

Flexible and Scalable Consistency Checking on Product Line Variability Models

Michael Vierhauser
Christian Doppler Lab for
Autom. Softw. Eng.
Linz, Austria
vierhauser@ase.jku.at

Paul Grünbacher
Systems Eng. and Automation
Johannes Kepler University
Linz, Austria
paul.gruenbacher@jku.at

Alexander Egyed
Systems Eng. and Automation
Johannes Kepler University
Linz, Austria
alexander.egyed@jku.at

Rick Rabiser
Christian Doppler Lab for
Autom. Softw. Eng.
Linz, Austria
rabiser@ase.jku.at

Wolfgang Heider
Christian Doppler Lab for
Autom. Softw. Eng.
Linz, Austria
heider@ase.jku.at

ABSTRACT

The complexity of product line variability models makes it hard to maintain their consistency over time regardless of the modeling approach used. Engineers thus need support for detecting and resolving inconsistencies. We describe experiences of applying a tool-supported approach for incremental consistency checking on variability models. Our approach significantly improves the overall performance and scalability compared to batch-oriented techniques and allows providing immediate feedback to modelers. It is extensible as new consistency constraints can easily be added. Furthermore, the approach is flexible as it is not limited to variability models and it also checks the consistency of the models with the underlying code base of the product line. We report the results of a thorough evaluation based on real-world product line models and discuss lessons learned.

Categories and Subject Descriptors

D.2.13 [Software Engineering]: Reusable Software—*Domain engineering, Reuse models*; D.2.4 [Software Engineering]: Software/Program Verification; D.2.6 [Software Engineering]: Programming Environments; D.2.8 [Software Engineering]: Metrics—*Performance measures*

General Terms

Experimentation, Measurement, Performance, Verification.

Keywords

Software product lines, variability models, model consistency, incremental consistency checking, performance, memory consumption, lessons learned.

published at
ASE'10, September 20–24, 2010, Antwerp, Belgium.

1. INTRODUCTION

Product line variability models are inherently complex. Regardless of whether feature-oriented [25], decision-oriented [10], or orthogonal [2] variability models are used, their size represents a major challenge in real-world product lines as they can easily contain thousands of elements with diverse and often complex dependencies. In the collaboration with our industry partner Siemens VAI – the world's leading company in engineering and plant-building for the iron, steel, and aluminum industries – we learned that engineers in practice face big challenges when maintaining the consistency of variability models. The models are subject to continuous evolution [16] and need to co-evolve with the actual system they represent to maintain consistency. The consistency of the models is also essential for deriving correct products. In product line engineering consistency constraints range from simple rules and well-formedness criteria (e.g., there must be no cycles in model element dependencies) to more sophisticated checks (e.g., each component in the variability model must exist in the product line code base and vice versa). Engineers need support for detecting and keeping track of such inconsistencies when modifying the product line's models or code base. More specifically, there are several challenges for consistency checking in model-based product lines:

C1. Dealing with variability and multiple levels of abstraction. Consistency needs to be checked within and between different spaces and various types of inconsistencies need to be considered. For instance, consistency checking needs to address the problem space describing variability using concepts from the problem domain, the solution space representing the variable system from a technical point of view, and the code base of the product line. While problem and solution space are typically represented by models, checking the code base requires different mechanisms. However, existing checkers are often limited to either models or to code and do not check across boundaries. Also, contrary to models describing single systems there can be optional or alternative model elements in product lines. Additional inconsistencies can occur when an element is optional or alternative.

C2. Addressing the scale and complexity of product lines. Several consistency checking mechanisms have been reported in the literature and have been applied to various types of models [7, 12, 34]. They are typically only capable of checking the consistency of entire models in a batch-oriented manner meaning that the relevant consistency constraints can be evaluated at certain points in time only (e.g., when saving a model) due to the complexity of

the models. For example, when using a batch-oriented approach in our DOPLER product line tools [18] we ran into significant performance problems when working with real-world variability models of our industry partner which contain thousands of model elements with non-trivial dependencies and mappings. The bad performance meant that feedback could not be provided to modelers immediately as launching the batch checker after each change to the model turned out to be infeasible. Special emphasis is thus necessary to design consistency checking algorithms such that immediate user feedback is possible.

C3. Providing support for domain-specific checks. Another drawback of many available consistency checkers is their lack of extensibility. It is complicated for modelers to add new consistency constraints and to remove or modify existing ones. We also faced this challenge with our batch-oriented checker as the various consistency checks were hard to identify and separate in the code.

To address these challenges, we started exploring the use of an existing incremental consistency checker [19] for product line variability modeling. Our goals were to extend the scope of consistency checking to model-to-code consistency and to improve performance to allow immediate feedback to modelers. We also aimed at checking whether the approach scales to industrial models of different sizes without performance and memory consumption problems. Another goal was to improve the extensibility by making it easy to add, remove, or modify consistency constraints.

In this experience paper we report results and lessons learned when applying existing consistency checking methods to a product line setting. The tool features we developed as part of this research are briefly summarized in a short workshop paper [35]. Here, we give an overview about the technical approach, discuss lessons learned during its implementation and report results of a thorough evaluation. We first introduce different types of possible inconsistencies on and among different layers of product line models. To illustrate the practical challenges, we describe consistency constraints we developed for the component-based product line of our industry partner. We then present our solution to incremental consistency checking on variability models and describe its integration in an existing product line engineering tool suite. We provide an evaluation with regard to performance and scalability and discuss lessons learned of applying the approach to product line models of our industry partner. We discuss related work and conclude the paper with an outlook on future work.

2. CONSISTENCY CHECKING IN MODEL-BASED PRODUCT LINES

Product line models cover both the problem space and the solution space [30] as shown in Figure 1. A problem space model describes the variability of a system using concepts from the problem domain, e.g., in the language of the users in a specific domain. A solution space model provides a representation of the variable system from a technical point of view, e.g., it describes optional or alternative artifacts and their dependencies [13]. Consistency needs to be maintained between these modeling spaces but also within each space. However, in order to be useful, product line models also need to be consistent with the actual code base (cf. challenge C1). Checking and maintaining consistency with the code base is particularly challenging as it largely depends on the implementation technology used. Furthermore, the code base is in practice often changed by developers outside the product line modeling environment.

Existing variability modeling approaches cover the upper two levels in Figure 1. Feature models, decision models, or orthogo-

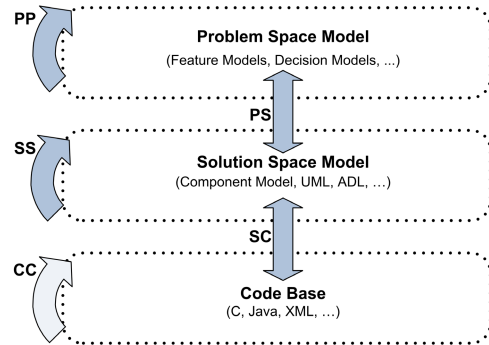


Figure 1: Modeling levels in product lines with types of intra- and inter-level inconsistencies.

nal models have been proposed for modeling the problem space [25, 10, 2]. Architecture description languages or general purpose languages like the UML can be used to define the solution space. There are also approaches such as DOPLER [17] or Clafer [3] that provide modeling support for both levels and mappings between them. Figure 1 shows the three modeling spaces together with five different types of consistency checks:

PP inconsistencies can exist within the problem space model. For instance, in a decision model cycles caused by decision dependencies need to be detected and prevented.

