

Automating the maintenance of nonfunctional system properties using demonstration-based model transformation

Yu Sun^{1,†}, Jeff Gray², Romain Delamare², Benoit Baudry³ and Jules White⁴

¹*Department of Computer and Information Sciences, University of Alabama at Birmingham, Birmingham, USA*

²*Department of Computer Science, University of Alabama, Tuscaloosa, USA*

³*IRISA / INRIA Rennes, Rennes, France*

⁴*Department of Electrical and Computer Engineering, Virginia Tech, Virginia, USA*

ABSTRACT

Domain-Specific Modeling Languages (DSMLs) are playing an increasingly significant role in software development. By raising the level of abstraction using notations that are representative of a specific domain, DSMLs allow the core essence of a problem to be separated from irrelevant accidental complexities, which are typically found at the implementation level in source code. In addition to modeling the functional aspects of a system, a number of nonfunctional properties (e.g., quality of service constraints and timing requirements) also need to be integrated into models in order to reach a complete specification of a system. This is particularly true for domains that have distributed real time and embedded needs. Given a base model with functional components, maintaining the nonfunctional properties that crosscut the base model has become an essential modeling task when using DSMLs.

The task of maintaining nonfunctional properties in DSMLs is traditionally supported by manual model editing or by using model transformation languages. However, these approaches are challenging to use for those unfamiliar with the specific details of a modeling transformation language and the underlying metamodel of the domain, which presents a steep learning curve for many users. This paper presents a demonstration-based approach to automate the maintenance of nonfunctional properties in DSMLs. Instead of writing model transformation rules explicitly, users demonstrate how to apply the nonfunctional properties by directly editing the concrete model instances and simulating a single case of the maintenance process. By recording a user's operations, an inference engine analyzes the user's intention and generates generic model transformation patterns automatically, which can be refined by users and then reused to automate the same evolution and maintenance task in other models. Using this approach, users are able to automate the maintenance tasks without learning a complex model transformation language. In addition, because the demonstration is performed on model instances, users are isolated from the underlying abstract metamodel definitions. Our demonstration-based approach has been applied to several scenarios, such as auto scaling and model layout. The specific contribution in this paper is the application of the demonstration-based approach to capture crosscutting concerns representative of aspects at the modeling level. Several examples are presented across multiple modeling languages to demonstrate the benefits of our approach. Copyright © 2013 John Wiley & Sons, Ltd.

Received 5 October 2011; Revised 24 May 2013; Accepted 3 June 2013

KEY WORDS: nonfunctional property maintenance; model transformation by demonstration; domain-specific modeling language

1. INTRODUCTION

Domain-Specific Modeling (DSM) [19] has become an important part of Model-Driven Engineering (MDE) [29] to address the complexity of software systems. A Domain-Specific Modeling Language (DSML) can be used to declaratively define a software system with specific domain concepts. Various software artifacts (e.g., source code, simulation scripts, and XML deployment description

*Correspondence to: Yu Sun, Department of Computer and Information Sciences, University of Alabama at Birmingham, Birmingham, USA.

†E-mail: yusun@cis.uab.edu

files) can be generated automatically from the models. By raising the level of abstraction, DSM helps to protect key intellectual assets from technology obsolescence, resulting in less effort and fewer low-level details to specify a given system, which can lead toward improvements for supporting end user development [11].

Although much effort has been focused on using DSMLs to specify the functional properties of a system, the correct and complete specification of nonfunctional properties is also critical in many domains such as distributed systems, embedded systems, or large-scale applications. The nonfunctional properties range from performance, timing, and QoS constraints to some behavior requirements such as logging and concurrency control. Although, several modeling tools (e.g., generic modeling environment [24], graphical modeling framework [16], and generic eclipse modeling system [GEMS] [15]) have been developed to allow users to define various types of DSMLs and instantiate model instances, the tools to support the maintenance of nonfunctional properties in a DSML are not well-developed. In most situations, manual maintenance has to be carried out through the basic operations provided by the editing environment, which is tedious, error-prone, and time-consuming, particularly when the base model is large and complex. A manual process for maintaining model changes often involves repeated and complicated changes and computations [17].

An alternative for automating the maintenance of nonfunctional properties with DSMLs is to use executable Model Transformation Languages (MTLs) [30], which can define a set of transformation rules that specify how to add, remove, or modify the nonfunctional properties in a base model. Although, MTLs are very powerful and expressive for handling a wide range of nonfunctional properties in diverse DSMLs, using a transformation language is not always the perfect solution. First, even though most MTLs are high-level declarative languages, they have a steep learning curve because of the complexity of their syntax, semantics, and other special features (e.g., Object Constraint Language (OCL) [26] specification is used in many MTLs). Moreover, model transformation rules are often defined at the metamodel level, rather than in the context of a concrete model instance, which exposes users to abstract metamodel concepts. Because many potential model users (e.g., requirements engineers and domain experts) are not necessarily software engineers or programmers, learning transformation languages, and understanding the formal syntax definitions may be beyond their capability. Such challenges may prevent some users from realizing nonfunctional property maintenance tasks for which they have extensive domain experience. Consequently, this leads to an irony that a technology (i.e., DSM) meant to enable end users to participate in the specification of functional components of a software system does not enable end users to manage the nonfunctional properties easily.

