element to be a linguistic instance (e.g. be an instance of **Class** in the upper linguistic meta-level) and also an instance of some domain concept (e.g. be an instance of **ProductType** in the upper ontological meta-level).

This paper presents METADEPTH, a new meta-modelling environment that allows modelling with an arbitrary number of ontological meta-levels. It implements the *potency* concept and permits a dual ontological and linguistic instantiation. It advances from other similar frameworks [2, 3] in that it supports advanced modelling concepts, like (OCL) constraints, derived attributes and transactions, and allows controlling whether ontological instance models can be linguistically extended. The purpose of the framework is to permit experimentation with this alternative way of meta-modelling, but at the same time provide a scalable, efficient system that permits its industrial use. Hence, METADEPTH can work in an interpreted mode, where a stack of models can be kept and work with, and then allows compilation to obtain specialized code (in the line of JMI [23]) and optimized performance. The framework is integrated with the Epsilon languages [8], which permits using the Epsilon Object Language (EOL) [14] as an action language to define behaviour for meta-models, as well as the Epsilon Validation Language (EVL) [15] for expressing constraints. Both EOL and EVL are extensions of OCL. To the best of our knowledge, no framework with similar characteristics exists nowadays. Moreover, the interplay of potency with constraints, actions, multiplicities and association ends has not been properly addressed in the literature, nor has been the mechanisms and benefits for controlling linguistic extensions. We also aim to contribute clarifying these issues.

Paper organization. Section 2 reviews multi-level meta-modelling and the concept of potency. Section 3 details the architecture of METADEPTH. Section 4 presents two case studies that show the benefit of our approach. In the first one, we define a multi-level language through a unique meta-model. For example, one can think that in UML it is natural that **Objects** are instances of **Classes**, and hence should belong to a lower meta-level (so that an object diagram is an instance of a class diagram). Our framework naturally allows this, with the benefit of a less complex language definition. In the

second case study, we solve the impedance mismatch arising when one needs to relate models at different meta-levels (a complicated technical issue in two-level frameworks such as EMF). Section 5 compares with related approaches and Section 6 concludes. A beta version of the tool can be downloaded from <http://astreo.ii.uam.es/~jlara/metaDepth/>.

2 Deep Meta-Modelling

Some authors have pointed out the limitations of considering only two meta-modelling levels at the same time [5, 9], either for language engineering or for domain modelling. A common example is the item-description or the type object pattern [16], where one needs to design a language containing both *ProductTypes* (e.g. Books) and *Products* (e.g. the book “Moby Dick”). In the classical meta-modelling approach, one would propose a two-level solution like the one to the left of Fig. 1. Although this solution is valid, it has some drawbacks. First, the user has to manually maintain the `type` links between each instance of `Product` and its `ProductType` at the model level. These links are indeed a (manually maintained) form of ontological instantiation relation for which the system does not provide automatic conformance checks. Should we have inheritance between types at the model level, it would have to be emulated manually too. Hence, this solution squeezes three meta-levels into two.

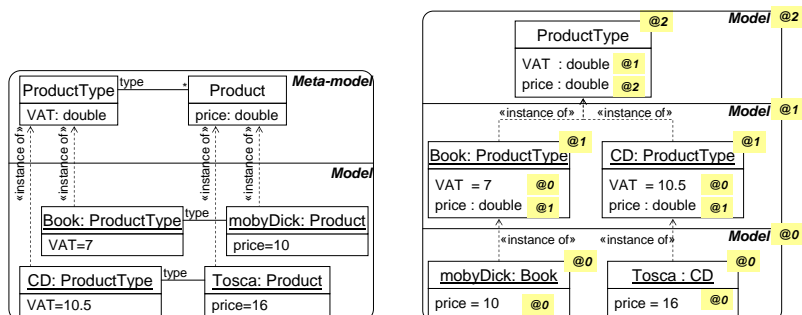


Fig. 1. A meta-model and a model including the *Type Object* pattern (left). The same system using deep meta-modelling (right), adapted from [16].

The solution to the right explicitly organizes the domain concepts into three levels. In this way, the `ProductType` is declared at the top-level, the different kinds of `ProductTypes` at the following meta-level, and the instances of these at the bottom level. This solution reduces accidental complexity, as one does not need the artificial class `Product` that the solution to the left introduced³. More-

³ One can still eliminate `Product` in the two-level solution by moving `Book` and `CD` to the top meta-model and setting them as subclasses of `ProductType`. However this solution is not valid if we need to add new kinds of `ProductTypes` at run time.

over, the instantiation of `ProductTypes`, `Books` and `CDs` is handled by the system thus enabling automatic conformance checks. Note that this pattern is ubiquitous in the definition of many languages, for example in UML (classes/objects), in web modelling languages (node types/node instances), role access control languages (user types/users) and so on. We call this style of meta-modelling, which considers more than two levels, *deep* or *multi-level meta-modelling*.

The solution meta-model to the left of Fig. 1 is able to control the attributes that instances of `Product` (`mobyDick`, `Tosca`) have. Hence, in deep meta-modelling we would expect the language designer to have the same level of control over *indirect* instances two or more meta-levels below. For this purpose, *potency* was proposed in [4] as a way to express how many times a property needs to be instantiated down the meta-levels before we get a *plain* instance and hence we have to assign it a value. The potency is a natural number that is assigned to properties, and which gets decremented each time we go down a meta-level. Hence, in our example, property `VAT` is assigned a value in the next meta-level, and `price` two meta-levels below. Not only properties can have potency, but also classes and associations. As we will see later, `METADEPTH` allows assigning potency to models, constraints and derived attributes as well.