PS inconsistencies can occur between the problem space model and the solution space model. For instance, a component in the solution space model can have dependencies to two contradicting decisions. This can result in an inconsistency if it is impossible to resolve both decisions during product derivation.

SS inconsistencies can exist within the solution space model such as contradicting dependencies between components. One component might require two other components which in turn exclude each other.

SC inconsistencies can appear between the solution space model and the code base it represents. Such inconsistencies can for example occur whenever an asset (or its dependencies or attributes) are changed in the solution space model but not in the code base.

CC inconsistencies can exist within the code base. They are however typically addressed by integrated development environments and are thus not further investigated in this paper.

The section on related work will discuss existing research for detecting and fixing specific types of inconsistencies. However, the major practical challenge lies in providing an integrated environment capable of dealing with all these types in a holistic manner.

2.1 Modeling a Component-based Product Line with DOPLER

In cooperation with our industry partner Siemens VAI, we have been creating variability models for the CC-L2 software product line controlling continuous casting machines in steel plants. The (simplified) meta-model for modeling the problem space, the solution space and the code base is shown in Figure 2. The product line has been modeled using the DOPLER approach which uses decision models to define the problem space, asset models defining reusable elements of different types and their dependencies for the solution space, as well as mappings among decisions and assets [17, 22].

In DOPLER variation points are represented using *Decisions* which have a unique name and a question that is asked to a user

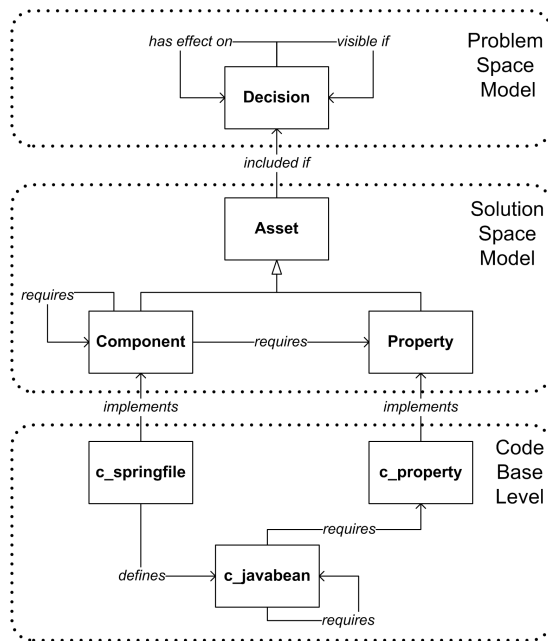


Figure 2: Partial Siemens VAI meta model with model elements for the problem space and solution space. The code base level defines additional elements for representing the component-based implementation of the product line.

during product derivation. Answering a question sets the value of a decision. Possible answers depend on the type of the decision (Boolean, enumeration, string, or number). The range of allowed values can be further restricted by validity conditions. Decisions can depend on each other hierarchically (if a decision needs to be taken before another decision becomes "visible") or logically (if taking a decision changes the value of another decision). Dealing with variability represents a specific challenge for consistency checking, as consistency constraints need to consider such dependencies (cf. challenge C1 and Table 1).

Assets represent the solution space in the product line (e.g., software components). *Assets* can depend on each other functionally (e.g., one component requires or excludes another component) or structurally (e.g., a component is part of a sub-system). DOPLER allows modeling assets at arbitrary granularity and with user-defined attributes and dependencies (based on a given set of basic types). Users can create domain-specific meta-models to define the types of assets, their attributes, and dependencies. In the Siemens VAI product line the asset types in the DOPLER variability models are components representing Spring XML component descriptions which in turn represent Java Beans and properties (key-value pairs). Diverse domain-specific dependencies have been defined. For example, a component can require another component or a property (cf. Figure 2). Dependencies between assets and decisions are explicitly modeled via inclusion conditions that define for an asset when it will be part of a derived product. Due to asset dependencies, not every asset needs to be related with a decision to be included in a derived product. For example, when an asset is included because of a decision, its *requires* dependencies might lead to the inclusion of other assets as well.

We decided to create a model image of the code base and to apply incremental transformation to synchronize it with the code. This allowed us to use the same checking mechanism for all different types

of artifacts. For that reason the meta-model was extended with a *code base level* containing the Spring files (*c_springfile*), the contained Java Beans (*c_javabean*) and their properties (*c_property*), as well as *defines* and *requires* relations among these elements (cf. Figure 2). Traceability to the solution space model is established by the *implements* relationship as we will show later.

2.2 Examples of Consistency Constraints

Figure 1 shows generic types of constraints (PP, PS, SS, SC, CC) in product line engineering. Here we describe a number of specific constraints based on these types needed in the component-based product line of our industry partner. Examples of the constraints are shown in Table 1. To allow arbitrary checks and freedom for the developer, the constraints are defined in the Java programming language. There are, however, no language restrictions and any constraint language could be used instead.

The problem space and problem to solution space constraints (PP and PS) are relevant in any DOPLER variability model. There are simple constraints that check whether decisions of type enumeration have at least two possible values (PP1) and that values for model elements with mandatory attributes have been defined (PP2). It is also essential to detect cycles in the decision model stemming from hierarchical and logical dependencies (PP3–5). PS1 is an example of a problem-to-solution space inconsistency. The constraint detects whether inclusion conditions exist in the solution space (assets) that never can evaluate to true because they refer to contradictory decisions in the problem space (cf. challenge C1).

Further constraints depend on the domain-specific meta-model for Siemens VAI (cf. Figure 2) and address intra-solution-space and model-to-code consistency. SS1 checks that each component modeled in the solution space (representing a Java Bean described in a Spring XML file) requires at least one property defining initialization parameters for that component. The most basic model to code constraint SC1 assures that each component modeled in the variability model exists in the code base of the product line. This constraint assures for example, that outdated components that are no longer available in the workspace are marked to be purged in the variability model as well. The constraints SC2 and SC3 cover the relations between components in the model, and the relations between Spring XML files in the file system (which in fact depend on relations between the Java Beans described in that Spring files). Both constraints assure that there are no unnecessary relations between components and that no relations are missing in the variability model. Constraints SC4-6 assure the consistency of variant type components which represent a particular characteristic of the Siemens VAI models. Variant types are used to group identical components implemented in different Spring XML files.

