

An Approach for Modeling Architectural Design Rules in UML and its Application to Embedded Software

ANDERS MATTSSON, Combitech AB, Sweden and University of Limerick, Ireland

BRIAN FITZGERALD, University of Limerick, Ireland

BJÖRN LUNDELL, University of Skövde, Sweden and University of Limerick, Ireland

BRIAN LINGS, University of Skövde, Sweden

Current techniques for modeling software architecture do not provide sufficient support for modeling architectural design rules. This is a problem in the context of model-driven development in which it is assumed that major design artifacts are represented as formal or semi-formal models. This article addresses this problem by presenting an approach to modeling architectural design rules in UML at the abstraction level of the meaning of the rules. The high abstraction level and the use of UML makes the rules both amenable to automation and easy to understand for both architects and developers, which is crucial to deployment in an organization. To provide a proof-of-concept, a tool was developed that validates a system model against the architectural rules in a separate UML model. To demonstrate the feasibility of the approach, the architectural design rules of an existing live industrial-strength system were modeled according to the approach.

Categories and Subject Descriptors: D.2.11 [Software Engineering]: Software Architectures—*Languages*; D.2.2 [Software Engineering]: Design Tools and Techniques—*Computer-aided software engineering (CASE)*; *Object-oriented design methods*; D.2.1 [Software Engineering]: Requirements/Specifications—*Methodologies*; *Tools*

General Terms: Design, Documentation, Human Factors

Additional Key Words and Phrases: Model-driven development (MDD), model-driven engineering (MDE), embedded software development

ACM Reference Format:

Mattsson, A., Fitzgerald, B., Lundell, B., and Lings, B. 2012. An approach for modeling architectural design rules in UML and its application to embedded software. *ACM Trans. Softw. Eng. Methodol.* 21, 2, Article 10 (March 2012), 29 pages.

DOI = 10.1145/2089116.2089120 <http://doi.acm.org/10.1145/2089116.2089120>

1. INTRODUCTION

A basic premise of model-driven development (MDD) [Schmidt 2006] is to capture all important design information in a set of formal or semi-formal models that are automatically kept consistent by tools. The purpose is to raise the level of abstraction at which the developers work and to eliminate time-consuming and error-prone manual work in maintaining consistency between different design artifacts such as UML diagrams and code. An important design artifact in any software development project

This research was financially supported by the ITEA project MoSiS (www.itea-mosis.org) through Vinnova (www.vinnova.se) and also by funding from Science Foundation Ireland to Lero (www.sfi.ie).

Authors' addresses: A. Mattsson, Hyacintvägen 4, SE-523 33 Ulricehamn, Sweden; email: ajmattsson@gmail.com; B. Fitzgerald, Lero – the Irish Software Engineering Research Centre, University of Limerick, Ireland; email: bf@ul.ie; B. Lundell and B. Lings, University of Skövde, P.O. Box 408, SE-541 28 Skövde, Sweden; email: {bjorn.lundell, brian.lings}@his.se.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permission may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 1049-331X/2012/03-ART10 \$10.00

DOI 10.1145/2089116.2089120 <http://doi.acm.org/10.1145/2089116.2089120>

is the software architecture [Bass et al. 2003]. The purpose of the architecture is to guide and control the design of the system so that it meets its quality requirements. A common way of capturing the architecture in MDD projects is to put the high-level structure in the form of packages and components with interfaces in the system model, together with a framework implementing a communication infrastructure used by the components [Mattsson et al. 2009]. This is, however, not enough; we also need to specify rules as to what kinds of component to put in different layers and how these are supposed to use the infrastructure. We call these rules architectural design rules [Mattsson et al. 2009]. The current state of practice is to express these rules in informal text for the developers to follow. This means that manual reviews have to be used to check that the rules have been followed during detailed design. If we could model architectural design rules in a form that could be interpreted by tools, and at the same time be easily understood by both architects and developers, we would be able to eliminate error-prone and time-consuming manual work.

In this article we present an approach to solving this problem by using the well-known modeling language UML [OMG 2009] to define architectural design rules at the meta-model level in an intuitive way. To verify that the approach can be automated, a tool has been built that checks that a system model conforms to architectural design rules modeled according to the approach. To demonstrate the applicability of the approach to real systems development, the architectural design rules of an already developed real-world embedded system have been modeled according to the approach.

The rest of this article is organized as follows. In the next section we present the background motivating the research. In Section 3 we introduce a fictional but realistic example to illustrate the problem of modeling architectural design rules and to introduce our proposed solution. Thereafter, we present the research approach adopted for the study. Following this, our findings are presented in three consecutive sections, covering the definition of the approach, tool support for automation, and the results from applying the approach to a real-world system. Finally we discuss our conclusions and the implications of the findings.

2. BACKGROUND

Our main research objective was to define an approach for modeling architectural design rules in an intuitive way while also stringent enough for automation. In order to motivate and validate our research objective, we conducted a literature review in line with Levy and Ellis [2006]. The review consisted of two consecutive phases, where the first phase focused on the role of architectural design rules in the context of MDD. The findings from this phase are presented in Section 2.1. Since we have reported these findings in Mattsson et al. [2009], where a detailed discussion can be found, this part is kept brief in the present article. Informed by the findings of the first phases the second phase of the literature review focuses on techniques for using UML to constrain UML modeling. The findings from this phase are presented in Section 2.2.

2.1 Architectural Design Rules and MDD

The purpose of the architecture is to guide and control the design of the system so that it meets its quality requirements. Bass et al. [2003] are unequivocal in stating the importance of an architectural approach:

“The architecture serves as the blueprint for both the system and the project developing it. It defines the work assignments that must be carried out by design and implementation teams and it is the primary carrier of system qualities such as performance, modifiability, and security – none of which can be achieved without a unifying architectural vision. Architecture is an artefact for early analysis to make sure

that the design approach will yield an acceptable system. Moreover, architecture holds the key to post-deployment system understanding, maintenance, and mining efforts. In short, architecture is the conceptual glue that holds every phase of the project together for all of its many stakeholders.”

A common understanding in architectural methods is that the architecture is represented as a set of components related to each other [Perry and Wolf 1992; Shaw et al. 1995]. The components can be organized into different views focusing on different aspects of the system. Different methods propose different views; there may be a view showing the development structure (e.g., packages and classes), a view showing the runtime structure (processes and objects), and a view showing the resource usage (processors and devices). In any view, each component is specified with the following:

- an interface that documents how the component interacts with its environment;
- constraints and rules that have to be fulfilled in the design of the component;
- allocated functionality; and
- allocated requirements for quality attributes.

A typical method of decomposition (see for instance, Bass et al. [2003], Wojcik et al. [2006] and Bosch [2000]) is to select and combine a number of patterns that address the quality requirements of the system and use them to divide the functionality in the system into a number of elements. Child elements are recursively decomposed in the same way down to a level where no more decomposition is needed, as judged by the architect. The elements are then handed over to the designers who detail them to a level where they can be implemented. For common architectural patterns such as model-view-controller, blackboard, or layers [Buschmann 1996], this typically means that we decompose our system into subsystems containing different kinds of classes (such as models, views, and controllers). However the instantiation into actual classes is often left to the detailed design, for two main reasons:

- (1) Functionality will be added later, either because it was missed or because a new version of the system is developed, so more elements will be added later that also have to follow the design patterns decided by the architect.
- (2) It is not an architectural concern. The concern of the architect is that the design follow the selected architectural patterns and not to do the detailed design.

This means that a substantial part of the architecture consists of design rules as to what kinds of elements, including behavioral and structural rules and constraints, should be in a certain subsystem.

The importance of architectural design rules is also highlighted in current research in software architecture that is focused on the treatment of architectural design decisions as first-class entities [Jansen and Bosch 2005; Jansen et al. 2007; Kruchten 2004; Kruchten et al. 2006; Tyree and Akerman 2005], where architectural design decisions impose rules and constraints on the design together with a rationale. However, there is not yet any suggestion on how to formally model these design rules. The current suggestion is to capture them in text and to link them to the resulting design. This may be sufficient for rules stating the existence of elements (“ontocrisis” in Kruchten [2004]) in the design, such as a subsystem or an interface, since the architect can put the actual element (i.e., a certain subsystem) into the system model at the time of the decision. It is however not sufficient for rules on potentially existing elements (“diacrisis” in Kruchten [2004]) such as rules as to what kinds of elements, including behavioral and structural rules and constraints, should be in a certain subsystem, since the actual elements are not known when the design decision is made.

Instead, the rule-based design occurs later in the detailed design phase, and involves other persons, potentially even in a different version of the system.

As previously reported [Mattsson et al. 2009], there is no satisfactory solution to how to model architectural design rules on potentially existing components in the current literature.

- Approaches to MDD, such as OMG’s MDA [OMG 2003], domain specific modeling (DSM) [Karsai et al. 2003; Tolvanen and Kelly 2005], and Software Factories [Greenfield and Short 2004] from Microsoft do not address the problem of how to represent architectural design rules.
- Numerous methods exist for architectural design such as ADD [Bass et al. 2003; Wojcik et al. 2006]; RUP 4+1 Views [Kruchten 1995, 2004]; QASAR [Bengtsson and Bosch 1998; Bosch 2000; Bosch and Molin 1999]; S4V [Hofmeister et al. 2000; Soni et al. 1995]; BAPO/CAFCR [America et al. 2004; van Der Linden et al. 2004]; and ASC [Ran 2000]. Also, current research in software architecture is focused on treating architectural design decisions as first class entities [Jansen and Bosch 2005; Jansen et al. 2007; Kruchten 2004; Kruchten et al. 2006; Tyree and Akerman 2005]. However, neither of these research streams provides any suggestion as to how architectural design rules should be modeled, other than as informal text.
- Architectural description languages (ADL) [Medvidovic and Taylor 2000; Medvidovic et al. 2002, 2007] (e.g., ACME, Aesop, C2, MeatH, AADL, SysML, and UML) do not provide sufficient means to specify constraints or rules on groups of conceptual components only partly specified by the architect where the actual components are intended to be identified and designed by developers in later design phases.

The state of the art in embedded software development [Mattsson et al. 2009] is to capture these rules in a text document. This means that we have to rely on manual reviews to ensure that the detailed design follows the architectural design rules. As a consequence, architectural enforcement becomes a bottleneck in MDD, where other design activities have been automated. As earlier report [Mattsson et al. 2009], this leads to a plethora of problems, including the following one.