Considering the solution to the right of Fig. 1, one realizes that the elements in the middle meta-level have both type and instance facets. This is so because they are instances of `ProductType` and, as they have potency bigger than zero, can be instantiated in its turn. The term *clabject* was coined in [4] to refer to elements with a dual type/instance facet.

The `«instance of»` relation between the elements in different meta-levels is *ontological*, as this is a relation within the domain (i.e. `mobyDick` is a `Book` and this a `ProductType`). At the same time, as we have to use a modelling language to build the models, we can argue that `Book` is an instance of `Class` and `mobyDick` an instance of `Object` (if such concepts exist in the language). One way to support this duality is to introduce another instantiation dimension, called *linguistic*, and to have a meta-model that governs the linguistic constructions used by the models at the different meta-levels. This situation is depicted in Fig. 2, where the linguistic meta-model contains concept `Clabject` with property `potency` to allow its instantiation in any meta-level. An important issue is that the union of the models in the three ontological meta-levels is a *strict* instance of the linguistic meta-model.

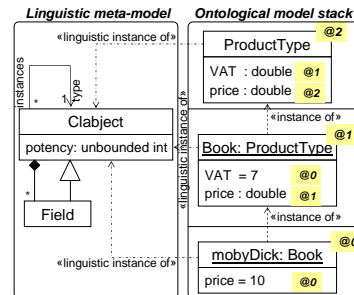


Fig. 2. Dual classification.

After having introduced the basics of deep meta-modelling, there are still missing details. For example, how could constraints be introduced in the different ontological models? Should these and other elements like association ends be given a potency? Finally, one may wonder whether, as elements `Book` and `CD` have a type facet, we could declare new attributes for them, or whether we

could *linguistically extend* a certain ontological model by introducing new types, instances of elements of the linguistic meta-model. In other words, do we demand strictness of the ontological $\ll\text{instance of}\gg$ relation? Next section introduces METADEPTH’s architecture, and discusses these issues.

3 The Architecture of MetaDepth

METADEPTH is a new meta-modelling system that we started to develop in 2008, based on the experience we gained with AToM³ [7] in previous years. AToM³ was a Python-based tool for the definition of the syntax of visual languages by meta-modelling and their semantics by graph transformation. METADEPTH is a completely rebuilt kernel, written in Java, which uses the deep meta-modelling approach presented in previous section. It can work in two ontological instantiation modes: *strict* and *extensible*, as shown in Fig. 3.

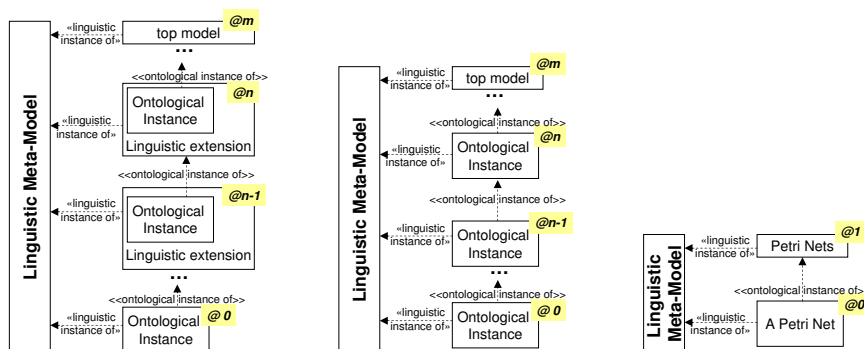


Fig. 3. METADEPTH instantiation schemes: extensible ontological instantiation (left) and strict ontological instantiation (center). Example of strict instantiation (right).

In the *extensible* case, each ontological instance model can be linguistically extended using the “horizontal” instantiation provided by the linguistic meta-model. Hence, instances of elements marked as **ext** can be extended with new attributes. A complete model can also be marked as **ext**, which means that it can be added new types and that all its elements (except those explicitly marked as **strict**) can be extended. This situation is shown to the left of Fig. 3. In all cases strictness is kept for the linguistic instantiation dimension.

The *strict* case is closer to standard meta-modelling environments, where the top-level meta-model hard-codes all language concepts and can be subsequently instantiated ontologically, but such instances cannot be linguistically extended. This situation is represented in the center of Fig. 3. In this mode one could use the highest meta-level to describe e.g. the MOF meta-model with potency 2, such model could be ontologically instantiated to describe meta-models for languages at potency 1, which in their turn could be instantiated to models of potency 0.

Hence, the strict mode is similar to most meta-modelling environments (although without restrictions on the number of meta-levels). The right of the figure shows a simple case where the linguistic meta-model is directly used to define a meta-model for Petri nets, which is instantiated into a Petri net model.

Note that allowing linguistic extensions adds extra flexibility to this meta-modelling framework in two senses. First, at any potency bigger than zero clabjects retain its *type* facet, and hence can be allowed for linguistic extension (i.e. to define attribute types). On the other hand, it is often convenient to extend models by allowing the introduction of new linguistic elements, e.g. to adapt languages to particular usages, as we will see throughout the paper.