We discuss constraint SC2 in detail to illustrate its high-level operation sequence: SC2 checks the *requires* relations between components. As illustrated in Figure 3 a *requires* relation between two components in the model is only needed if it is based on an existing dependency in the product line code base. Each component is defined by a spring file which in turn is realized by one or more Java Beans. If at least one Java Bean contained in the Spring file (*Deburrer.xml*) requires a Bean defined in the second Spring file (*Caster.xml*) then the relation is needed on component level. Otherwise the consistency check will reveal an obsolete relation between the two components. Although SC2 is fairly simple it is important to note that such constraints have to be evaluated many times in complex models. For example, SC2 needs to be evaluated for each *requires* relationship among two components (and there are thousands of such relationships in our models).

We developed a batch-oriented checker that worked well with

Table 1: Examples of constraints in component-based product lines.

<i>Constraint</i>	<i>Description</i>
PP1: Enumeration decision	An enumeration decision must have at least two options to choose from.
PP2: Mandatory attribute	Mandatory attributes of model elements must not be empty.
PP3: Decision effect cycle	There must be no cycles caused by logical decision dependencies.
PP4: Visibility condition cycle	There must be no cycles caused by hierarchical decision dependencies (visibility).
PP5: Visibility condition self reference	A Visibility Condition must not contain a reference to itself.
PS1: Inclusion condition exclusion	An Inclusion Condition must not contain two contradictory decisions (to avoid assets that will never be included in a derived product).
SS1: Component properties	Each component (representing a Java Bean) requires at least one property defining initialization parameters for the Bean.
SC1: Component matching	Each component in the variability model must exist in the product line code base.
SC2: Component relation	Relations between components in the variability model must also exist in the product line code base.
SC3: Java Bean relation	A relation between Java Beans must be represented in the variability model as a component relation.
SC4: Variant type relation	Variant types must not have <i>requires</i> relations.
SC5: Variant type occurrence	If two or more components are identical all of them must contribute to the same variant type component.
SC6: Variant type consistency	Only identical components must contribute to a single variant type component.

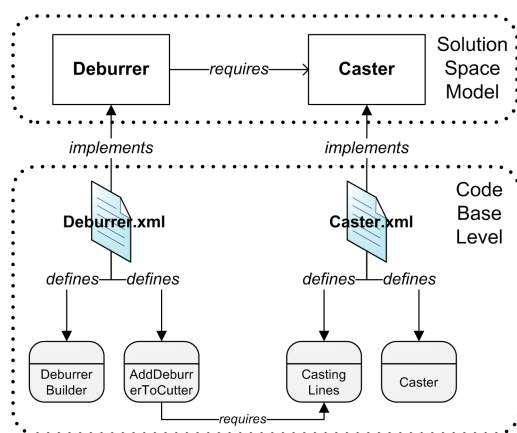


Figure 3: Schematic view of Constraint SC2.

small variability models and a limited number of constraints in earlier research. However, this approach did not scale well enough for larger models and high number of required consistency constraints that we encountered in industrial practice (cf. challenge C2). In particular, modelers requested immediate feedback regarding consistency after changes to the model. While technically it would have been possible to realize immediate feedback with a batch checker, the computational cost made it impractical to report inconsistencies after each change. Previously, our approach thus only reported inconsistencies after a user invoked the checker when saving the model in the tool. Typically many changes were made by modelers between two invocations of the checker and multiple new errors were reported at once. This made it difficult for modelers to relate the errors to the changes they made to the model or in the code base – changes that were made minutes or even hours earlier. Also, our batch-oriented checker was difficult to extend or adapt because the consistency constraints were woven together

into a long and complex algorithm. This was a big obstacle to the practical use of consistency checking because new constraints were increasingly hard to integrate in the existing algorithm.

3. INCREMENTAL CONSISTENCY CHECKING ON VARIABILITY MODELS

We thus pursued an incremental strategy to consistency checking to help modelers with detecting and tracking inconsistencies correctly and quickly after every change. We adapted an existing incremental checker [19] to support the fast checking of model changes. The checker identifies all model elements that affect the truth value of any given consistency constraint. A consistency constraint needs to be re-evaluated if one of these model elements changes. We refer to this set of affected model elements as the scope of a consistency constraint. The elements contained in the scope are stored in the scope database. Our incremental checker computes the change impact scope of a constraint instance automatically by observing the run-time behavior of consistency constraints during their evaluation. To this end, we developed the equivalent of a profiler for consistency checking. The profiling data is used to establish a correlation between model elements and consistency constraints. Based on this correlation, it then decides when to re-evaluate consistency constraints and when to display inconsistencies – thus allowing an engineer to quickly identify all inconsistencies that pertain to any part of the model of interest at any time.

However, thus far the checker had only been used and evaluated for UML models. Although our models might seem similar to UML on the surface, they exhibit quite a range of differences which required significant extensions we will describe in Section 4. For example, as discussed in Section 2 and shown in Table 1 many of our consistency constraints go beyond the checking of models and also include source code and other artifacts. Since the existing approach was limited to checking models, we had to incrementally transform code fragments into model fragments to enable their checking. Thus, our extended approach combines incremental transformation with incremental consistency checking.

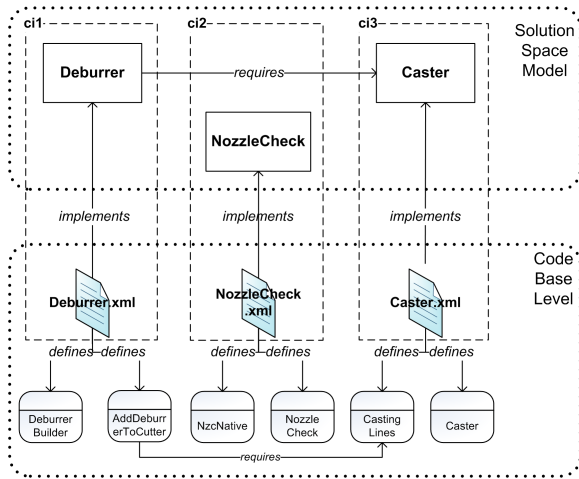


Figure 4: Incremental Evaluation of SC1 and SC2.

We illustrate the incremental consistency checker using consistency constraint SC1 (cf. Table 1) evaluating whether a modeled component is implemented in the source code. Figure 4 depicts a (simplified) excerpt of one model with three components in the top and a range of code-level constructs at the bottom. Consistency constraint SC1 is written from the perspective of a component – i.e., a given component is considered consistent if it is implemented by a code-level XML file. Since there are three components in Figure 4, each component has to be evaluated with regard to SC1. Our approach thus instantiates the consistency constraint SC1 three times, once for each component (ci1, ci2, ci3). We do this because incremental consistency checking needs to react to model changes and each instance of the consistency constraint is affected differently by model changes. This means that we need to compute how each consistency constraint or instance thereof is affected by a model change. For this purpose, we need to know the complete change impact scope for every constraint instance. Figure 4 indicates that the three SC1 instances ci1, ci2, and ci3 access distinct model elements when they are evaluated.