- (1) *Stalled detailed design.* The design teams have to wait for the architects to review their overall design before they can dig deeper into the design.
- (2) *Premature detailed design.* Design teams commence a detailed design before their overall design is approved by the architect, with the risk that they will have to redo much work after the review.
- (3) *Low review quality.* Time pressures lead to a low quality of review, leading to problems later in the project.
- (4) *Poor communication of architecture.* The architects have no time to handle the communication with the design teams regarding architectural interpretations or problems; problems are “swept under the carpet.”

An architectural style (also known as architectural pattern) [Shaw and Garlan 1996] is an idiomatic pattern of system organization. It is comparable to the solution part of a certain kind of design pattern [Gamma 1995], specifying system wide design rules, categorized as architectural patterns in Buschmann [1996].

The problem of modeling design rules has much in common with the problem of modeling architectural styles or the solution part of a design pattern, in so far as it is basically about specifying rules to follow in the design. There are a number of suggestions on how to formally model design pattern specifications and architectural styles [Bayley 2007; Eden 2002; France et al. 2004; Lauder and Kent 1998; Mak et al. 2004;




Fig. 1. Defining a stereotype Data_Class in a UML profile.

Mikkonen 1998; Pahl et al. 2007; Zdun and Avgeriou 2005]. While some approaches use mathematical formalisms such as predicate logic and set theory, others use UML, applied at the meta-model level. Based on our experience, we believe that, in order to be successful in practice, it is essential that architectural design rules are modeled in such a way that they are both amenable to automatic enforcement of the detailed design and easy to understand and use by both architects and developers. The latter is important in order to avoid increasing the work of developing the rules; otherwise there is a risk that the work burden is increased instead of decreased, even though the enforcement is automated. Another important issue is that it should be possible to use current mainstream modeling tools to model both the architectural design rules and the system model so as to make it widely adoptable. Given that UML is probably the most widely used modeling language in the embedded software industry, our choice would therefore be to use UML to model architectural design rules for UML models. Our approach is therefore based on the same idea as in Zdun and Avgeriou [2005] and France et al. [2004], namely to use UML on the meta-model level to restrict the use of UML in a system model. However, instead of using it to specify patterns, we use it to specify architectural design rules.

2.2 Architectural Design Rules in an MDD context Using UML

The purpose of architectural design rules is to provide the necessary constraints for the detailed design. In an MDD context where the detailed design is made in UML, this means that the architectural design rules must be modeled in such a way that they restrict how UML is used. Furthermore, to suit our purpose, it must be possible to automatically enforce the restrictions on the detailed design or to automatically check that the restrictions are followed in the detailed design. Within UML, a profile provides a mechanism to restrict the use of UML [Fuentes-Fernández and Vallecillo-Moreno 2004]. A UML profile contains a number of stereotypes where each stereotype extends one or more UML meta-classes with new properties and constraints. The stereotype can then be applied to model elements of the extended meta-class in a model using the profile. In Figure 1, an example is given where we define a stereotype Data_Class that extends the UML meta-class Class. The stereotype adds the constraint that classes with the stereotype Data_Class cannot have any operations. The constraint is expressed in OCL [OMG 2003], a language for specifying constraints and queries on models in UML and other MOF-based [OMG 2006] languages defined by OMG (MOF is a subset of UML intended for meta-modeling). The application of this stereotype is shown in Figure 2 where we define a class Position with the stereotype Data.Class.

There are, however, at least two problems with defining profiles in this way. The first is that it requires detailed knowledge of the UML meta-model (the model defining the abstract syntax of UML), which is quite complex and beyond what can be expected from a typical architect or developer; it would likely impede widespread adoption of the approach [Conboy and Fitzgerald 2010]. For example, the very simple constraint in Figure 1 requires the knowledge that operation has the role name ownedOperation




Fig. 2. Defining a class position of the stereotype Data_Class in a UML model.

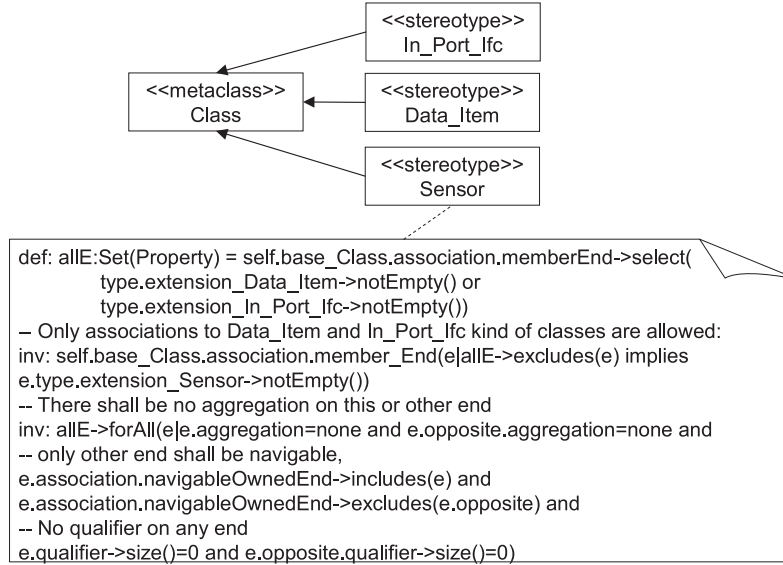


Fig. 3. Definition of the architectural design rule example using OCL.

in its association to Class in the UML meta-model. The second problem is that the OCL expressions become quite complex, even for quite simple constraints. Consider the following example rule (rule S4 in the illustrating example in Section 3).

A sensor may only have associations to In.Port.Ifcs and Data.Items. These associations shall only be navigable from the sensor.

Using the standard approach for defining profiles we get the constraint definition for the Simulator stereotype shown in Figure 3. As can be seen, this involves a great deal of detailed knowledge of the UML meta-model.

Another possibility is to make a new meta-model with classes that extend the classes in the UML meta-model through generalizations. But, as illustrated in Figure 4, this is very similar to the approach using a profile, in that we still need to specify almost the same OCL constraints as when using a profile. The only benefit is that we avoid the navigation between the stereotypes and the elements in the meta-model (e.g., self.base_Class and type.extension_Data_Type).

What is needed is a technique to specify the constraints in a more intuitive way. In Fuentes-Fernández and Vallecillo-Moreno [2004], a technique using a meta-model as a precursor to a UML profile specification is suggested. According to this approach, stereotypes are defined by classes in a meta-model where the relations between the classes impose constraints on the stereotypes. Using an approach like this, the above example would be expressed according to Figure 5.

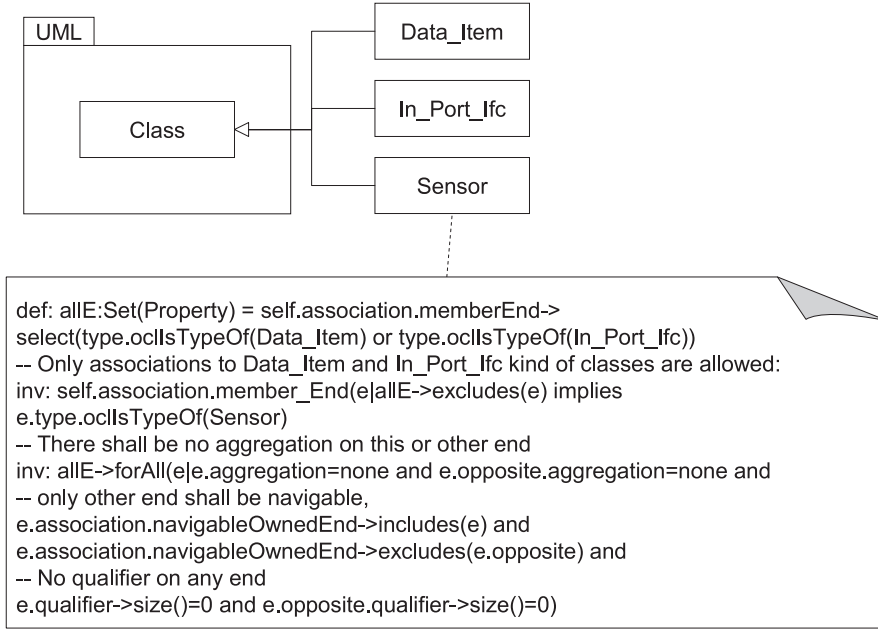


Fig. 4. Definition of the architectural design rule example using specialisation of the UML meta-model.



Fig. 5. Capturing the architectural design rule example using the approach suggested in Fuentes-Fernández and Vallecillo-Moreno [2004].

This approach has the benefit that it is more intuitive—it is both easier to model and to understand. Another benefit is that it does not contain any details from the UML meta-model, so it does not require any knowledge of that. The drawback of this approach is that it lacks rigor on how to transform it to a UML profile. In the context, addressed in Fuentes-Fernández and Vallecillo-Moreno [2004], this is not a problem because the purpose of the model is just to aid in the process of designing a profile, not to be automatically transformed into a profile. For our purpose, a detailed specification as to how it may be transformed to a UML profile is necessary, to the level where it could be implemented in a tool. To that purpose, we have defined a set of transformation rules, described in Section 5.

3. AN ILLUSTRATIVE EXAMPLE

In this section we introduce a fictional but realistic example to illustrate the problem of modeling architectural design rules and to introduce our proposed solution to that problem.

Our fictional system is a product line of washing machines. The product line consists of a wide variety of washing machines, from simple cheap machines with a minimum of features to advanced machines with user access control, monitored and controlled over the internet for industry and public self-service laundries. Since there is a high

degree of functional commonality between different machines, it was decided to build a common model from which software for all machines (existing and future) can be generated. With this goal, there are a number of nonfunctional requirements that must be addressed by the architectural design, such as the following.

- (1) *Performance scalability.* In simple machines it should be possible to run the software in a microcontroller with very limited performance and memory, while the more advanced machines have fully featured CPU's with hundreds of megabytes of memory.
- (2) *IO hardware variability.* Since the availability and price of IO hardware varies over time, change of IO hardware should require minimal effort.
- (3) *Communication protocols variability.* Since different machines use different protocols for communication with external systems now and in the future, change of communication protocols should require minimal effort.
- (4) *Functional scalability.* Since the functionality is highly variable, adding, removing, and changing functionality, including beyond what is considered currently, should require minimal effort.
- (5) *User interface variability.* There is high variability in how the user controls the machines, from simple variants with knobs and LEDs to touch screens for the most advanced machines. Therefore, it should require minimal effort to change the interface for the user, beyond the controls existing currently.
- (6) *Sensor and actuator variability.* While some machines use actual sensors to monitor water temperature and water level, others use time to estimate these values. There are also different scalings between the sensor output and measured values for different physical sensors and different machines, (e.g., for water level). Depending on the functionality of the machine, there are also different kinds of sensor and actuators, for different machines (e.g., if the machine also has a tumble-drying functionality or dirt-sensing capabilities). To cater for this variability, adding, removing, or changing sensors and actuators should require minimal effort.