3.1 The Linguistic Meta-Model

METADDEPTH's linguistic meta-model took MOF as inspiration, but we have modified it to accommodate an arbitrary number of meta-levels, deep instantiation and potency. Fig. 4 shows a fraction of it, where the uncoloured concrete classes are those the designer typically instantiates when building a model (i.e. Model, Node, Edge, Field and DerivedField).

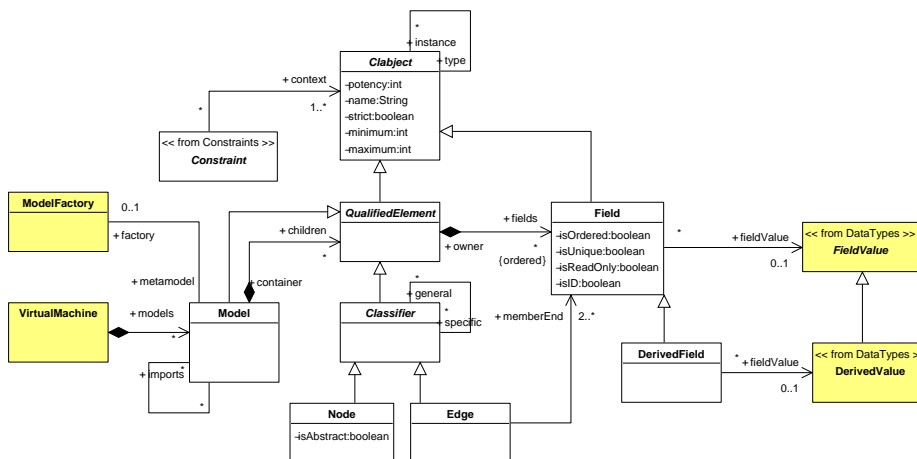


Fig. 4. METADDEPTH's linguistic meta-model, partially shown.

The root class **Clabject** takes responsibility of handling the dual type/object facet of elements. As such, it holds a potency value, as well as links to its type and instances. The potency can be unlimited, so that such clabject can be instantiated by a clabject of arbitrary potency (included unlimited). **Clabjects** also define a minimum and maximum multiplicity to control the cardinality of its instances within a given context. **Constraints** can be attached to clabjects, have a potency, and can specify on which events they should be evaluated (e.g. when creating or deleting the clabject). **QualifiedElements** are **Clabjects** owning some

field. `Models`, `Nodes` and `Edges` are all `QualifiedElements`. `Classifiers` are a special kind of `QualifiedElement` that can form general/specific hierarchies, and are refined into `Nodes` and `Edges`. The latter has two or more association ends modelled as `Fields`. Finally, derived fields are fields of which their value is automatically calculated.

METADEPTH's modularity mechanism is based on the notion of `Model`, which can be nested, as shown by the composite pattern used. `Models` are `QualifiedElements` and hence can own `Fields` and have associated `Constraints`. Each model with potency bigger than zero has an associated `ModelFactory`, in charge of instantiating the clobjects defined in such model. All working models are managed by a `VirtualMachine` container, which is a *singleton* object.

The framework supports the usual *atomic* data types, like integers, floating point numbers, strings, etc.; user-defined enumerations, as well as ordered and unordered collections with unique or non-unique elements.

3.2 Tool Support, Compiled and Interpreted Modes

METADEPTH models can be built through the provided Java API, or through a `CommandShell` and a textual syntax – similar to the Human Usable Textual Notation (HUTN) [19] – that we have built with ANTLR [21]. Loading and storing models is also done in this format. As an example, Listing 1 shows how the three models in Fig. 2 are defined with the textual syntax. The top-most model `Store` is assigned potency 2, which means it can be instantiated in the following two lower meta-levels. All elements defined inside `Store` have the potency of the container clobject, hence we only need to explicitly declare potencies different from 2 (field `VAT` in this case). Besides we give an initial value to the fields of `ProductType`. Although all elements in the listing are given an explicit name, it is not mandatory: we can declare *anonymous* clobjects like `Book{price=10;}`, and the system assigns them a unique UUID-based identifier.

```

1 Model Store@2 {
2   Node ProductType {
3     VAT@1 : double = 7.5;
4     price : double = 10;
5   }
6 }
7 Store Library {
8   ProductType Book { VAT = 7; }
9 }
10 Library MyLibrary {
11   Book mobyDick { price=10; }
12 }
```

Listing 1: A simple METADEPTH three-model stack.

The framework is fully integrated with Epsilon, a family of languages built on top of the Epsilon Object Language (EOL), which extends OCL with expressions permitting secondary effects such as assignments and methods. The integration was possible because EOL communicates with the models through a connectivity layer. Thus EOL can work with EMF models, but also with any

other model technology that implements the interface of this connectivity layer. We implemented such interface and provided support to make EOL aware of the multiple ontological levels. The solution is very practical as one can use EOL programs e.g., to build models as in Listing 2. The listing shows a typical interaction with the command line interpreter. In line 1 we enter in the context of model *MyLibrary* (we could have created a new model as well). In line 2 we begin writing the EOL program (which could be loaded from a file as well). Then the program inserts 1000 new books in the model and initializes their price.

```

1 context MyLibrary                               3 for (i in Sequence{1..1000}) {
  ::entering context MyLibrary                    4   var b: new Book;
2 # EOL                                           5   b.price:=10+i/500;
  :: entering eol execution mode                  6 }

```

Listing 2: A simple EOL program to populate a METADEPTH model.