For example, the first instance *ci1* starts its evaluation at component *Deburrer*, then navigates along the *implements* relationship, and finally accesses *Deburrer.xml*. The constraint instance is satisfied because the component does have an associated code-level XML file. Our approach recognizes that from the entire model only these accessed model elements were needed to evaluate instance *ci1*. This instance must be re-evaluated only if one of those elements changes. Our approach thus maintains a table of instances of consistency constraints and the model elements they accessed during their evaluation to understand which instances must be re-evaluated when the model changes. It is also important to note that a model element may be accessed by multiple instances of one or more consistency constraints. For example, constraint SC2 evaluates a component with regard to the *requires* relationship. It makes sure that a component-level *requires* relationship is matched by a code-level *requires* relationship. Figure 4 depicts a *requires* relationship from *Deburrer* to *Caster* and during evaluation an instance of SC2 will access the components, their code-level XML files *Deburrer.xml* and *Caster.xml* and some of the *defines* and *requires* relationships underneath. The change impact scope of this instance of SC2 is also automatically observable through the model profiler. This scope is also larger than the scope of SC2 and encompasses

many of the model elements underneath *Deburrer* and *Caster* but not *NozzleCheck*. Some model elements, such as *Deburrer.xml* thus affect multiple instances of consistency constraints. Only these related instances must be re-evaluated if *Deburrer.xml* changes.

4. USING THE INCREMENTAL CHECKER IN THE DOPLER TOOL SUITE

A major goal when developing our incremental checker was to increase usability by providing immediate feedback to modelers about the detected inconsistencies (cf. challenge C1) as part of a development environment. We thus decided to seamlessly integrate the incremental consistency checker in the Eclipse-based DOPLER tool suite [18] which supports product line variability modeling and product derivation. In particular, we integrated the checker with DOPLER's modeling tool that allows defining decision and asset models as discussed in Section 2.

A modeler can define assets in the **Variability Model Editor** as shown in Figure 5. The Asset Overview shows an outline of already available software components. This view allows adding or removing components. Selected components can be modified in the Asset Detail View which provides information about its attributes and relations to other components. The modeler can simply add or remove relations to other components via drag and drop in the detail view. Manipulating components in the editor has an immediate effect. For instance, after adding a relation to another component all involved constraints (and only those!) are re-evaluated. Feedback about detected errors is provided within milliseconds after the user action leading to an inconsistency. The Error View provides information about all inconsistencies found in the model. The Error View also provides details about the source model elements causing the problem. This helps the modeler to resolve the problem, for instance, by removing an unneeded relation between two components from the model.

We illustrate our tool architecture by discussing the chain of events and data flow (cf. Figure 6). The incremental consistency checker relies on tracking changes in the Eclipse workspace. For that purpose, our **Eclipse Change Tracker** [24] observes arbitrary changes to the variability model and the Eclipse workspace. For instance, it listens to the "delete component" change to the variability model after the modeler deletes a component in the editor and propagates the change to the **Constraint Manager** which is responsible for initializing, managing, and storing constraints. Using the **Scope Database** the Constraint Manager determines the constraint instances to be re-evaluated after the deletion of a certain component. Our architecture makes use of the Eclipse extension point mechanism to allow adding new constraints flexibly and easy for new domains. Constraint definitions can easily be added or removed from the evaluation process by simply activating or deactivating plugins. The list of constraint instances is used by the **Incremental Consistency Checker** which applies incremental checking to variability models by controlling the execution of all affected constraint instances. The constraint instances do not directly access and query the **Variability Model**. Instead, they use the **Model Access Tracker** – a model adapter component which monitors and records all read access events to model elements for each single constraint instance that is evaluated. This fine-grained "model profiling" helps improving performance, because with each level of detail fewer constraint instances need to be evaluated eventually. The Model Access Tracker also builds the Scope Database ensuring that only necessary constraints are re-evaluated after changes to the model (not shown in Figure 6). The incremental consistency checker sends all newly found inconsistencies to the **Error Man-**

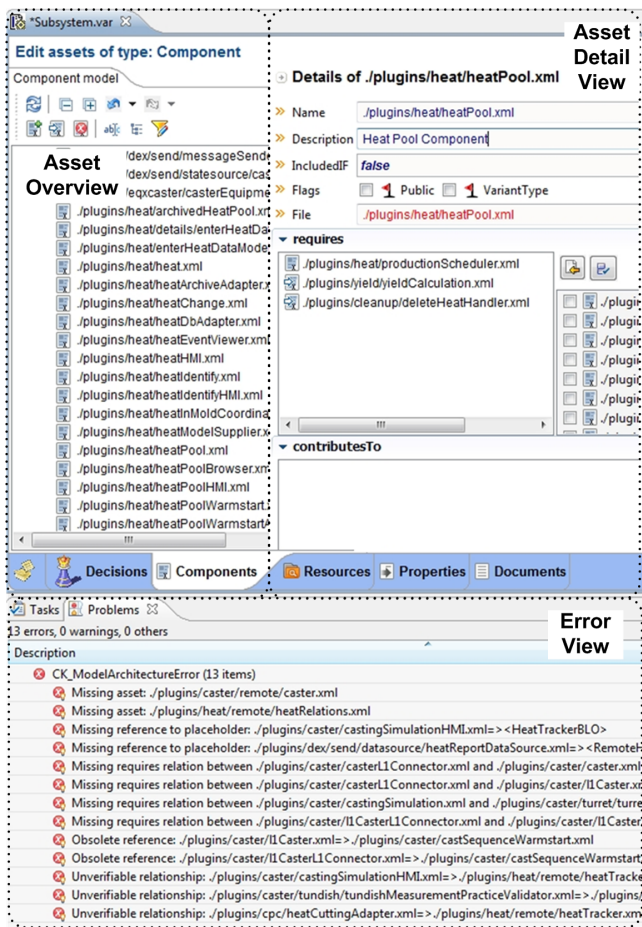


Figure 5: The DOPLER Variability Model Editor. The Incremental Consistency Checker identifies errors and displays them in an Error View.

ager which manages a list of errors for the user. It translates abstract constraint evaluation results into error messages meaningful for a modeler and performs basic filtering functions to avoid information overload. Finally, the **Error Viewer** (see also Figure 5) displays the detected inconsistencies for the evaluated constraint. We use the Eclipse marker mechanism for the easy management of errors and their presentation in a viewer.

5. EVALUATION

We define our evaluation goal using the Goal-Question-Metric (GQM) approach [5]. A goal in GQM consists of a purpose describing the actual aim of the evaluation, the issue of interest, a specific object to be measured and the viewpoint which describes the stakeholder perspective from which the goal is defined. Goals are refined using questions targeting explicit problems of interest. Finally, metrics are defined to obtain data addressing those questions. *Our evaluation goal is to assess (purpose) the performance (issue) of incremental consistency checking (object) from the viewpoint of the modeler (viewpoint).* We explore three questions regarding performance, memory usage, and scalability of the approach (cf. challenge C2.)