To handle these requirements, the following design principles have been decided on by the architect.

- (1) *Performance scalability.* No heavyweight functionality is required. For example, it might have been sensible to use a database with remote accessibility to store the data items, since this would have eliminated the need for implementing support for remote accessibility. However, this would have made it impossible to run the software on a microcontroller.
- (2) *IO hardware variability.* The IO hardware is only accessible through a small stable set of IO interfaces. These interfaces are then implemented for the different hardwares by different IO_Ports.
- (3) *Communication protocol variability.* This is handled by the same design principle as used for handling the IO hardware variability. Different protocols are handled by different IO_Ports towards the same stable interface.
- (4) *Functional scalability.* This is handled by not allowing any dependencies on or between applications (e.g., washing program, remote monitoring, access control). An application reads and writes to Data_Items and IO_Interfaces and may act as an observer [Gamma 1995] on Data_Items reacting to changes for these.
- (5) *User interface variability.* This is handled by using the same principles for user interface controllers as for the applications described above. A user interface controller provides a mapping between a physical user interface and Data_Items.

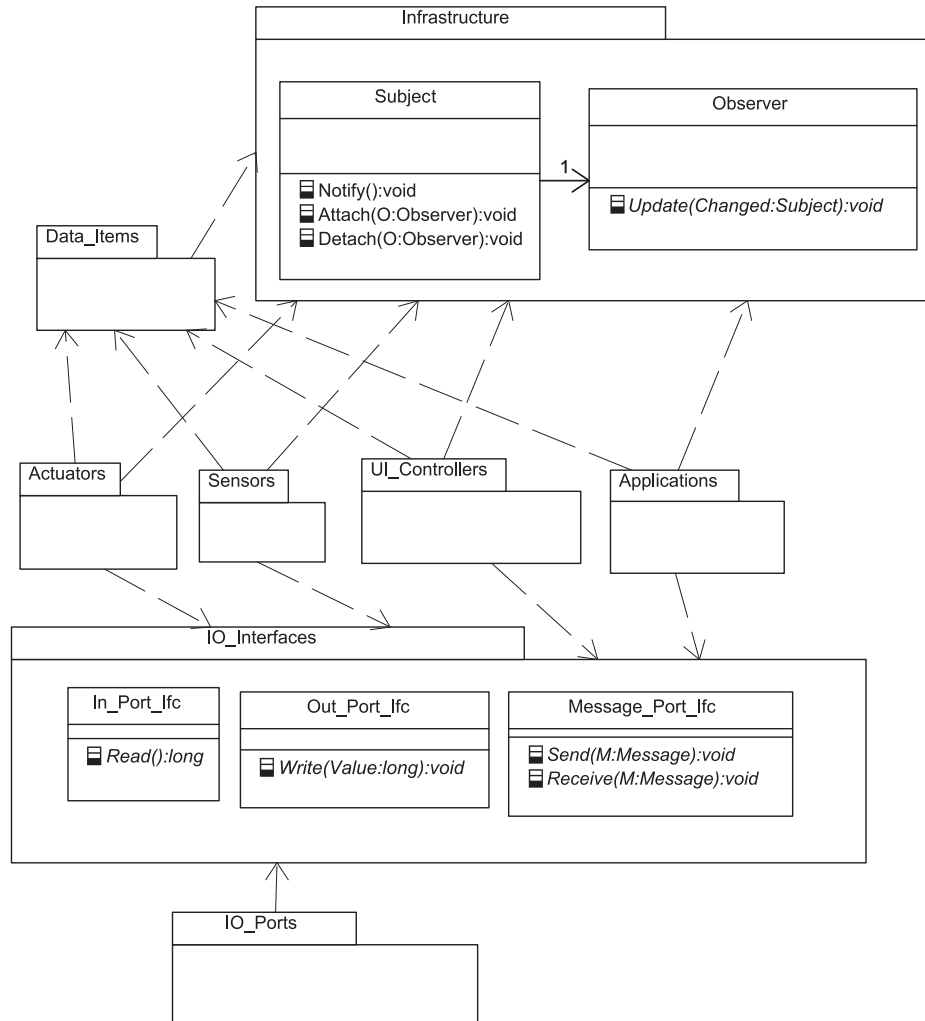


Fig. 6. High-level structure and framework in UML using the traditional way of modeling architecture.

- (6) *Sensor and actuator variability.* This is handled by using the same principles for sensors and actuators as for the applications described above. Sensors and actuators provide a mapping between physical sensors or actuators and Data Items

In the following two sections we first show how an architecture that captures these design principles would have been modeled in a traditional way and then using our modeling approach.

3.1 Traditional Way of Modeling the Architecture

Traditionally, the architect would have documented the architecture according to these design principles with a high-level structure and a support framework in UML, together with a set of rules expressed informally in text. This is exemplified with the UML model in Figure 6, accompanied by a set of textual architectural design rules, such as those below the figure.

The following are some rules for Data_Items.

- D1. A Data_Item is a class that reflects the state of the system or its environment that is needed by an application. The intention is that the set of Data_Items will be stable over time.
- D2. A Data_Item must inherit Infrastructure::Subject.
- D3. A Data_Item must be defined in the Data_Items package.
- D4. The only public operations of a Data_Item must be set and get operations to read and write data stored by the class.
- D5. A Data_Item may be a composition of Data_Items.
- D6. A set operation for a Data_Item must always end by calling its Notify operation.

And these are some rules for sensors.

- S1. A sensor is typically responsible for reading the value from a physical sensor scaling it and writing the value to a Data_Item. Some sensors may however not be connected to a physical sensor, but use indirect measures such as heating effect and time to estimate a value to write to the Data_Item.
- S2. A sensor must be defined in the Sensors package.
- S3. A sensor may inherit Infrastructure::Observer to be able to react to changes in Data_Items, for instance to activate or deactivate itself.
- S4. A sensor may only have associations to In_Port.Ifcs and Data_Items. These associations must only be navigable from the sensor.
- S5. A sensor may not have any public operations or attributes.
- S6. A sensor must periodically update its Data_Item.

In addition, there would be corresponding rules for Actuators, UI.Controllers, Applications, and IO.Ports.

3.2 Modeling the Architecture According to our Approach

In our approach, instead of using informal text, the architectural design rules are modeled in UML. Since UML (and any other OO language) is well suited to define structural relationships (such as, for instance, “*every country has one capital city*” used in many introductory courses in OO), this can be done in a straightforward way for rules such as the ones in the previous section. Using this approach the architectural rules for Data_Items above can be modeled as in Figure 7. In the figure it is indicated how each rule is modeled by the Dx labels. For example, the rule D2, stating that a Data_Item will inherit Subject, is modeled by associating a Data_Item to one Subject with an association stereotyped with <<Generalization>>. A major principle is that nothing that is not explicitly allowed is forbidden, so a Data_Item may not have any association other than compositions to other Data_Items, and may not inherit anything except a subject class.

We call the model where we model the architectural rules the *architectural rules model*. It is important to realize that the classes in this model are at the meta-level of the classes of the system model; that is, they define different kinds of classes and constrain them. For instance, the association between Subject and Observer in Figure 7 means that a Subject kind of class will have any number of Observer kinds of classes. An operation or an attribute in a class in this model means that a class of the corresponding kind in the system model must have an operation or attribute with the same characteristics as this operation or attribute. To allow for variations, wild cards can be used in attribute and operation definitions, where “@”



In the system model we use UML stereotypes to show the kind, corresponding to the classes in the architectural rules model, of an element. For example, in Figure 8 it can be seen that the class Subject has the stereotype <<Subject>>, meaning that it has to comply with the constraints defined by the Subject class in the architectural rules model. Figure 8 shows the high-level structure and framework classes modeled in the system model by the architect. The only difference to the one in Figure 6 modeled according to the traditional approach, is the stereotypes attached to the packages and classes. Since the architect models the high-level structure of the system, the rules restrict the developers as to which stereotypes to use in which package. For instance, in a package with stereotype <<Data.Items>, all classes must have the stereotype <<Data.Item>>. Figure 9 shows a number of Data.Items modeled in the system model following the rules of the architectural rules model in Figure 7. Any violations to the rules are automatically detected and reported by the tool built as part of our case study, described in Section 6.

ACM Transactions on Software Engineering and Methodology, Vol. 21, No. 2, Article 10, Publication date: March 2012.

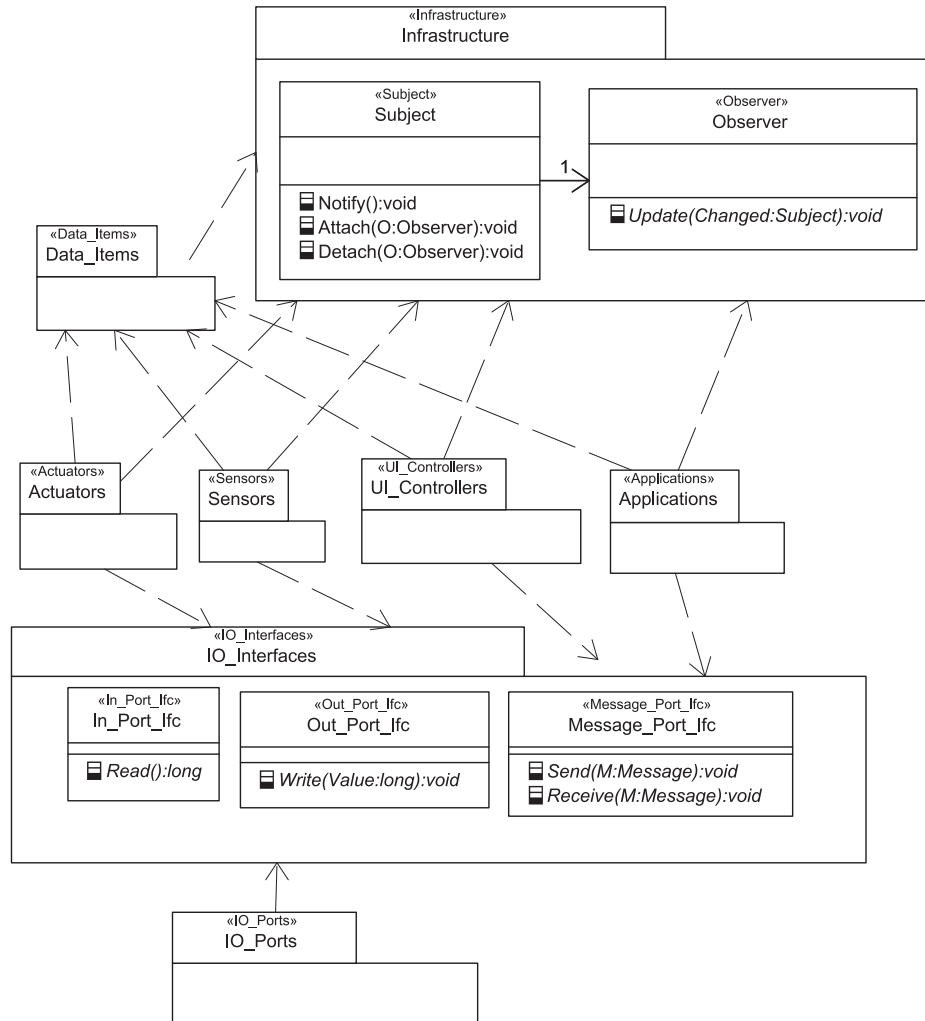


Fig. 8. High-level structure of the system model in UML using our approach.