We can use EOL not only for initializing models, but also to define its behaviour. As an example, Listing 6 shows a Petri net simulator we have built for Petri net models. Moreover, we can use the rest of the Epsilon languages with our METADEPTH models so that it is possible e.g., to transform models with the Epsilon Transformation Language [8].

Another feature of METADEPTH concerns undoability of actions. Using the *Command* pattern, all API calls are recorded in an event list, and each command provides an appropriate undo function. This allows undo/redo of any action on models, and permits integration with the transaction syntax of EOL.

By default, METADEPTH works in interpreted mode. This allows for flexible modelling and is useful for rapid prototyping of languages, as one can evolve models and meta-models at the same time. One can also create several independent models in the same *VirtualMachine* (i.e. models do not need to be related through instantiation). Once the meta-models are ready, they can be *compiled* if so desired. We have built a code generator that produces specialized classes inheriting from the classes in the meta-model of Fig. 4, as well as interfaces declaring getter, setter and creation methods that follow the JMI specification [23]. This enables interface compatibility with applications that handle JMI meta-data (like those of the EMF), and improves performance as we generate optimized code which improves, e.g., constraint evaluation, objects creation and field access. However, compiled meta-models are less flexible because they can no longer be modified, even though all its properties are readily accessible. The compilation we have implemented is more complex than in normal two-level meta-modelling frameworks, since compiling a model with certain potency implies compiling all direct and indirect models above in the ontological meta-level hierarchy. The compilation also generates a specialized command shell that initializes the *VirtualMachine* with the compiled meta-models and allows their instantiation.

3.3 Constraints and Derived Attributes

Constraints and actions can be defined using Java or EOL, nonetheless METADEPTH's design makes easy to plug in additional languages. Constraints and derived attributes have an assigned potency that governs the meta-level at which they have to be evaluated. For example, Listing 3 modifies the running example by adding a few constraints and a derived attribute on the top-level model.

```

1 Model Store@2 {
2   Node ProductType@2 {
3     VAT@1      : double = 7.5;
4     price@2    : double = 10;
5     discount@2: double = 0;
6     minVat@1  : $self.VAT>0$
7     minPrice@2: $self.price>0$
8     maxDisc@2 : $self.VAT*self.price
9               *0.01+self.price<self.discount$
10              /finalPrice@2: double =
11                $self.VAT*self.price/100
12                +self.price-self.discount$;

```

Listing 3: Constraints and derived attributes in METADEPTH.

The previous listing adds property `discount` to `ProductType`, declares three constraints in lines 6, 7 and 8, and defines a derived field in line 9. Constraints are specified between two “\$” symbols, preceded by their identifier, and can be declared inside the context of a clobject (as done in this case), or be declared outside and then explicitly assigned to one or more clobjects, promoting reusability. The constraint in line 6 has potency 1, therefore it will be evaluated in the next meta-level below. This constraint cannot access the value of fields with bigger potency, like `price`, as these may not have a value⁴. The default language for constraints is EOL, but one can also use Java. For example, the equivalent Java code to the constraint in line 6 is `minVat[Java]@1: $((Integer)self.getValue("VAT"))>0$`, which is more verbose but permits interacting with external Java programs.

Constraint `maxDisc` is more interesting as it uses fields with potency 1 and 2. This is allowed as, from the point of view of action and constraint languages, fields whose value is given in a type are accessed in the same way as fields whose value is given in the instance. In our example, `mobyDick` interprets `VAT` like a static field for which its value was set at the upper meta-level. This feature simplifies writing constraints spanning several meta-levels.

Finally, `ProductType` defines the derived field `finalPrice`. Its declaration is similar to a normal field, but it is preceded by a backslash, and includes a calculation function in EOL or Java that can use fields with a lower potency. Our current implementation calculates derived field values in a lazy way, whenever they are accessed by some getter function. This works well in textual modelling environments, but we foresee the need for a change propagation algorithm in case some exogenous observer (e.g. a graphical visualization) needs the value.

⁴ In our case `price` does have a value as it has been initialized with the value 10, but this is not the general case.

3.4 Controlling Linguistic Extensions

METADEPTH supports both *strict* and *extensible* ontological instantiation, the latter being the default. Linguistic extension is interesting to permit unforeseen extensions to Domain Specific Languages (DSLs) spawning more than one level, as our running example. In these languages, the top-most meta-model is usually highly generic, and hence extensions at lower levels are often required.

Listing 4 shows an extension of the running example, where an extensible instantiation of model `Store` is used to define a `Library`. In this usage scenario we are interested in associating an author to `ProductType` instances (i.e. to `Books`). Thus, we add to the library a new node `Author`, instance of `Node` in the linguistic meta-model. For the sake of illustration, `Author` is provided with the constraint `nonRep` that forbids replicating names. This shows that `allInstances` effectively returns all ontological instances of `Author`. Anyhow, we could have just assigned the modifier `{id}` to the field `name` to obtain the same behaviour. Please note that `Library` is still a strict instance of the linguistic meta-model.

```

1 Store Library {
2   ProductType Book {
3     VAT      = 7;
4     title   : String;
5     author  : Author;
6   }
7   Node Author {
8     name    :String;
9     nonRep@1:$Author.allInstances().
10
11     forAll(x|x<>self implies
12       x.name<>self.name)$
13     books : Book[1..*]{unique};
14   }
15   Edge writer (Book.author,
16               Author.books) {
17     year : int;
18   }

```

Listing 4: Linguistic extensions and associations in METADEPTH.