Question 1.1 "What is the impact of the incremental consistency checker on the start-up performance?" is further refined using two

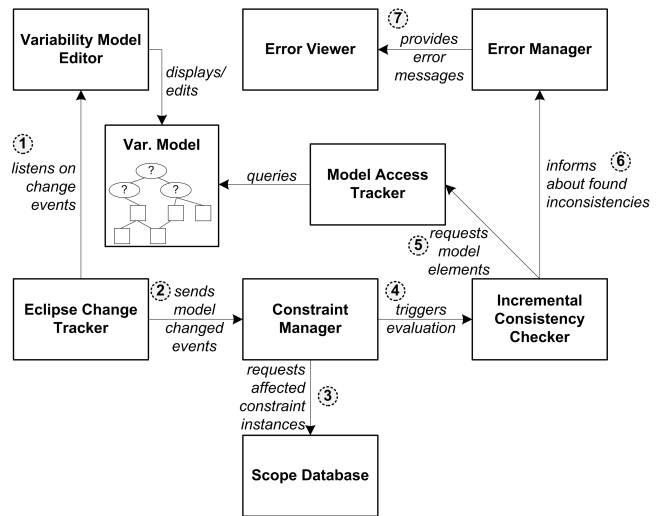


Figure 6: The Incremental Consistency Checker re-evaluates constraint instances after changes to the model.

metrics: M 1.1.1 is the time needed to initialize the incremental consistency checker (measured in milliseconds). It measures the overall start-up time when opening a variability model and includes library initializations as well as the creation of the scope database. The startup time is critical as engineers frequently open models in the DOPLER tools. M 1.1.2 measures the initial memory footprint of the incremental consistency checker after its start-up expressed by the number of objects created (i.e., constraint instances, scope elements). Understanding memory consumption is important as developers modeling real-world product lines have to work with very large models containing thousands of model elements.

Question 1.2 "What are the specific performance characteristics of the incremental consistency checker during modeling?" is further defined with metric M 1.2.1 measuring the time needed (in milliseconds) for consistency checking of key atomic modeling tasks such as adding, modifying, or deleting model elements. This metric is important as it shows the feasibility of the approach during everyday modeling activities. In particular, it shows whether the approach can provide "immediate feedback" to modelers.

Question 1.3 "How well does the approach scale?" investigates performance and memory consumption of our approach when applied to industrial models of different sizes. We compute the metrics M 1.1.1., M 1.1.2, and M 1.2.1 for real-world models of different sizes to address this question. In addition we use another metric to investigate this question. M 1.2.2 measures the number of constraint instances that need to be evaluated for the atomic modeling actions. This metric checks whether the approach scales regardless of the specific modeling actions performed by the user.

5.1 Evaluation Setup

We used variability models of the Siemens VAI CC-L2 product line to perform our evaluation. Regarding question 1.3 we evaluate the incremental checker with different model sizes. We thus created models of different sizes by incrementally merging the available variability models using the merging approach described in [16]. The smallest variability model I represents a single subsystem and consists of 57 components. We use two other models II and III for evaluation with 200 and 401 components. The largest model IV consists of 768 components. The underlying code base used for

Table 2: Siemens VAI solution space models used for evaluation.

Included subsystems		# C's	# CI's	# SE's
I	Cutting	57	1,044	2,084
II	Model I subsystems + Caster, Heating	200	1,902	3,385
III	Model II subsystems + Optimizer, JAMP, Simulator	401	3,108	5,180
IV	Model III subsystems + Analysis, DefectTracking, VAIQfeeder, Warmstart	768	5,310	8,175

evaluation contains 1,956 code model elements (702 XML spring files and 1,254 Java Beans) and is used in all four models I-IV.

Table 2 provides a detailed overview of the four models regarding the number of components (# C's), number of instantiated constraints (# CI's), and the total number of instantiated scope elements (# SE's) at startup. A constraint is instantiated for each model element only if needed as explained in Section 3. The scope database contains all scope elements defining the constraint instances that need to be re-evaluated after a change. The actual number of model elements is much higher in these models as we only show assets of type of Component in the table and do not consider the number of relations and model element attributes in Table 2.

5.2 Evaluation Results

We present the results gained by comparing the four different models in terms of startup and runtime performance when applying the incremental consistency checker.

The mechanism for incremental consistency checking relies on a complete execution of all constraints during startup to initialize the constraint instances and the scope database (scope elements need to be created). Figure 7 provides an overview of the initialization costs of the incremental checker for the four different models (lower solid line). The additional time needed to open a variability model lies between 2.5 seconds for a model with about 200 assets, and 5 seconds for a model with 400 assets. We used a PC with a Intel Pentium Dual Core 1.8 Ghz processor and 3 GB of main memory for the evaluation. This time is clearly acceptable for typical model sizes at our industry partner (100 to 400 components per single model). Also, this task is performed as a background thread when opening a model and does not block the modeler (even in the case of model IV where the initialization takes about 14 seconds).

The lower solid line is essential for the evaluation of the incremental checker. However, we made an additional interesting observation. Re-structuring the constraints already implemented in our legacy batch checker into the new format required by the incremental checker already had a significant impact on performance. We compared the complete time of executing all constraint instances with the initialization time of our legacy batch-checking approach. The dashed line demonstrates the scalability issues we had experienced with the batch-checking approach. Please note, however, that the legacy checker is not fully comparable in terms of the constraints it checks. However, the results show the order of magnitude of improvement gained by using the incremental checker.

Besides the initialization time, we also took a closer look at the memory costs (M 1.1.2). Memory consumption depends on the number and type of constraints and the number of model elements. The maximum number of constraint instances is created if every constraint is instantiated once for every asset in the model and every asset in the code base (this is only a theoretical case). Instead, constraints are typically instantiated for selected types of asset only

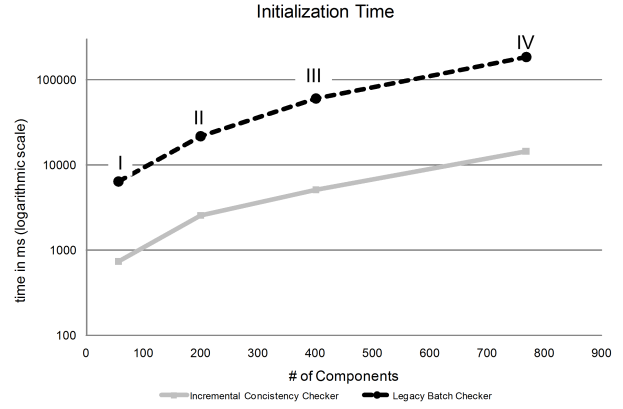


Figure 7: Metric M 1.1.1. Initialization time of the new incremental checker and evaluation time of legacy batch checker for models I-IV.