4. RESEARCH APPROACH

There were three objectives of this research motivated by the lack of a satisfactory solution on how to model architectural design rules and the need for automating the enforcement of these in a practical situation, as discussed in Section 2:

- (1) Define an approach for modeling architectural design rules in UML.
- (2) Verify that the approach is stringent enough to be automated.
- (3) Demonstrate that the approach is applicable to a real development project.

To achieve the first objective, that of defining the approach, a systematic literature review presented in Section 2 was performed. The approach adopted was based on the findings of the literature review, and was refined based on the activities undertaken to achieve the second and third objectives above.

The second objective, to verify that the approach could be automated, was addressed by developing a tool to automatically check that a system fulfilled the architectural

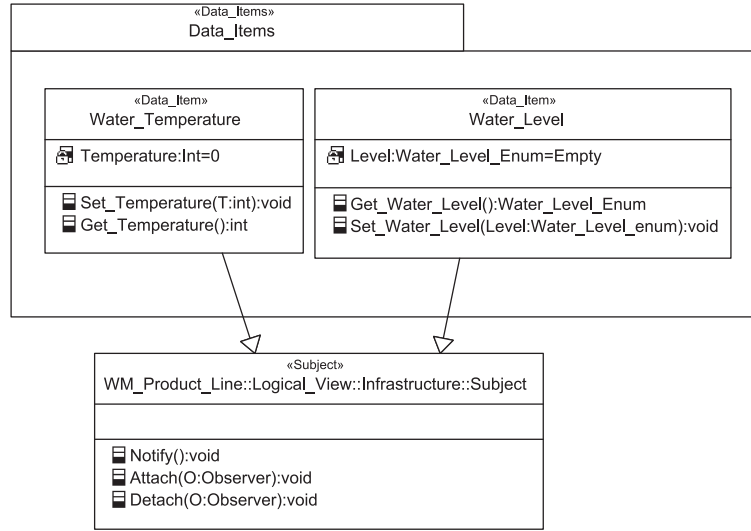


Fig. 9. Part of the detailed design for the Data_Items package in the system model.

rules specified according to the approach. In addition, a MOFScript transformation that transforms architectural rules defined according to our approach to OCL constraints in an architectural rules profile has been defined. This MOFScript transformation can be found in Appendix A.

To achieve the third objective, that of demonstrating the applicability of the approach, the architectural design rules from an existing, previously developed system, were modeled according to the approach. The goal of this activity was to establish the degree to which it enabled modeling of the architectural design rules. Specifically, we searched for answers to the following two questions.

- (1) To what extent could the specified rules be modeled?
- (2) Were there certain kinds of rules that could not be modeled and if not, why not?

The system was selected based on the following criteria.

- (1) The system had to have been developed using MDD.
- (2) The system had to be an existing real system of significant size and with a sufficient functionality to make it generally representative as a real-world embedded system.
- (3) The architecture, including the architectural design rules, had to be documented to a level where it could be interpreted by the research team.
- (4) The research team had to have good access to people who had first-hand knowledge of the architecture, to be able to see beyond the documentation and to be able to resolve any ambiguities.

The selected system, fulfilling these criteria, was a software platform for digital TV set-top boxes for the DVB¹ standard. The system had been developed by a project that had been studied in an earlier case study by the research team, reported in Mattsson et al. [2009], which meant that the team had good insight into the case. It was developed using the modeling tool Rhapsody (version 4.x) from Telelogic [Telelogic Rhapsody modeling], with all code generated from UML models in the tool, using C++

¹Digital Video Broadcasting, <http://www.dvb.org>.

as the action code language. The size of the software platform was approximately 350,000 eLOC in C++ and the effort to develop it was about 100 person years over a 24 month period. The architecture was documented partly in the system model and partly in one manually written document. The system model contained a high-level package structure and a framework of classes supporting the architectural design rules. The document contained the architectural design rules. Finally, the researchers had first-hand knowledge of the architecture, since the primary author of this article was the technical manager of the project, responsible for work practices and tools. The architecture was, however, developed by two other persons acting as architects.

The study was conducted by a systematic walkthrough reviewing the rules from the architectural document in several iterations, gradually transforming them to modeling constructs according to our approach.

5. AN APPROACH TO MODELING ARCHITECTURAL DESIGN RULES

In this section we present the definition of the approach for modeling architectural design rules that was developed in response to our first research objective.

As motivated in Section 4, our approach is based on transforming design rules, modeled in an architectural rules model, using UML, to a UML profile, applied to the system model. The implication is that our approach to modeling architectural design rules can be reduced to a set of transformations from constructs in the architectural rules model to stereotypes with constraints in a UML profile, hereafter referred to as the architectural rules profile. Therefore our approach is defined using such a set of transformations. In this section we present these transformations in an informal descriptive way, a formal definition in the form of a MOFScript transformation to a UML profile and OCL constraints can be found in Appendix A1, available in the ACM Digital Library.

The transformations are divided into two subsets, a general, complete transformation set and an additional UML-specific transformation set. The first transformation set is general, in the sense that it is applicable to any meta-model modeled in MOF, and not only to UML models. It is also complete, in the sense that it allows us to constrain any construction of any modeling language defined in MOF. However, using only these general transformations, it is still hard to model certain types of architectural rules commonly needed for UML models—for example, rules restricting UML associations. To ease the modeling of such rules, the additional UML, specific set of rules is needed. This transformation set is, however, not complete, so the fundamental set is still needed for completeness.

All examples illustrating the transformations in this section are taken from the washing machine example introduced in Section 3.

5.1 General Transformations

This section defines a set of transformations from an architectural rules model to an architectural rules profile defining constraints for types of classes in a system model. The transformations are applicable to all MOF-based languages, not only UML. The definitions refer to the generic architectural rules model in Figure 10, where C1 and C2 are replaced by class names; M1 and M2 replaced by stereotypes; R1 and R2 are replaced by role names; SR1 and SR2 are replaced by stereotypes; and Mu1 and Mu2 are replaced by multiplicities. In the transformations which follow, these conventions are used.

—References to terms defined in the generic architecture model in Figure 10 are in italics.



Fig. 10. A generic architectural rules model used in the definition of the transformations.

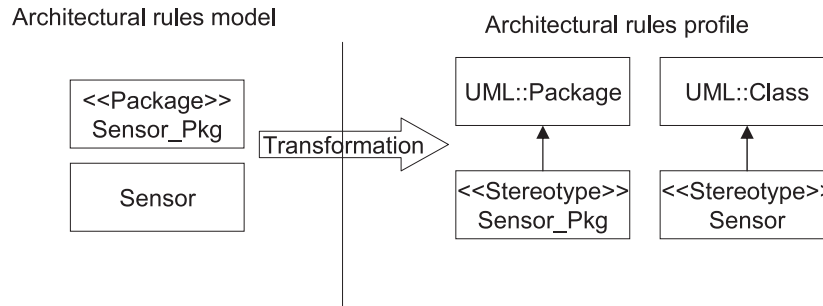


Fig. 11. Example of transformation according to transformation T1.

- The phrase “<<Cx>> element” shall be interpreted “element of stereotype Cx” or, if Mx equals “metaclass,” “element of metaclass Cx” (where x is 1 or 2).
- The term “metaclass” in the transformations refers to a metaclass of the modeling language that is constrained, for instance the metamodel for UML.
- The term “metamodel” in the transformations refers to the metamodel of the modeling language that is constrained, for instance the metamodel for UML.

The transformations are as follows.

- T1. A class named *C1* with the stereotype *M1* is transformed into a stereotype named *C1*, extending the metaclass *M1* unless transformation number T2, (below) applies. If *M1* is undefined, then “Class” is assumed; see Figure 11 for an example.
- T2. If *M1* equals “metaclass” then *C1* represents the class *C1* in the language metamodel (i.e., the UML metamodel) and is not transformed into anything in the profile. This can be used to specify constraints in other stereotypes in respect to these meta-classes; see Figure 12 for an example.
- T3. If *SR2* is the role in the language metamodel on the far end of an association from the metaclass of *C1* to the metaclass of *C2* then the multiplicity of *R2* for a << C2 >> element shall be constrained to *Mu2* in stereotype << C1 >>.

An example is shown in Figure 12, where a <<Sensor>> class is constrained to only have one <<SamplingPeriod>> attribute and no other attributes.

It is allowed to have several association ends matching the same meta-model association end. In that case the multiplicity of the end with the most narrow type scope is applied for a certain << C2 >> element. In the example in Figure 12 the multiplicity is “1” for an attribute with the stereotype <<SamplingPeriod>>, since this multiplicity is only applicable to attributes with stereotype <<SamplingPeriod>> and the multiplicity of 0 is applicable to all attributes.

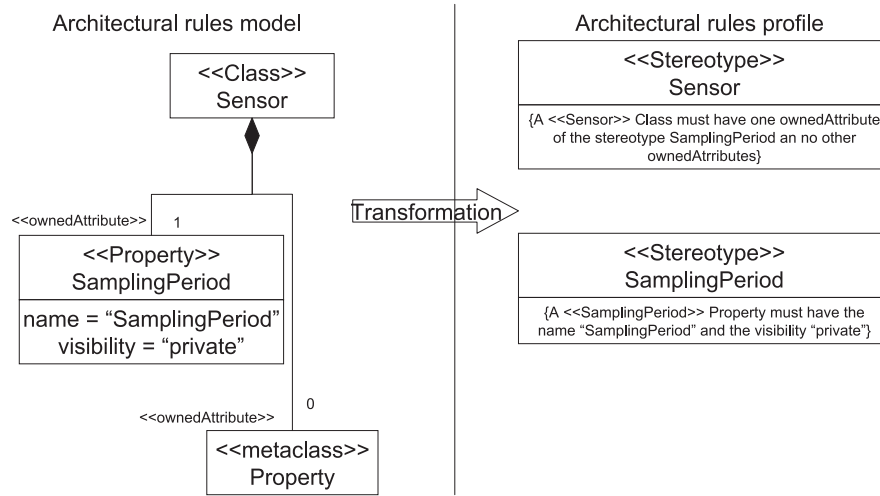


Fig. 12. An example of using transformation T2, T3, and T4.