`Authors` are related to one or more `Books`, which is modelled through their field `books`. The `{unique}` modifier ensures that a given `Author` is not related to the same `Book` twice. Other supported modifiers are `id` (ensures uniqueness of values among all cljects in the same context), `ordered` (retains the order of elements) and `readOnly` (forbids changing the value).

Associations can be provided with fields (i.e. similar to association classes) by explicitly defining an `Edge` between their association ends. An example is shown in lines 12–14 of Listing 4, where the `Edge` relating books and authors includes the year in which the book was written. As in UML, declaring such `Edge` has the effect of allowing the navigation from an `Author` to all its edges through `self.writer`, while the direct navigation from an `Author` to its `Books` is done by `self.books`. In the context of the `writer` edge, it is possible to navigate to the `Author` and `Books` through the `author` and `books` ends.

In the example we have made reference to the ontological types `Author` and `Book` to declare associations. However, in specific situations, it is useful to refer to linguistic types, like `Node`, when defining association ends. This makes sense if

we want to specify that a certain association end is to be taken by any (linguistic) instance of `Node`. As next section will show, this is especially useful if we want to relate ontological models of different potency.

Listing 5 shows an example of *strict* meta-modelling. It is a meta-model for Petri nets containing `Places` and `Transitions` (both inheriting from `NamedElement`), as well as weighted arcs. All these elements inherit the *strict* modifier from their container model. In the example, `NamedElement` and its children have the same potency 1, but `METADEPTH` also allows clabjects in a hierarchy to define different potencies. A clabject keeps the biggest potency of all its ancestors.

```

1 strict Model PetriNets@1{
2   abstract Node NamedElement { name : String{id}; }
3   Node Place : NamedElement {
4     tokens      : int = 0;
5     outTrans    : Transition[*] {ordered,unique};
6     inTrans     : Transition[*] {ordered,unique};
7     minTokens   : $self.tokens>0$
8   }
9   Node Transition : NamedElement {
10    inPlaces    : Place[*] {ordered,unique};
11    outPlaces   : Place[*] {ordered,unique};
12  }
13  Edge ArcPT(Place.outTrans, Transition.inPlaces){ weight : int = 1; }
14  Edge ArcTP(Transition.outPlaces, Place.inTrans){ weight : int = 1; }
15  minWeight(ArcPT, ArcTP) : $self.weight>0$
16  minPlaces : $Place.allInstances()->size()>0$
17 }

```

Listing 5: A meta-model for Petri nets.

The Petri net meta-model defines several constraints. In the context of the model, `minPlaces` restricts nets to have at least 1 place (line 16). Most meta-modelling approaches do not allow global constraints, but some constraints (like `minPlaces`) are inherently global and do not fit in the context of any class in the meta-model. As all `METADEPTH` elements have built-in cardinalities (see Fig. 4), we can obtain the same restriction as `minPlaces` by replacing line 3 with “`Node Place[1..*] : NamedElement {`”. Constraint `minWeight` is also defined globally, but it is assigned to both kinds of arcs to enforce a positive weight (line 15). This has the advantage of promoting reusability as the constraint does not have to be defined twice.

For the sake of completeness, Listing 6 shows a simulator for Petri nets specified through EOL. EOL allows adding operations on meta-classes, and we have used this feature to define the operations `enabled` and `fire` on node `Transition`. The first one contains pure OCL code, which checks if the transition is enabled. For this purpose it iterates through all incoming arcs checking that the number

of tokens in the pre-place is bigger or equal than the arc's weight. Operation `fire` has secondary effects: the removal and addition of tokens to the pre- and post-places of the transition. The main simulation loop is defined in line 1, which iterates on all transition instances while there is some enabled, and then fires it.

```

1 while (Transition.allInstances()->exists(t |
1     t.enabled() and t.fire())) {}
2 operation Transition enabled() : Boolean {
3     return self.ArcPT->forAll(arc | arc.inPlaces.tokens>=arc.weight);
4 }
5 operation Transition fire() : Boolean {
6     for (arc in self.ArcPT)
7         arc.inPlaces.tokens := arc.inPlaces.tokens-arc.weight;
8     for (arc in self.ArcTP)
9         arc.outPlaces.tokens := arc.outPlaces.tokens+arc.weight;
10    return true;
11 }

```

Listing 6: A simulator for Petri nets defined with EOL.

4 Case Studies

This section presents two case studies that show the usefulness of `METADEPTH` and help illustrating some of its distinguishing features, such as the use of linguistic types or the interplay of potency, multiplicities and association ends.

4.1 Defining Multi-level Languages

The first example shows the use of deep meta-modelling for defining DSLs spawning more than one meta-level. This is the case of many languages that implement the *Type Object* pattern. For example, UML defines class and object diagrams as two different structural diagrams. However, UML defines both in the same meta-level, with the drawback that one has to maintain explicit relations between objects and their classifiers and ensure that they remain consistent. Instead, one can use deep meta-modelling to simplify the language definition and to automate the maintenance of consistency between classes and objects.