(e.g., components) and the real number of constraint instances per model is therefore much lower than the theoretical maximum.

The maximum number of scope elements can be calculated as the number of assets times the number of attributes per asset. Scope elements are created only once, but can be used in multiple constraint instances. Figure 8 shows the linear increase of scope elements and constraint instances for the four variability models demonstrating the scalability of our approach.

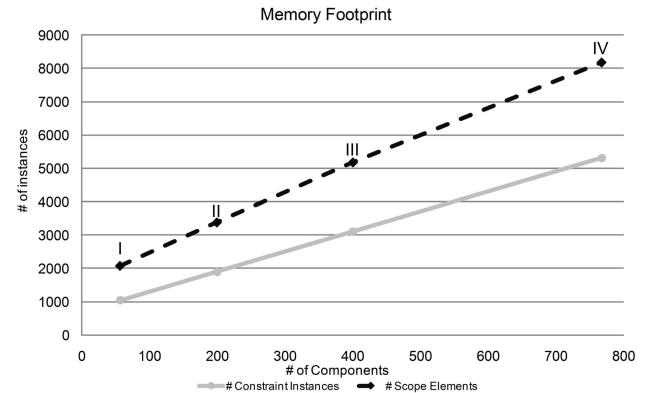


Figure 8: Metric M1.1.2. Number of constraint instances and scope elements after startup (models I-IV).

In Section 4.1 we described a number of basic tasks when working with the DOPLER variability modeling tool. We use these typical actions on variability models to gain data concerning the time needed to evaluate model changes at runtime. We simulate creating and deleting assets, as well as adding and removing relations for our live evaluation (metric M 1.2.1) and the analysis of affected constraints during one evaluation cycle (metric M 1.2.2).

For live evaluation analysis each single task has been performed 100 times to outweigh effects such as unpredictable background tasks. The mean values gained are shown in Figure 9. The costs for adding assets, as well as adding and removing relations to/from an asset remains constant in terms of model size. A single action performed on a variability model took only about 5 to 10 ms which

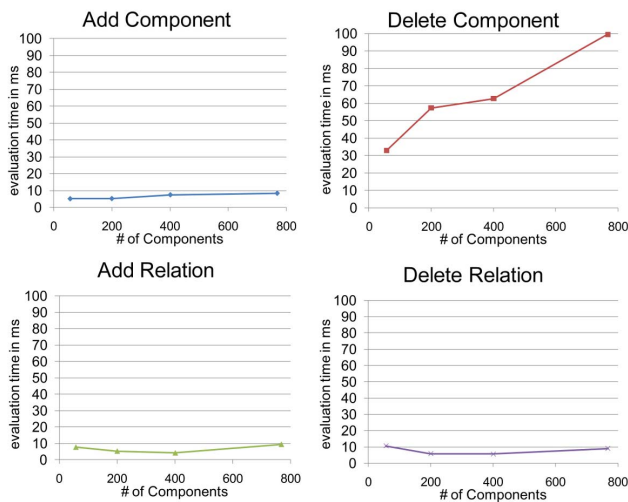


Figure 9: Metric M 1.2.1. Performance of basic modeling actions for models I-IV.

demonstrates that our approach indeed provides *immediate feedback*.

Only the task of deleting an asset from a model increases linearly. This is because constraints that were instantiated for the asset to be deleted, need to be found and removed from the consistency checking instance. We expect however that this problem can be easily fixed by using a different data structure.

While constraint instances and scope elements in total increase depending on model size (as shown in Figure 8), the number of constraint instances that need to be re-evaluated during live evaluation stays constant. Depending on the type of change very few constraints need to be evaluated (see Figure 10). The number of constraints that need to be evaluated remains constant when adding assets in models of different size. Adding and deleting relations as well as deleting component moves between 2.5 and 3.2 constraints evaluated per change event, again tested on 100 change samples.

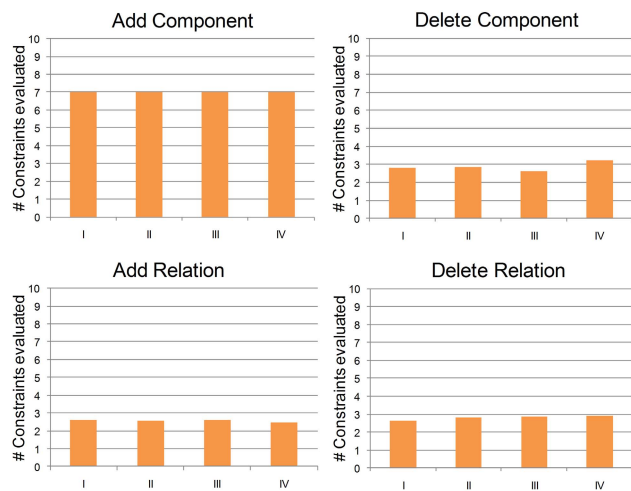


Figure 10: Metric M 1.2.2. Number of constraint instances affected by a change for models I-IV.

5.3 Lessons Learned

We derive the following lessons from our experiences:

Support different types of consistency checks. Our aim was to develop a holistic approach capable of checking various consistency constraints within and among different levels. In this paper we distinguished five different types of inconsistencies depending on the modeling level (problem space vs. solution space vs. code base; cf. challenge C1) and presented both generic constraints (applicable to variability modeling with DOPLER) and domain-specific constraints (applicable to the specific technical platform of Siemens VAI). Domain-specific constraints (which depend on the technology used; mainly SS and SC types) can greatly complicate consistency checking but are very important for the usefulness and acceptance of the checker in practice (cf. challenge C3).

Use a model-based representation to check consistency. Existing consistency checking mechanisms either work on model level or on code level but not on both and among them (cf. challenge C1). We learned that the benefits of an explicit representation of the code in models outweighs the overhead of increased memory consumption and slightly increased startup time. However, allowing the checker to operate only on the level of models makes the implementation more elegant and easier to understand. This solution has the advantage that the checker is more resilient to technical changes as only the transformation algorithm has to be updated, for example, if the component technology language is replaced.

Decide whether to allow inconsistencies or to avoid them in the first place. An interesting issue is that while it is important to know about inconsistencies it is often too distracting to resolve them right away. The notion of "living with inconsistencies" [4] advocates that there is a benefit in temporarily allowing inconsistencies in design models. While our approach provides inconsistencies instantly, it does not require the engineer to fix them instantly. Our approach tracks all presently-known inconsistencies and lets the engineer explore inconsistencies according to his/her interests in the model. However, there are situations where it is beneficial to avoid inconsistencies in the first place. When developing constraints for the different types we noticed that the number of generic constraints (mainly PP and PS types) can be kept small by doing trivial plausibility checks (on all levels from PP to CC) in the variability modeling tool. This lesson is related with challenges C1 and C2 defined in the introduction.