T4. If the name of an attribute A matches the name of an attribute of class $M1$ in the meta-model, then it is transformed into a constraint on that attribute on allowed values. The value of the attribute is constrained to match a regular expression specified as the default value of the attribute.

An example is shown for the attribute name in Figure 12 where the name of the `<<SamplingPeriod>>` attribute is constrained to be “SamplingPeriod.”

T5. If no match is found for A , then A is transformed into an attribute A of the stereotype (tag-definition), thus defining a tagged value to be set in the model element where the stereotype is applied.

T6. Any OCL constraint in the context of a class $C1$ is copied into the architectural rules profile with the context of stereotype $C1$. This means that the constraints is written the same way as when defining stereotypes directly in the profile.

Even though OCL expressions, as discussed in Section 3, are not suitable for modeling architectural design rules in general, there is a need for them to express, for instance, constraints on combinations of rules. For example, if we would like to specify that a `<<Sensor>>` class has either a `<<Sample>>` operation or a `<<Trig>>` operation, it could be done like this:

```
context Sensor
inv: self.base_Class.ownedOperation.extension_Sample.size()=1 xor
self.base_Class.ownedOperation.extension_Trig.size()=1
```

T7. A generalization relationship from a class $C3$ to a class $C1$ in the architectural rules model is transformed to a generalization from stereotype `<<C3>>` to stereotype `<<C1>>` in the architectural rules profile as exemplified in Figure 13. This means that stereotype `<<C3>>` inherits all constraints from stereotype `<<C1>>` and that a `<<C3>>` class is also to be regarded as a `<<C1>>` class.

This set of transformations is general and complete in the sense discussed in what follows.

—*The transformation set is general.* These transformations allow us to use a subset of UML to constrain the usage of any modeling language defined in MOF, since




Fig. 13. An example of using transformation T7.

the transformations only assume that the modeling language is defined using MOF and do not assume anything about the content in the meta-model (i.e., the UML metamodel).

- *The transformation set is complete.* A model is an instance of its metamodel, which means that any model element is an instance of a class in the metamodel. The only things that may vary between two models of the same metamodel defined in MOF is the number of instances of each metaclass, the values and multiplicities of the metaclass attributes, and the links between the instances. Since these transformations allow us to constrain allowed values and multiplicities for attributes and constrain the types and multiplicities of associations, the set of transformations is complete in the sense that it allows us to constrain anything that can vary between different models of a certain meta-model defined in MOF.

By using only these transformations it is, still too complex to model constraints on some common UML constructs such as associations, attributes, operations, and state machines. For example, Figure 14 shows how the simple example rule S4 introduced in Section 3 is modeled according to these transformations.

To overcome this problem we have defined a set of additional transformations that makes it considerably simpler to specify certain constraints on UML models, common within the embedded software domain; they are described in the next section.

5.2 Additional UML-Specific Transformations

This section defines a set of transformations, in addition to the general ones defined in the previous section. The purpose of these transformations is to make it simpler to describe frequently needed architectural rules on UML models that are hard to describe using only the general transformations. These transformations override the general transformations in cases where both a general and an additional UML-specific transformation apply. The definitions refer to the generic architectural rules model in Figure 15. In the definitions, the following conventions are used.

- References to terms defined in the generic architecture model in Figure 15 are in *italics*.
- The phrase “<<Cx>> element” where x is 1 or 2 should be interpreted as “element of stereotype Cx” or, if Mx equals “metaclass,” “element of metaclass Cx”.

Constraints on stereotype C1 is defined according to the following.

- T8. If M1 equals “Package” and aggregation of R2 is “composite”: A << C1 >> package is constrained to contain Mu2 number of << C2 >> elements. The visibility of these elements must be the visibility of Mu2. Also, a << C1 >>

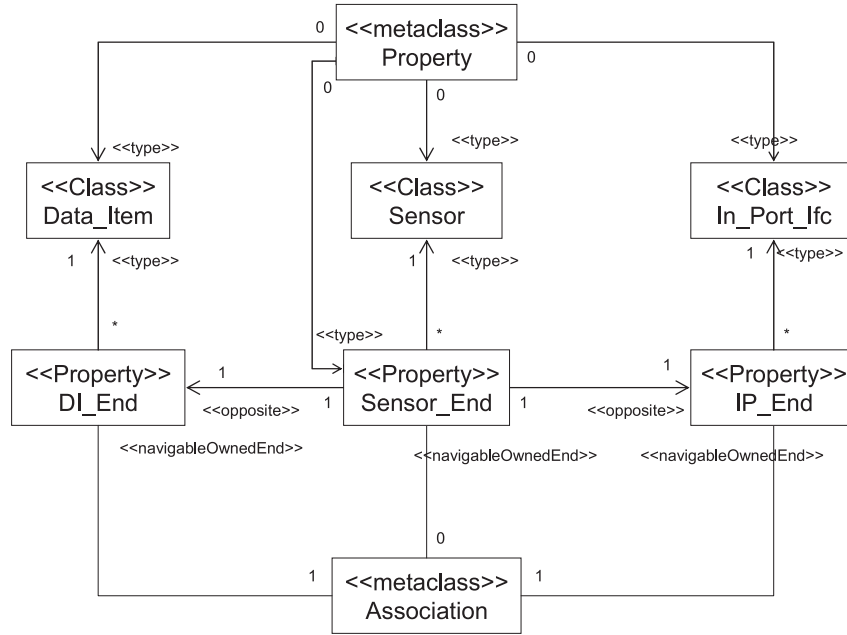


Fig. 14. Capturing an association constraint in a meta-model using the general transformations.




Fig. 15. A generic architectural rules model used in the definitions of the transformations.

Architectural rules model

Architectural rules profile




Fig. 16. Example of rules on package containment.

package is not allowed to have any packagedElements unless explicitly allowed in the model. This transformation makes it easy to model rules on package containment. An example is shown in Figure 16.

- T9. << C1 >> elements are only allowed to have the associations, dependencies, generalizations, and realizations explicitly allowed.
- T10. If *MA* equals “Association”: A << C1 >> element must be associated with *Mu2* number of << C2 >> elements. The association ends must have the same navigability, aggregation (none, shared, or composite) and visibility as *R1* and

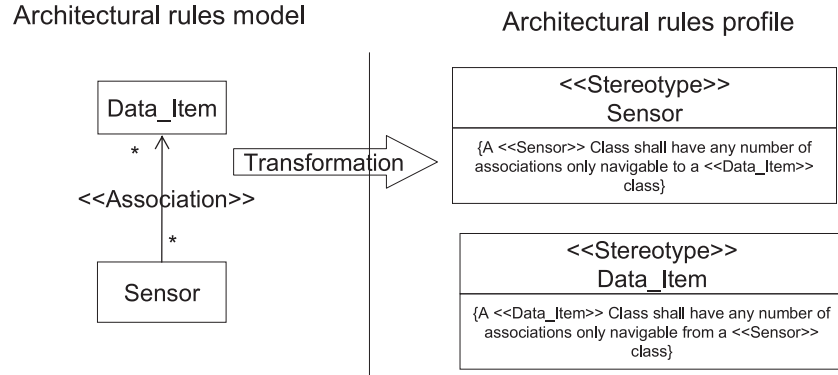


Fig. 17. Example of rules on associations.

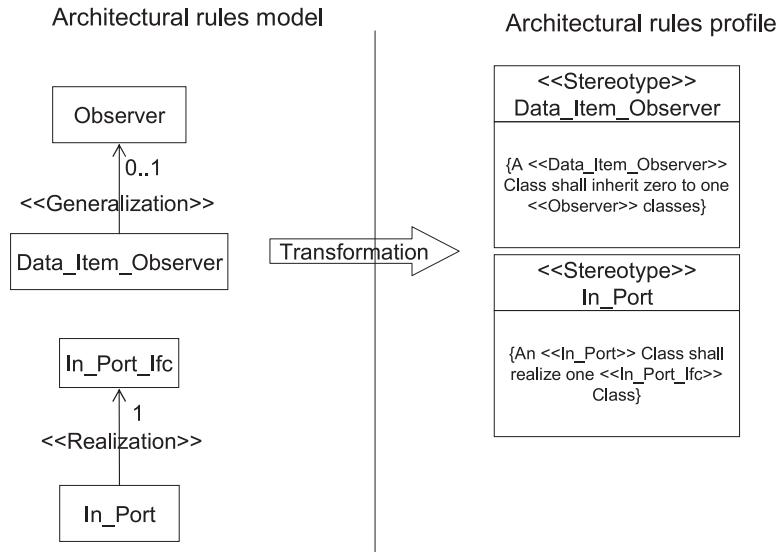


Fig. 18. Example of rules on generalizations, dependencies, and realizations.

R2. The association ends must also have qualifiers according to the qualifiers of *R1* and *R2*. The name and type of these must be according to the transformations for attributes specified in the following. The association must have the stereotypes *S1* to *Sn*. This transformation makes it easy to formulate rules on associations; as, for instance, the example rule introduced in Section 4 can now be modeled as shown in Figure 17. Contrast this with the model in Figure 14 to see the difference from modeling using only the fundamental transformations.

- T11. If *MA* equals “Dependency,” “Generalization,” or “Realization,” and the association is only navigable from *C1* to *C2*: A in << *C1* >> element must have a relationship according to *MA* to *Mu2* number of << *C2* >> elements with stereotypes *S1* to *Sn*. Examples of these kinds of transformation are shown in Figure 18.
- T12. If there are attributes *A* of *C1* that starts with \$ then, the following hold.
- all parts of the definition of an attribute of a << *C1* >> class must match the corresponding part of an *A*, where the wild card characters “@” and “%” in any part of the definition of *A* can be replaced with any character sequence.




Fig. 19. Example of rules on attributes.