Listing 7 shows how a simple language containing class and object diagrams is defined in `METADEPTH`. The idea is specifying a three-level meta-modelling architecture where the top-most level contains the definition of class diagrams and potency 2 (Model `ClassDiagram` in Listing 7). In this way, in the next meta-level we can build class diagrams (e.g., `Zoo` in Listing 8), and in the bottom meta-level we can build object diagrams and the very meta-modelling infrastructure handles the type checking w.r.t. the class diagram that the object diagrams

instantiate. In this way, a stack of two languages is defined with just the model in Listing 7. This model is *strict* to avoid the creation of new linguistic types in class and object diagrams, and permit only the creation of **Classes** and **Assoc**s. On the contrary classes and associations are *extensible* to allow their instances to define new fields in them. Node **Class** contains field **isAbstract** to designate whether the class is abstract or not, and constraint **noAbsObjects** ensures that object diagrams (two levels below) do not contain objects whose class is abstract.

```

1 strict Model ClassDiagram@2 {
2   ext Node Class {
3     isAbstract@1 : boolean = false;
4     in           : Class[*];
5     out          : Class[*];
6     noAbsObjects : $self.isAbstract=false$
7   }
8   ext Edge Assoc(Class.out,Class.in);
9 }

```

Listing 7: A meta-model for class and object diagrams (meta-level 2).

Listing 8 shows an instance of **ClassDiagram**, namely a class diagram named **Zoo** which declares two classes and one association. Class **Person** declares one field (**name**) and one association end (**pet**). The field is a linguistic extension of **Person** (i.e. it is not an instance of any feature in the upper ontological meta-level), whereas the association end is an ontological instance of the association end **out** defined for **Class** in the upper meta-level (indicated as a modifier after its declaration). The listing also shows an instance of the defined class diagram called **myZoo**.

```

1 ClassDiagram Zoo {
2   Class Person {
3     name : String {id};
4     pet  : Animal[*] {out};
5   }
6   Class Animal {
7     kind : String {id};
8     owner : Person[1..*] {in};
9   }
10  Assoc hasPet (Person.pet,
11                Animal.owner)
12    { since : int; }
13  }
14  Zoo myZoo {
15    Person p{ name="Juan"; }
16    Animal a{ kind="monkey"; }
17    hasPet(p,a){since=2010;}
18  }

```

Listing 8: A class and an object diagram (meta-levels 1 and 0, respectively).

Coming back to our language specification in Listing 8, we can see that fields **in** and **out** have an (unbounded) multiplicity. This multiplicity constrains the

meta-level immediately below, and hence one can have an unbounded number of instances of them in a class diagram. In their turn, these instances can declare their own multiplicity, which affects their instantiation in object diagrams. Anyhow `in` and `out` have potency 2 and are also available in object diagrams, storing the content of all their instances. For example, `p.out` evaluates to `[a]` because `p.pet` evaluates to `[a]`. This shows that the instantiation semantics for fields is coherent with the way of handling indirect instances of all other elements, like `Nodes` and `Edges`, as e.g., both `p` and `a` are indirect instances of `Class`.

4.2 Relating Models at Different Meta-levels

Next we illustrate how to relate models at different meta-levels in `METADEPTH`. This is of practical importance as, in standard approaches like in `EMF`, models cannot be meta-models at the same time. Models can be treated as meta-models by passing through a transformation called *promotion*, thus making difficult to link elements of models with elements of meta-models. As an example we show a multi-level language to define the graphical concrete syntax of meta-models. The purpose of the example is neither showing advanced concrete syntax concepts, nor discussing how instances would be graphically rendered, but only show how models can be naturally put in relation with meta-models. In order to keep the example simple we restrict to the visual representation of `Nodes` as rectangles.

```

1 Model Graphics@2 {                               8 }
2   abstract Node Figure {                       9   Node Rectangle : Figure {
3     x: int;                                     10     width@1 : int;
4     y: int;                                     11     height@1 : int;
5     rotation : double = 0;                     12   }
6     scale : double = 1;                         13 }
7     refsTo : Node;

```

Listing 9: The Graphics meta-model.

Listing 9 shows the meta-model for the two-level language called `Graphics`. When instantiated in the next meta-level it allows defining visualizations for a meta-model `M`. Then, by instantiating it again we obtain instances with the rendering information about the instances of `M`. This situation is depicted in Fig. 5, which shows that `Graphics` models are associated to the definition of a language, and the rendering information to models of this language. The `Graphics` definition contains an abstract class `Figure` with fields `x` and `y` that store the absolute

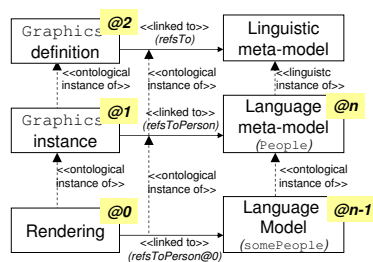


Fig. 5. Scheme of the example.

position of instances of figures two meta-levels below, as well as its rotation and scaling. Figures also contain field `refersTo` to point to the `Node` that the figure is representing. A real-world model for concrete syntax would include several types of figures (circles, text, etc.) that could be composed to form the visual concrete syntax of a `Node`. In our case, for space constraints, we consider just `Rectangles` which can be configured with their dimension.