Go for incremental checking. While our incremental consistency checking support for product line models has some disadvantages (initialization time, memory consumption, redundant representation of code in models and the required transformation) the benefits outweigh these drawbacks as shown in our evaluation. In particular, we have learned that the run-time performance enabling immediate user feedback greatly contributes to the usefulness and acceptance in practical settings (cf. challenge C2).

Support the evolution of constraints. Constraints evolve over time based on experiences made as well as due to changing and evolving technology. Both product line tools as well as the product lines they describe can evolve. In our case, the product line tools were under constant development during the early phases of the project and the product line of Siemens VAI was refactored and evolved in parallel. This is why we developed an extensible architecture for our consistency checker to allow adding rules incrementally and to allow easily replacing and modifying existing rules. Our batch-oriented checker was difficult to extend or adapt because the consistency constraints were woven together into a large and complex algorithm. This was a big obstacle to the practical use of consistency checking because new constraints were increasingly hard to integrate in the existing algorithm. Thus, the extensible ar-

chitecture turned out to be very important due to this continuous evolution (cf. challenge C3).

6. RELATED WORK

We structure our discussion of related work according to the types of inconsistencies defined in Section 2.

PP. Batory [6] describes an approach combining feature models, grammars, and propositional formulas to allow arbitrary propositional constraints to be defined among features and checked using satisfiability solvers. Other authors have addressed the automatic analysis of feature models [8]. For example, the FaMa approach [33] allows using different solvers in the back-end to perform analysis operations on feature models.

PS. Heidenreich [23] explores possibilities for checking the well-formedness of models describing features in the problem space and their realizations in the solution space to ensure their correctness. Several authors have presented work on type systems to guarantee that only well-typed programs are generated: Apel *et al.* have developed such a type system on the basis of a formal model of a feature-oriented Java-like language [26, 1]. Similarly, Delaware *et al.* [14] present an extension to Lightweight Java with support for features together with a constraint-based type system.

SS. Several approaches are based on identifying inconsistencies between different design models by direct comparison. Some of these approaches also perform incremental consistency checking. xLinkIt [29] allows evaluating the consistency of XML-based documents. The approach can check the consistency of entire UML models but can also handle incremental consistency checking by only evaluating changes to versions of a document. ArgoUML [31] detects inconsistencies in UML models based on annotated consistency rules that also enable incremental consistency checking. The approach implements two consistency checking mechanisms: consistency rules without annotations are placed into the queue which is continuously evaluated in the background using a batch-checker at 20% CPU time. Consistency rules with annotations are evaluated using an incremental checker. It has been demonstrated that ArgoUML's type-based consistency checking produces good performance but it is not able to keep up with an engineer's rate of model changes in very large models [21]. Similar to our approach Blanc *et al.* [9] address the issue of incremental consistency checking from the perspective of model changes. However, their consistency rules need to be defined explicitly in terms of their impact on changes. If done correctly this leads to good performance. However, since writing these annotations may easily cause errors, they are no longer able to guarantee the correctness of incremental consistency checking. While researchers generally agree on the importance of consistency checking, the methods on how to detect them vary widely. For example, Tsiolakis and Ehrig [34] present an approach for checking the consistency between class and sequence diagrams based on a common graph structure. Van der Straeten *et al.* [32] use description logic to detect inconsistencies between sequence and state chart diagrams. Campbell *et al.* [11] use a model checker to evaluate inconsistencies within and across UML diagrams. Zisman and Kozlenkov [36] use a knowledge base and express consistency rules using patterns and axioms.

SC. Several papers address the issue of consistency between models and code in evolving product lines. Murta *et al.* [28] present an approach for ensuring consistency of architectural models to implementation during evolution. Their approach support arbitrary evolution policies and is based on recording changes in a configuration management system. Product line evolution support becomes critical in model-based approaches to ensure consistency after changes to meta-models, models, and actual development arti-

facts. Mende *et al.* [27] describe tool support for the evolution of software product lines based on the "grow-and-prune" model. They support identifying and refactoring code that has been created by copy & paste and which might be moved from product to product line level. Deng *et al.* [15] describe a model-driven product line approach that addresses the issue of domain evolution in product line architectures with model transformations.

7. CONCLUSIONS AND FUTURE WORK

We presented experiences, lessons learned, and evaluation results of applying an incremental consistency checker on product line variability models. The incremental consistency checker works on and across different levels of variability models and also checks consistency between variability models and source code. The incremental consistency checker is independent from the domain-specific DOPLER meta-model and can be easily used with arbitrary meta-models. As our event tracking mechanism allows to identify changes down to the level of model element attributes only few constraints need to be evaluated during incremental consistency checking. Our evaluation with large-scale models demonstrates the performance and scalability of the approach. It is fast enough to provide immediate feedback to users and identifies errors within 5 to 10 ms for typical modeling actions. Furthermore, our approach allows adding new constraints in a flexible manner via Eclipse extension points. Constraints can be activated and deactivated as needed in certain domains.

While this paper focused on the practical experiences and lessons learned we plan to report on the technical details of our approach in a further paper. In particular, we will describe the incremental code-to-model transformation approach we developed as part of this work. We will also experiment with more and other types of constraints and plan to investigate how dependencies among constraints can be exploited during constraint evaluation. Finally, we will extend our tools to support fixing identified inconsistencies as already demonstrated in [20].

Acknowledgements

This work has been supported by Siemens VAI Metals Technologies, the Christian Doppler Forschungsgesellschaft, and the Austrian FWF grant P21321-N15. We are also grateful to Martin Lehofer and Deepak Dhungana for their support.

8. REFERENCES

- [1] S. Apel, C. Kästner, A. Größlinger, and C. Lengauer. Type safety for feature-oriented product lines. *Automated Software Engineering – An International Journal*, 17:251–300, 2010.
- [2] F. Bachmann, M. Goedicke, J. Leite, R. Nord, K. Pohl, B. Ramesh, and A. Vilbig. A meta-model for representing variability in product family development. In F. van der Linden, editor, *5th Int'l WS on Sw. Product-Family Engineering*, volume LNCS 3014, pages 66–80, Siena, Italy, 2003. Springer.
- [3] K. Bak, K. Czarnecki, and A. Wasowski. Feature and class models in Clafer: Mixed, specialized, and coupled. Technical report, University of Waterloo, 2010.
- [4] R. Balzer. Tolerating inconsistency. In *Int'l Conf. on Software Engineering (ICSE)*, pages 158–165, 1991.
- [5] V. Basili, G. Caldiera, and D. Rombach. Goal/question/metric paradigm. In J. Marciniak, editor, *Encyclopedia of Sw. Eng.*, pages 528–532, New York, 1994. John Wiley and Sons.