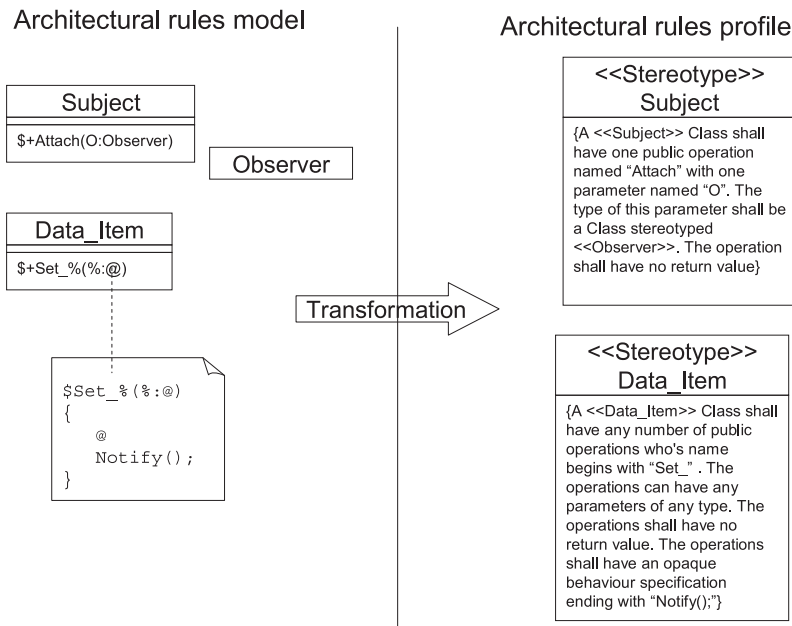


Fig. 20. Example of rules on operations.

Parts of A not specified (as for instance, default value for Sampling_Period in Figure 19) are unconstrained.

- (b) Every A must be matched by one attribute in a << C1 >> class. An exception to this is if the name of A contains the wild card character "%"; in this case any number of matches (including zero) is allowed.
- (c) if the name of a type of A is identical to the name of a class C in the architectural rules model, then the type of a matching attribute must be a <<C>> element.

This transformation is exemplified in Figure 19.

T13. If there are operations O of C1 that start with \$ then the following hold.

- (a) All parts of the definition of an operation of a << C1 >> class must match the corresponding part of an O, where, for each part of the definition, the wild card characters "@" and "%" can be replaced with any character sequence. Parts of O not specified (as, for instance, parameter directions for operations in Figure 20) are unconstrained.

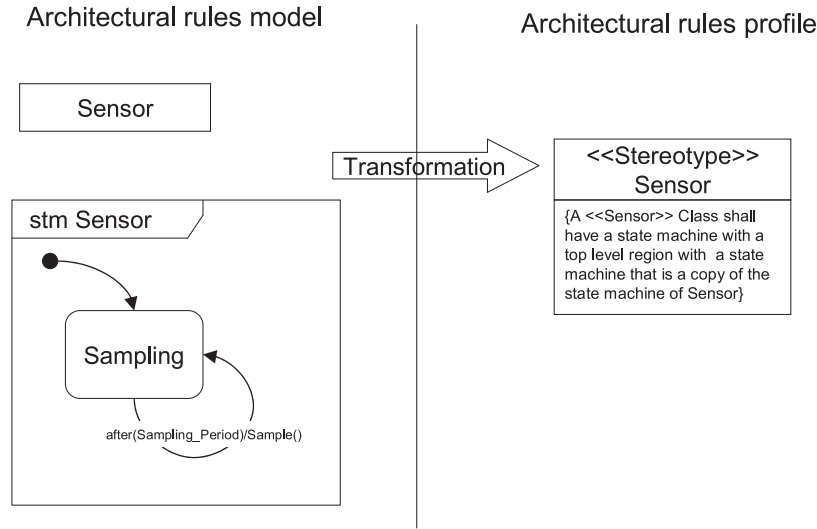


Fig. 21. Example on rules on state machines.

- (b) This requirement holds for all parts of the definition of *O* defined in the UML meta-model, such as, for instance, opaque behavior specified for the operation.
- (c) The character “%” in a parameter name means that the definition of this parameter can be repeated any number of times, including zero. In these parameter definitions “%” can be replaced with any character sequence.
- (d) If the name of the type of *O* or a parameter of *O* is identical to the name of a class *B* in the architectural rules model, then the type of matching operations or parameters in the << *C1* >> class must be of a <<*B*>> Class.
- (e) All *O* must be matched by one operation in a << *C1* >> class. An exception to this is if the name of *O* contains the wild card character “%”; in this case any number of matches (including zero) is allowed.

This transformation is exemplified in Figure 20.

- T14. If *C1* has a state machine, then a << *C1* >> class must have a state machine where, for each region in *C1*, there must be an identical region in the <<*C1*>> class. The wild card character “@” may be used in the transition definitions in *C1* and will then be matched with any text string in the corresponding transition in the state machine of a << *C1* >> class. It is allowed to have additional regions in the state machine of a << *C1* >> class.

This transformation is exemplified in Figure 21. In this example a <<Sensor>> class is constrained to have a top region exactly matching the state machine for *Sensor* in the architectural rules model, which in effect forces *Sensor* classes to call the operation *Sample()* periodically with the period specified by the attribute *Sampling_Period*. A <<Sensor>> class may have additional behavior specified in parallel regions to the one specified in *Sensor*.

These additional UML-specific sets of transformations make it easy to specify constraints on, for instance, how different kinds of classes may be associated. To illustrate, let us revisit the previously used example in Section 2.2.



Fig. 22. Capturing the architectural design rule using the specialized additional transformations.

A sensor may only have associations to In_Port_Ifc and Data_Items. These associations shall only be navigable from the sensor.

This rule may now be modeled according to Figure 22,² which is very close to the simple (but only indicative) model in Figure 5, and significantly less complex than the model in Figure 14, where only the general transformations were used.

These additional transformations also make it simple to specify other common constraints, such as on package structure and on interfaces and the behavior of classes. This is further illustrated in Section 7.

6. AUTOMATING THE APPROACH

In this section we present the tool for automating enforcement of architectural design rules that was developed in response to our second research objective (which was to verify that the approach was stringent enough to be automated).

To provide a proof-of-concept of the feasibility of automating the approach, a tool was built making it possible to automatically check that a system conforms to rules modeled according to the approach.

There were several options when considering tool support for the approach.

- The rules could be enforced as a separate test, reporting violations.
- The rules could be continuously enforced during modeling, giving the possibility of guiding the developer during development and, if desired, preventing the modeler from breaking the rules.
- In both cases the modeled rules could either first be transformed into OCL constraints in a UML profile, which would then be enforced, or they could be directly enforced on the system model.

An important thing to consider was how to make it as easy as possible for an organization to adopt the new method and tool. For an organization that is already using modeling tools it would be a big advantage if they did not have to change their modeling tools. New tools would incur cost in purchase, training, and transferring models to the new tools. In our case the organization was using the Rhapsody modeling tool. This was also the tool that had been used in building the system for which we intended to re-model the architectural design rules. Hence we needed a tool that could take Rhapsody models both for the architectural design rules and for the system models. Considering this we built the tool as a stand-alone checker to the Rhapsody tool, validating the system model directly against the modeled rules for the following reasons.

- (1) To make a plug-in to the modeling tool that continuously checks the model, harder to make a stand-alone checker, would be harder to move to another tool, and would risk increasing the response time when modeling.

²Actually, in our example introduced in Section 3, Sensor would instead inherit the rule so as to be allowed to have associations to <<Data_Item>> from Data_Item.Observer shown in Figure 7.

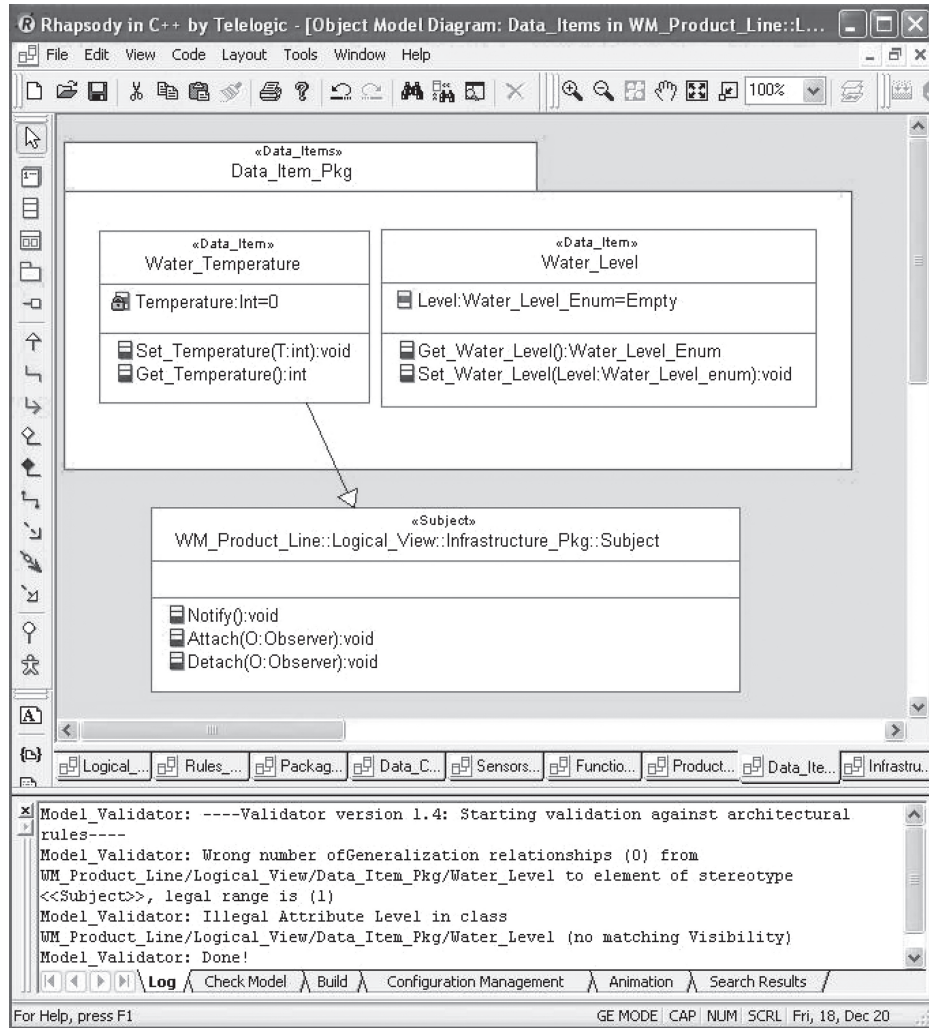


Fig. 23. Screenshot of tool.

- (2) Although open-source stand-alone tools for OCL checking are available, making an OCL generator and integrating an OCL checker in Rhapsody was considered at least as hard as our current approach, and in addition would increase the risk, since we would be relying on another tool.