Then, consider the simple language defined by the meta-model `People` in Listing 10. We can instantiate the `Graphics` meta-model to define a visual representation for `Person` in the language meta-model. Moreover, since meta-model `Graphics` is generic, we can also define a visualization for the class diagram `Zoo` in Listing 8 by just replacing “`People`” by “`Zoo`” in line 4. This shows the flexibility of `METADEPTH` to relate models at different potencies and meta-levels, and the advantages of treating uniformly models at different meta-levels.

```

1 Model People@1 {                               6     width = 15;
2   Node Person { name:String; }                 7     height = 10;
3 }                                               8     refsToPerson: Person{refsTo};
4 Graphics CS imports People {                   9   }
5   Rectangle iconPerson {                       10  }
```

Listing 10: Defining the concrete syntax for a language meta-model.

Finally, we can instantiate the `People` meta-model and its associated concrete syntax as shown in Listing 11.

```

1 People somePeople {                             6     x = 10; rotation = 90;
2   Person e { name = "Esther"; }                 7     y = 10; scale = .5;
3 }                                               8     refsToPerson = e;
4 CS cs1 imports somePeople {                     9   }
5   iconPerson {                                  10  }
```

Listing 11: Instantiating the meta-model and its concrete syntax.

5 Related Work

There are two main lines of related research: those works following a MOF-like way of meta-modelling, where only two adjacent levels are considered at the same time, and those following a deep meta-modelling approach.

Regarding the first group, MOF [20] is the OMG’s language to specify meta-models and the most adopted approach in practice. The MOF specification is divided in two parts: a basic one called Essential MOF (EMOF), and a more

advanced one called Complete MOF (CMOF). The specification claims that it can be used with as many meta-levels as users demands. However there are conceptual problems, e.g. when one needs to introduce data types instances, as the basic data types would have to be replicated across the different meta-levels. Moreover, current implementations only allow handling two levels at the same time. Even though MOF specifies a set of reflective services, the specification neglects that considering three or more meta-levels at the same time requires some entities to simultaneously have both a type and an object facets. The main implementation of MOF is integrated within EMF [22] and is called *Ecore*. It forces a tree-based edition of models, and only supports EMOF, therefore lacking useful constructs like a proper concept of association enabling e.g. the definition of associative classes. Neither EMOF nor CMOF specify how to define constraints or actions to calculate derived attributes.

Many current meta-modelling research efforts revolve around MOF. For example, KM3 [12] is a DSL to specify meta-models based on MOF. KM3 has a textual front-end for meta-modelling frameworks, and as such can be compared to the textual notation we have developed for METADEPTH, but does not introduce new meta-modelling concepts. Kermeta [13] is another textual language whose purpose is to specify behaviours for EMF meta-models. Hence, its role for EMF meta-models is equivalent to our use of EOL to specify behaviour.

Concerning deep meta-modelling, several efforts can be recently found directed to a practical test of the seminal ideas of Atkinson and Kühne [4]. For example, DeepJava [16] is an extension of Java with the concept of potency and, as such, it cannot be considered a meta-modelling framework. It provides methods with potency, but has to use special keywords to navigate up the type hierarchy in order to find attribute values. On the contrary, our constraints and computations for derived attributes can access type fields in a uniform way. This is similar to considering that a type attribute value is like a *static attribute* with respect to an instance, and has the advantage that one does not need to know exactly how many meta-levels up the given field was given a value, and facilitates the integration with constraint and action languages.

The work in [3] is another recent proposal for deep meta-modelling. The tool is currently being developed, based on *Ecore*. They consider multi-level constraints, and propose extending OCL to cope with multiple ontological meta-levels. This is similar to our approach, but we assign *potency* to constraints, making them easier to define. This is so because potency layers constraints, and hence they do not have to, e.g., explicitly invoke *allInstances()* a predefined number of times, it is enough to implement in the OCL interpreter the ability to recognise indirect ontological instances of clobjects, and to interpret fields of ontological types similar to static attributes in Java. Another difference concerns relations, as they do not consider association ends, but add this information inside the relation class itself. The main motivation for this is graphical visualization in concrete syntax and uniformity of structure between ontological meta-levels. On the contrary, the design of our framework was not driven by concrete syntax issues, as we foresee building systems supporting graphical

syntax as well as more sophisticated textual syntaxes, probably posing different challenges. We retain association ends (similar to MOF and UML), as this allows us to reuse the multiplicity semantics of structural features. Most importantly, it makes easier the practical integration with navigation languages like OCL. We also agree on the importance of uniformity of relations at the different levels, and hence we retain association ends (i.e. fields) at all levels. Finally, they do not consider linguistic extensions, transactions, derived attributes, nor an action language, whereas we can use the Epsilon languages, which enable manipulation and transformation of models. Our dual working scheme interpreted/compiled allows rapid prototyping of languages on the one hand, and enables the generation of stand-alone, efficient domain specific tools on the other hand.

Nivel [2] is a deep meta-modelling framework based on the weighted constraint rule language (WCRL). It implements the concept of potency and the dual linguistic/ontological classification. It brings some interesting ideas from conceptual modelling, like the possibility of several classes to implement an association role. Nivel's semantics is given by its translation into WCRL, which allows some form of automated reasoning, but the kind of reasoning and its usefulness was not shown in [2]. The language lacks constraints and action languages (except WCRL itself), which hinders its use in practical MDE.