- [6] D. S. Batory. Feature models, grammars, and propositional formulas. In *Proc. Int'l Conf. Sw. Product Lines*, pages 7–20, 2005.
- [7] B. Belkhouche and C. L. Olalde. Multiple view analysis of designs. In *Proc. of the 2nd Int'l Software Architecture WS (ISAW-2) and Int'l WS on Multiple Perspectives in Software Development on SIGSOFT '96 WSs*, pages 159–161, New York, NY, USA, 1996. ACM.
- [8] D. Benavides, S. Segura, and A. Ruiz-Cortes. Automated analysis of feature models 20 years later: a literature review. *Information Systems (in Press)*, 2010.
- [9] X. Blanc, I. Mounier, A. Mougnot, and T. Mens. Detecting model inconsistency through operation-based model construction. In *Proc. ICSE Conf.*, pages 511–520, 2008.
- [10] G. H. Campbell, S. R. Faulk, and D. M. Weiss. Introduction to Synthesis. Technical report, Software Productivity Consortium, Herndon, VA, USA, 1990.
- [11] L. A. Campbell, B. H. C. Cheng, W. E. McUmber, and K. Stirewalt. Automatically detecting and visualising errors in UML diagrams. *Requir. Eng.*, 7(4):264–287, 2002.
- [12] B. H. C. Cheng, E. Y. Wang, and R. H. Bourdeau. A graphical environment for formally developing object-oriented software. In *ICTAI*, pages 26–32, 1994.
- [13] K. Czarnecki. Variability modeling: State of the art and future directions (keynote presentation). In *Proc. 4th Int'l WS on Variability Modelling of Software-intensive Systems (VaMoS 2010)*, page 11, Linz, Austria, 2010. ICB-Research Report 37, Univ. of Duisburg Essen.
- [14] B. Delaware, W. R. Cook, and D. S. Batory. Fitting the pieces together: a machine-checked model of safe composition. In *Proc. ESEC/FSE Conf., 2009, Amsterdam*, pages 243–252, 2009.
- [15] G. Deng, J. Gray, D. Schmidt, Y. Lin, A. Gokhale, and G. Lenz. Evolution in model-driven software product-line architectures. In P. Tiako, editor, *Designing software-intensive systems*, pages 1280–1312. Idea Group Inc (IGI), 2008.
- [16] D. Dhungana, P. Grünbacher, R. Rabiser, and T. Neumayer. Structuring the modeling space and supporting evolution in software product line engineering. *Journal of Systems and Software*, 83(7):1197–1122, 2010.
- [17] D. Dhungana, P. Heymans, and R. Rabiser. A formal semantics for decision-oriented variability modeling with DOPLER. In *Proc. 4th Int'l WS on Variability Modelling of Software-intensive Systems (VaMoS 2010)*, pages 29–35, Linz, Austria, 2010. ICB-Research Report 37, Univ. of Duisburg Essen.
- [18] D. Dhungana, R. Rabiser, P. Grünbacher, and T. Neumayer. Integrated tool support for software product line engineering. In *Proc. ASE Conf.*, pages 533–534, Atlanta, Georgia, USA, 2007. ACM.
- [19] A. Egyed. Instant consistency checking for the UML. In *28th Int'l Conf. on Software Engineering*, pages 381–390, New York, NY, 2006. ACM.
- [20] A. Egyed, E. Letier, and A. Finkelstein. Generating and evaluating choices for fixing inconsistencies in UML design models. In *Proc. of the 2008 23rd IEEE/ACM Int'l Conf. on Automated Sw. Eng.*, pages 99–108, Washington, DC, USA, 2008. IEEE Computer Society.
- [21] A. Egyed and D. S. Wile. Support for managing design-time decisions. *IEEE TSE*, 32(5):299–314, 2006.
- [22] P. Grünbacher, R. Rabiser, D. Dhungana, and M. Lehofer. Model-based customization and deployment of Eclipse-based tools: Industrial experiences. In *Proc. ASE Conf.*, pages 247–256, 2009.
- [23] F. Heidenreich. Towards systematic ensuring well-formedness of software product lines. In *FOSD '09: Proc. of the First Int'l WS on Feature-Oriented Software Development*, pages 69–74, New York, NY, USA, 2009. ACM.
- [24] W. Heider, R. Rabiser, D. Dhungana, and P. Grünbacher. Tracking evolution in model-based product lines. In *1st Int'l WS on Model-driven Approaches in Software Product Line Engineering (MAPLE 2009)*, *Proc. (vol 2) of the 13th Int'l Software Product Line Conf. (SPLC 2009)*, pages 59–63, San Francisco, CA, USA, 2009. SEI CMU.
- [25] K. Kang, S. Cohen, J. Hess, W. Nowak, and S. Peterson. Feature-oriented domain analysis (FODA) feasibility study. Technical report, SEI-CMU, USA, 1990.
- [26] C. Kästner and S. Apel. Type-checking software product lines – a formal approach. In *23rd IEEE/ACM ASE Conf., L'Aquila, Italy*, pages 258–267, 2008.
- [27] T. Mende, F. Beckwermer, R. Koschke, and G. Meier. Supporting the grow-and-prune model in software product lines evolution using clone detection. In *12th European Conf. on Sw. Maintenance and Reengineering*, pages 163–172, Washington, DC, USA, 2008. IEEE CS.
- [28] L. G. P. Murta, A. van der Hoek, and C. M. L. Werner. ArchTrace: Policy-based support for managing evolving architecture-to-implementation traceability links. In *Proc. ASE Conf.*, pages 135–144, 2006.
- [29] C. Nentwich, L. Capra, W. Emmerich, and A. Finkelstein. xlinkit: a consistency checking and smart link generation service. *ACM Trans. Internet Techn.*, 2(2):151–185, 2002.
- [30] K. Pohl, G. Böckle, and F. van der Linden. *Software Product Line Engineering: Foundations, Principles, and Techniques*. Springer, 2005.
- [31] J. Robins et al. ArgoUml, <http://argouml.tigris.org/>. Technical report.
- [32] R. V. D. Straeten, T. Mens, J. Simmonds, and V. Jonckers. Using description logic to maintain consistency between UML models. In *Proc. 6th Int'l UML Conf., San Francisco, CA, USA*, pages 326–340, 2003.
- [33] P. Trinidad, D. Benavides, A. Ruiz-Cortés, S. Segura, and A. Jimenez. FaMa framework. In *Proc. Sw. Product Lines Conf.*, page 359, 2008.
- [34] A. Tsiolakis and H. Ehrig. Consistency analysis of UML class and sequence diagrams using attributed graph grammars. In *Proc. of the Graph Transformation and Graph Grammars (GRATA), Berlin, Germany*, pages 77–86, 2000.
- [35] M. Vierhauser, D. Dhungana, W. Heider, R. Rabiser, and A. Egyed. Tool support for incremental consistency checking on variability models. In *Proc. 4th Int'l WS on Variability Modelling of Software-intensive Systems (VaMoS 2010)*, pages 171–174, Linz, Austria, 2010. ICB-Research Report 37, Univ. of Duisburg Essen.
- [36] A. Zisman and A. Kozlenkov. Knowledge base approach to consistency management of UML specification. In *Prof. ASE Conf.*, pages 359–363, 2001.