The tool is built in C++, and is currently limited to reading Rhapsody [Telelogic Rhapsody modeling] models, both for architectural rules and for the system model. The tool is designed so that there are no dependencies to the model reader component from any other parts in the tool. This makes it a relatively small task to adapt the tool to another modeling tool. The total effort to build the tool was approximately 200 hours; the estimated effort to build another model reader is about 40 hours. A screenshot of the tool is shown in Figure 23 where the output from the validator is shown in the text window in the bottom. The text refers to the violations in the Water_Level class in respect to the architectural rules in Figure 7.

Table I. A Subset of the Full Table (which had mappings between all the original architectural design rules and resulting modeling constructs)

Id	Original rule (Quotation)	Modelreference	Used Transformations
3.2	“Functionality specific to a PAPI requirement shall be kept in this layer unless it is reusable for another PAPI or applicable to DVB standard. In this case it shall be placed in CMP or CMD.”	-	-
4.1	“All coupling between <i>arcComponents</i> shall be loose in the sense not statically linked”	Handled by only allowing associations from a component to an Interface that is realized by an <i>arcComponent</i> .	T9, T10, T11
4.2	“All associations between <i>arcComponents</i> shall be navigable from the client to the server (user to the resource)”	User/Resource association from <i>mComponent</i> to <i>mCompIfc</i>	T10
6.7	“In the case of a component locked to a specific <i>arcComponentUser</i> , it is the responsibility of the locker to allow only one thread at a time to access the component.”	This is ensured by the implementation of the enforced implementation of the operations of the <i>mLockableComponent</i> .	T13
8.1	“All locked components shall inherit the same base class, <i>arcLocked</i> .”	Generalization from <i>mLockableComponent</i> to <i>marcLockableComponent</i> , there is only one instance of <i>marcLockableComponent</i> allowed, in an <i>Architecture_Pkg</i> and finally there is only one <i>Architecture_Pkg</i> with only one <i>marcLockableComponent</i> class in the <i>Systemmodel</i> .	T8, T11
9.16	“Transmission events and exceptions initiated by a <i>Write()</i> shall be reported back to the <i>arcPortUser</i> via the <i>TxDone()</i> call.”	A <i>Write</i> operation is forced to always end with a call to <i>TxDone</i>	T13

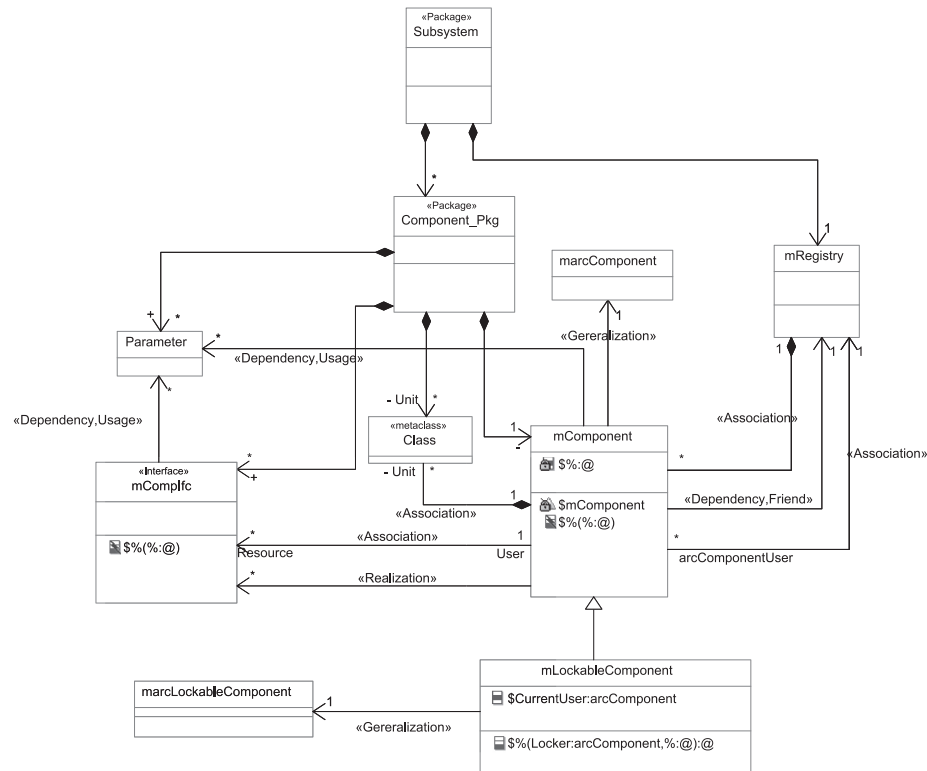
7. MODELING ARCHITECTURAL DESIGN RULES OF AN INDUSTRIAL-STRENGTH SYSTEM

In this section we present the findings of a case study performed in response to our third research objective (which was to demonstrate that the approach was applicable to a real development project).

To demonstrate the applicability of our approach to a real problem, we modeled the architectural design rules from an already developed system according to our set of transformations.³ The mapping between the original architectural rules and the models was documented in a text table. A part of this table is shown in Table I. The rules could be classified into three categories: structural, behavioral, and judgmental. Structural rules specified structural constraints such as rule 4.1, 4.2, 6.7, and 8.1 in the table. Behavioral rules specified constraints on behavior such as rule 9.16 in the table. Judgmental rules were rules where the developer had to exercise judgment to follow the rule; 3.2 is an example of such a rule in the table. There were 66 rules

³Note that in a real case the architectural rules model would be modeled as a natural part of the architectural design and not as a separate activity. Normally, there would not even be any textual expression of the rules, only the architectural rules model.

	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	T12	T13	T14
Usage frequency	N.A.	2	3	2	0	2	1	10	N.A.	13	22	6	17	5
Violations (%)	N.A.	3,37	4,46	2,92	0,00	3,19	1,55	15,38	N.A.	20,73	34,84	8,55	25,29	9,57



in total; eight of these could not be modeled. These rules were all judgmental, and therefore inherently impossible to formalize. The rules typically consisted of one or two sentences, where the sample rules in Table II are representative.

Both the architectural rules model and the architectural parts of the system model were captured in the Rhapsody modeling tool (version 7.2). Figure 24 to Figure 26 show parts of these models. Figure 25 shows the subsystems modeled as packages in the system model. This level of the system model is owned by the architects. The stereotypes of

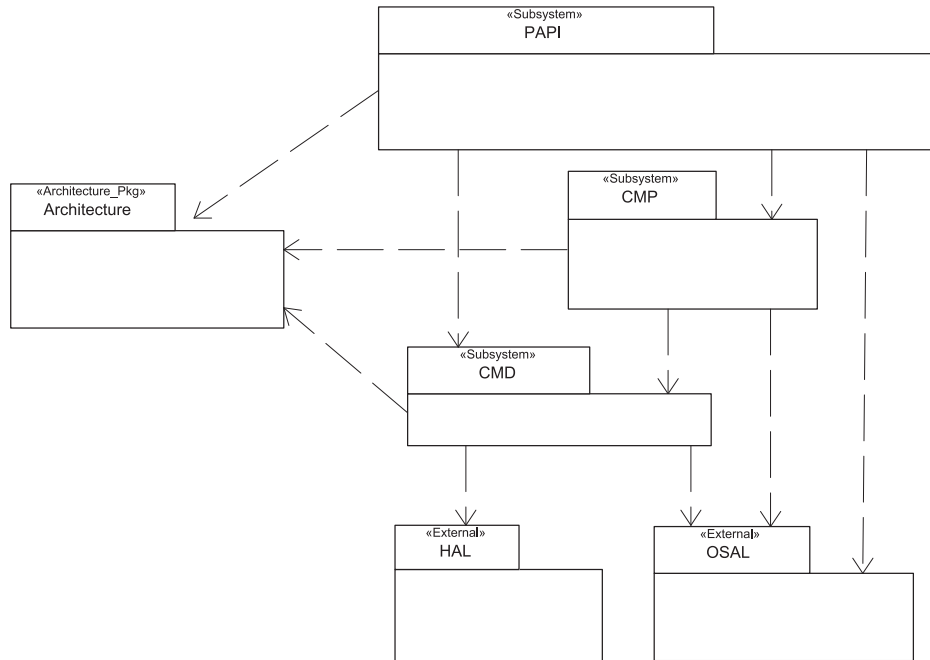


Fig. 25. Top level of the system model.

these packages are defined in the architectural rules model partly shown in Figure 24. In this model we can see, for instance, the architectural rules that a `<<Subsystem>>` package (that is, a package with the stereotype `<<Subsystem>>`) must contain a number of `<<Component_Pkg>>` packages and one `<<mRegistry>>` class. We can also see that a `<<Component_Pkg>>` must contain exactly one `<<mComponent>>` class that must inherit a `<<marcComponent>>` class (defined by the architects in the architecture package in the system model.). In Figure 26 an example of a small component in the system model is shown, following the architectural rules defined in the architectural rules model.

8. DISCUSSION AND CONCLUSIONS

Architectural design rules are an important part of the architecture and there are no adequate solutions in the current body of literature on how to model them. The inability to formalize the architectural design rules leads to a need for error-prone and time-consuming manual tasks to enforce them. The approach developed in this study addresses this problem by providing a technique for modeling architectural design rules in a way that is formal enough to allow automation. An important property of the approach is that the architectural design rules are modeled using UML at a high abstraction level, without requiring detailed knowledge of the UML meta-model. That the rules are modeled at an abstraction level close to that of the rule itself is required for the models to be easily understandable for architects and developers, an issue of paramount importance for the usability of the approach. The use of UML reduces the required investment in tools and training, since architects and developers benefit from previous knowledge in UML and are able to use their current UML tools for modeling; to provide automation only requires an additional tool that checks the system model against the architectural model according to our defined transformations. Our

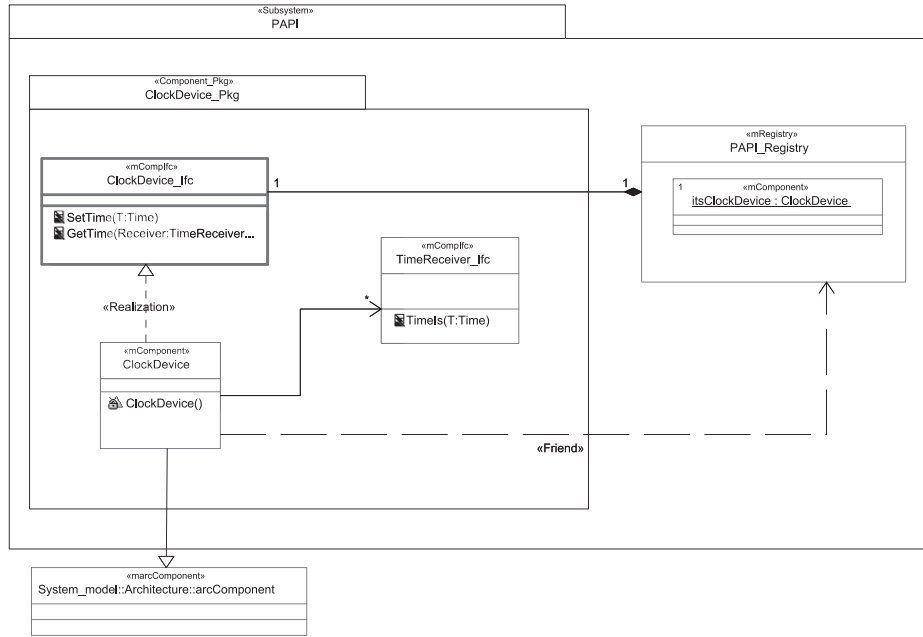


Fig. 26. The ClockDevice in the system model.

effort for building such a tool for the Rhapsody modeling tool using its COM API was approximately 200 man-hours, so this should be a relatively small task.