Other works that have influenced METADEPTH include Amulet [18], whose prototype/instance concept is similar to our linguistic clobject extensions.

6 Conclusions and Future Work

In this paper we have presented METADEPTH, a novel framework for deep meta-modelling. The tool supports the concept of potency, allowing an arbitrary number of ontological meta-levels. It provides advanced features like multi-level constraints, derived attributes and linguistic extensions at lower meta-levels. The framework can work either in interpreted or compiled modes, favouring flexibility or efficiency. The current implementation offers a textual syntax, inspired by HUTN, and is integrated with the Epsilon languages. In particular, we have shown the use of EOL for specifying actions and EVL for constraints.

After more than one year of development, we are excited about the possibilities opened by METADEPTH, and we will continue to improve it in the near future. For example, we would like to allow the framework to run in client/server mode, so that the kernel can be accessed through web services. We would also like to build a system to support a graphical concrete syntax, in the spirit of the old AToM³ [7], but allowing interaction through the web navigator. Even though we can use ETL now, the plan is to incorporate a formal model transformation language into it, and for this purpose we are working on an implementation of our pattern-based transformation language [11]. It could be interesting to study the implications of deep meta-modelling for model transformation, and for this we would need a formalization of the framework. Finally, we are also enthusiastic about deep meta-modelling and the new possibilities it offers for MDE. We are

currently exploring idioms, and identifying good practices and patterns for deep meta-modelling and multi-level language engineering.

References

1. J. M. Álvarez, A. Evans, and P. Sammut. Mapping between levels in the metamodel architecture. In *UML'01*, volume 2185 of *LNCS*, pages 34–46. Springer, 2001.
2. T. Asikainen and T. Männistö. Nivel: a metamodeling language with a formal semantics. *SoSyM*, 8(4), 2009.
3. C. Atkinson, M. Gutheil, and B. Kennel. A flexible infrastructure for multilevel language engineering. *IEEE Trans. Soft. Eng.*, 35(6):742–755, 2009.
4. C. Atkinson and T. Kühne. Rearchitecting the UML infrastructure. *ACM Trans. Model. Comput. Simul.*, 12(4):290–321, 2002.
5. C. Atkinson and T. Kühne. Reducing accidental complexity in domain models. *SoSyM*, 7(3):345–359, 2008.
6. D. S. Batory. Multilevel models in model-driven engineering, product lines, and metaprogramming. *IBM Systems Journal*, 45(3):527–540, 2006.
7. J. de Lara and H. Vangheluwe. AToM³: A tool for multi-formalism and meta-modelling. In *FASE'02*, volume 2306 of *LNCS*, pages 174–188. Springer, 2002.
8. Epsilon. <http://www.eclipse.org/gmt/epsilon/>, 2009.
9. C. González-Pérez and B. Henderson-Sellers. A powertype-based metamodeling framework. *SoSyM*, 5(1):72–90, 2006.
10. C. González-Pérez and B. Henderson-Sellers. *Metamodeling for Software Engineering*. Wiley, 2008.
11. E. Guerra, J. de Lara, and F. Orejas. Pattern-based model-to-model transformation: Handling attribute conditions. In *ICMT'09*, volume 5563 of *LNCS*, pages 83–99. Springer, 2009.
12. F. Jouault and J. Bézivin. KM3: A DSL for metamodel specification. In *FMOODS'06*, volume 4037 of *LNCS*, pages 171–185. Springer, 2006.
13. Kermeta. <http://www.kermeta.org/>.
14. D. S. Kolovos, R. F. Paige, and F. Polack. The Epsilon Object Language (EOL). In *ECMDA-FA'06*, volume 4066 of *LNCS*, pages 128–142. Springer, 2006.
15. D. S. Kolovos, R. F. Paige, and F. Polack. On the evolution of OCL for capturing structural constraints in modelling languages. In *Rigorous Methods for Software Construction and Analysis*, 2007.
16. T. Kühne and D. Schreiber. Can programming be liberated from the two-level style? – Multi-level programming with DeepJava. In *OOPSLA'07*, pages 229–244. ACM, 2007.
17. S. J. Mellor, K. Scott, A. Uhl, and D. Weise. *MDA Distilled*. Addison-Wesley Object Technology Series, 2004.
18. B. Myers, R. McDaniel, R. Miller, A. Ferrenzy, A. Faulring, B. Kyle, A. Mickish, A. Klimovitski, and P. Doane. The Amulet environment: new models for effective user interface software development. *IEEE Trans. Soft. Eng.*, 23(6):347–365, 1997.
19. OMG. HUTN. <http://www.omg.org/cgi-bin/doc?formal/2004-08-01>, 2009.
20. OMG. MOF 2.0. <http://www.omg.org/spec/MOF/2.0/>, 2009.
21. T. Parr. ANTLR. <http://www.antlr.org>, 2010.
22. D. Steinberg, F. Budinsky, M. Paternostro, and E. Merks. *EMF: Eclipse Modeling Framework, 2nd Edition*. Addison-Wesley Professional, 2008. See also <http://www.eclipse.org/modeling/emf/>.
23. Sun. Java Metadata Interface. <http://java.sun.com/products/jmi/index.jsp>.