In applying our approach to modeling the architectural design rules of an industrial strength system, we found that of the original 66 rules only eight could not be modeled. This means that we would have relieved the architects of a large part of their enforcement effort; only 12% of the rules would have been left for manual enforcement. The rules that could not be modeled were all rules where the developer was supposed to exercise judgement, which made them inherently impossible to formalize. The following is a typical example of such a rule.

“Functionality specific to a PAPI requirement shall be kept in this layer unless it is reusable for another PAPI or applicable to the DVB standard. In this case it shall be placed in CMP or CMD.”

These are rules that need a lot of interaction between the developers and the architects in order to develop a common understanding of what the rules really mean. It is very important to get this right at the same time as it is impossible to finalize and formalize them at an early stage in the project. This is where the focus of the architects should be, and our approach gives the architects the time to do that. Other benefits are that modeling eliminates ambiguities and redundancy in the rules, which should make them easier to understand and give less room for erroneous interpretations.

Although the approach has only been tested on one system, two factors suggest that the results should, to a large extent, be transferable to other systems and organizations in the embedded software domain.

- (1) The defined transformations are based on raising the general modeling constructs of UML to the meta-model level, not on the specific needs of the system used for the test.

- (2) It is a real-world embedded system of significant size with functionality quite common in this domain.

There is a need for further research to study the implications when adopting the approach in other application domains. Factors to investigate include the ease with which architects, developers and other stakeholders can learn the approach and accommodate their working practices to it.

9. ELECTRONIC APPENDIX

The electronic appendix to this article is available in the ACM Digital Library.

REFERENCES

- AMERICA, P., ROMMES, E., AND OBBINK, H. 2004. Multi-view variation modeling for scenario analysis. In *Software Product-Family Engineering*. F. Vanderlinden Ed., Springer, Berlin, 44–65.
- BASS, L., CLEMENTS, P., AND KAZMAN, R. 2003. *Software Architecture in Practice*. Addison-Wesley, Reading, MA.
- BAYLEY, I. 2007. Formalising design patterns in predicate logic. In *Proceedings of the 5th IEEE International Conference on Software Engineering and Formal Methods*. IEEE, Los Alamitos, CA, 25–36.
- BENGTSOON, P. AND BOSCH, J. 1998. Scenario-based software architecture reengineering. In *Proceedings of the 5th International Conference on Software Engineering*. 308–317.
- BOSCH, J. 2000. *Design and Use of Software Architectures : Adopting and evolving a Product Line Approach*. Addison-Wesley, Reading, MA.
- BOSCH, J. AND MOLIN, P. 1999. Software architecture design: Evaluation and transformation. In *Proceedings of the IEEE Conference and Workshop on Engineering of Computer-Based Systems (ECBS'99)*. IEEE, Los Alamitos, CA, 4–10.
- BUSCHMANN, F. 1996. *Pattern-Oriented Software Architecture: A System of Patterns*. Wiley, New York.
- CONBOY, K. AND FITZGERALD, B. 2010. Method and developer characteristics for effective agile method tailoring: A study of expert opinion. *ACM Trans. Softw. Eng. Methodol.* 20, 1.
- EDEN, A. H. 2002. A theory of object-oriented design. *Inf. Syst. Frontiers* 4, 4, 379–391.
- FRANCE, R. B., KIM, D. K., SUDIPTO, G., AND SONG, E. 2004. A UML-based pattern specification technique. *IEEE Trans. Softw. Eng.* 30, 3, 193–206.
- FUENTES-FERNÁNDEZ, L. AND VALLECILLO-MORENO, M. 2004. An introduction to UML profiles. *Euro. J. Inform. Profess.* V, 2.
- GAMMA, E. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, MA.
- GREENFIELD, J. AND SHORT, K. 2004. *Software Factories: Assembling Applications with Patterns, Models, Frameworks, and Tools*. Wiley, New York.
- HOFMEISTER, C., NORD, R. AND SONI, D. 2000. *Applied Software Architecture*. Addison-Wesley, Reading, MA.
- JANSEN, A. AND BOSCH, J. 2005. Software architecture as a set of architectural design decisions. In *Proceedings of the 5th Working IEEE/IFIP Conference on Software Architecture (WICSA'05)*. 109–120.
- JANSEN, A., VAN DER VEN, J., AVGERIOU, P., AND HAMMER, D. K. 2007. Tool support for architectural decisions. In *Proceedings of the 6th Working IEEE/IFIP Conference on Software Architecture (WICSA'07)*. 44–53.
- KARSAI, G., SZTIPANOVITS, J., LEDECZI, A., AND BAPTY, T. 2003. Model-integrated development of embedded software. *Proc. IEEE* 91, 1, 145–164.
- KRUCHTEN, P. 2004a. An ontology of architectural design decisions in software intensive systems. In *Proceedings of the 2nd Groningen Workshop on Software Variability*. 54–61.
- KRUCHTEN, P. 2004b. *The Rational Unified Process: An Introduction*. Addison-Wesley, Reading, MA.
- KRUCHTEN, P., LAGO, P., AND VAN VLIET, H. 2006. Building up and reasoning about architectural knowledge. In *Quality of Software Architectures*. Springer, Berlin, 43–58.
- KRUCHTEN, P. B. 1995. The 4+1 view model of architecture. *IEEE Softw.* 12, 6, 42–50.
- LAUDER, A. AND KENT, S. 1998. Precise visual specification of design patterns. In *Proceedings of the 12th European Conference on Object-Oriented Programming*. Springer, Berlin.
- LEVY, Y. AND ELLIS, T. J. 2006. A systems approach to conduct an effective literature review in support of information systems research. *Inf. Sci. J.* 9, 181–212.

- MAK, J. K. H., CHOY, C. S. T. AND LUN, D. P. K. 2004. Precise modeling of design patterns in UML. In *Proceedings of the 26th International Conference on Software Engineering*. 252–261.
- MATTSSON, A., LUNDELL, B., LINGS, B., AND FITZGERALD, B. 2009. Linking model-driven development and software architecture: A case study. *IEEE Trans. Softw. Eng.* 35, 1, 83–93.
- MEDVIDOVIC, N., DASHOFY, E. M., AND TAYLOR, R. N. 2007. Moving architectural description from under the technology lamppost. *Inf. Softw. Technol.* 49, 1, 12–31.
- MEDVIDOVIC, N., ROSENBLUM, D. S., REDMILES, D. F., AND ROBBINS JASON, E. 2002. Modeling software architectures in the Unified Modeling Language. *ACM Trans. Softw. Eng. Methodol.* 11, 1, 2–57.
- MEDVIDOVIC, N. AND TAYLOR, R. N. 2000. A classification and comparison framework for software architecture description languages. *IEEE Trans. Softw. Eng.* 26, 1, 70–93.
- MIKKONEN, T. 1998. Formalizing design patterns. In *Proceedings of the International Conference on Software Engineering Forging New Links*. 115–124.
- OMG. 2003. MDA Guide version 1.0.1 OMG.
- OMG. 2003. UML 2.0 OCL specification.
- OMG. 2006. Meta Object Facility (MOF) core specification.
- OMG. 2009. Unified modeling language: Superstructure.
- PAHL, C., GIESECKE, S., AND HASSELBRING, W. 2007. An ontology-based approach for modeling architectural styles. In *Software Architecture*. 60–75.
- PERRY, D. E. AND WOLF, A. L. 1992. Foundations for the study of software architecture. *SIGSOFT Softw. Eng. Notes* 17, 4, 40–52.
- RAN, A. 2000. ARES conceptual framework for software architecture. In *Software Architecture for Product Families Principles and Practice*, M. Jazayeri, et al. Eds. Addison-Wesley, Reading, MA, 1–29.
- SCHMIDT, D. C. 2006. Model-driven engineering. *IEEE Computer* 39, 2, 25–31.
- SHAW, M., DELINE, R., KLEIN, D. V., ROSS, T. L., YOUNG, D. M., AND ZELESNIK, G. 1995. Abstractions for software architecture and tools to support them. *IEEE Trans. Softw. Eng.* 21, 4, 314–335.
- SHAW, M. AND GARLAN, D. 1996. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, Upper Saddle River, NJ.
- SONI, D., NORD, R. L., AND HOFMEISTER, C. 1995. Software architecture in industrial applications. In *Proceedings of the IEEE 17th International Conference on Software Engineering*. IEEE, Los Alamitos, CA, 196–207.
- TELELOGIC RHAPSODY MODELING. <http://www.telelogic.com/products/rhapsody/>.
- TOLVANEN, J. P. AND KELLY, S. 2005. Defining domain-specific modeling languages to automate product derivation: Collected experiences. In *Proceedings of the 9th International Conference on Software Product Lines (SPLC'05)*. Lecture Notes in Computer Science, Springer, Berlin, 198–209.
- TYREE, J. AND AKERMAN, A. 2005. Architecture decisions: Demystifying architecture. *IEEE Softw.* 22, 2, 19–27.
- VAN DER LINDEN, F., BOSCH, J., KAMSTIES, E., KANSALA, K., AND OBBINK, H. 2004. Software product family evaluation. In *Proceedings of the 3rd International Conference on Software Product Lines (SPLC'04)*. Lecture Notes in Computer Science, vol. 3154, Springer, Berlin, 110–129.
- WOJCIK, R., BACHMANN, F., BASS, L., CLEMENTS, P., MERSON, P., NORD, R. L., AND WOOD, B. 2006. Attribute-driven design (ADD), Version 2.0., Software Engineering Institute, Carnegie Mellon University.
- ZDUN, U. AND AVGERIOU, P. 2005. Modeling architectural patterns using architectural primitives. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages, and Applications*. ACM, New York.

Received July 2009; revised June 2010; accepted September 2010