To overcome the problem and provide an end-user approach to support automating the maintenance of nonfunctional properties in DSMLs, we have been investigating and extending the idea of Model Transformation By Demonstration (MTBD) [31], which applies a demonstration-based approach to enable users to specify the desired maintenance process of nonfunctional properties and to execute the process to any model instance through an execution engine. The goal of the approach is to assist general model users (e.g., domain experts and nonprogrammers) in realizing the automated maintenance of nonfunctional properties in DSMLs without knowing a specific model transformation language or the metamodel definitions.

This paper provides the following contributions to the study of nonfunctional system property maintenance in DSMLs:

- We describe how MTBD is extended and adapted to support the demonstration of nonfunctional system property maintenance tasks in a DSML (including a user refinement mechanism and an updated transformation engine to enable users to accurately specify *where* to apply the properties and *how* to apply them) without using MTLs or exposing metamodel definitions to end users.
- We present how MTBD can generalize and infer a transformation pattern that can be used to automate the whole maintenance process through an execution engine, as well as how users can refine the pattern.
- We select two typical nonfunctional system property maintenance tasks (i.e., adding a logging mechanism and applying QoS strategies) in DSMLs that have already been solved using MTLs and present the new solutions using MTBD.

- We present an evaluation by comparing efforts (based on operational effort to perform the same tasks) between the two approaches to solve the exact same tasks. This evaluation identifies the key advantages of using MTBD over MTLs in the maintenance of nonfunctional system property tasks.

The remainder of the paper is organized as follows: two nonfunctional property maintenance scenarios in two different DSMLs are presented in Section 2 as motivating examples; Section 3 explains the MTBD concept through two motivating examples. In Section 4, a comparison of MTBD and the use of traditional model transformation languages is given in the context of the two motivating examples. Related works are discussed in Section 5, and Section 6 offers concluding remarks and summarizes the directions of future work.

2. MOTIVATING EXAMPLES OF NONFUNCTIONAL PROPERTIES IN MODELS

This section presents two examples that motivate the need for automating the maintenance of nonfunctional system properties in different DSMLs. For each of the examples in Sections 2.1 and 2.2, background information about the specific application domain and the context of the DSML will be given, followed by an illustration using a concrete model instance. Then, we present a typical nonfunctional system property maintenance scenario in the domain and describe the desired model instance after the maintenance of these properties. The challenges of accomplishing each task will be summarized afterwards. We have used both of these examples in previous efforts focused on MTLs [17, 18]. In this paper, we compare our new approach that is focused on end-user demonstration to the more traditional approach that uses MTLs. The solution of using MTBD to address the exact same problems will be given in Section 3, so that a better comparison can be made between the traditional MTL approach and the new MTBD approach.

2.1. Maintaining nonfunctional system properties in Embedded Systems Modeling Language models

The Embedded Systems Modeling Language (ESML) [22] is a DSML used to graphically model real-time mission computing embedded avionics systems. ESML is specifically designed for modeling component-based avionics application deployment and distribution middleware infrastructure, which allows users to model the system from several different aspects such as *Interfaces*, *Events*, *Components*, *Interactions*, and *Configurations*.

Figure 1 shows an excerpt of an ESML model specified for one usage scenario. The top of Figure 1 illustrates the interaction among components in a mission-computing avionics application. The set of components collaborate to process a video stream that provides a pilot with real-time navigational data. The bottom of Figure 1 shows the internal representation of two components (*Display_Device* and *AirFrame_SynchProxy*), revealing the data elements and other constructs used to describe the infrastructure of component deployment and the distribution middleware. An event-driven model is implemented for the infrastructure that enables different components to update and transfer data to each other through event notification and callback methods.

2.1.1. Maintaining nonfunctional properties in Embedded Systems Modeling Language. The need to maintain nonfunctional properties in ESML models arises from the requirement to apply nonfunctional properties to each *Data* in the components. For instance, different logging mechanisms (e.g., *LogOnRead* and *LogOnWrite*) can be added to each data element to realize the recording policy of a black-box flight data recorder; specific constraints are occasionally required to attach the data element as preconditions to assert a set of valid values when a client invokes the component at runtime; in some cases, concurrency elements (e.g., *Internal Locking* and *External Locking*) specifying the synchronization strategy should be distributed across the components.

As the first example used in this paper, we choose the task of adding a logging mechanism for each data element, which is defined as follows:

Example 2.1: In each implementation Component (i.e., the name of the component ends with 'impl'), if an Action exists in the Component, a *LogOnRead* logging element should be attached to every Data element in the Component.

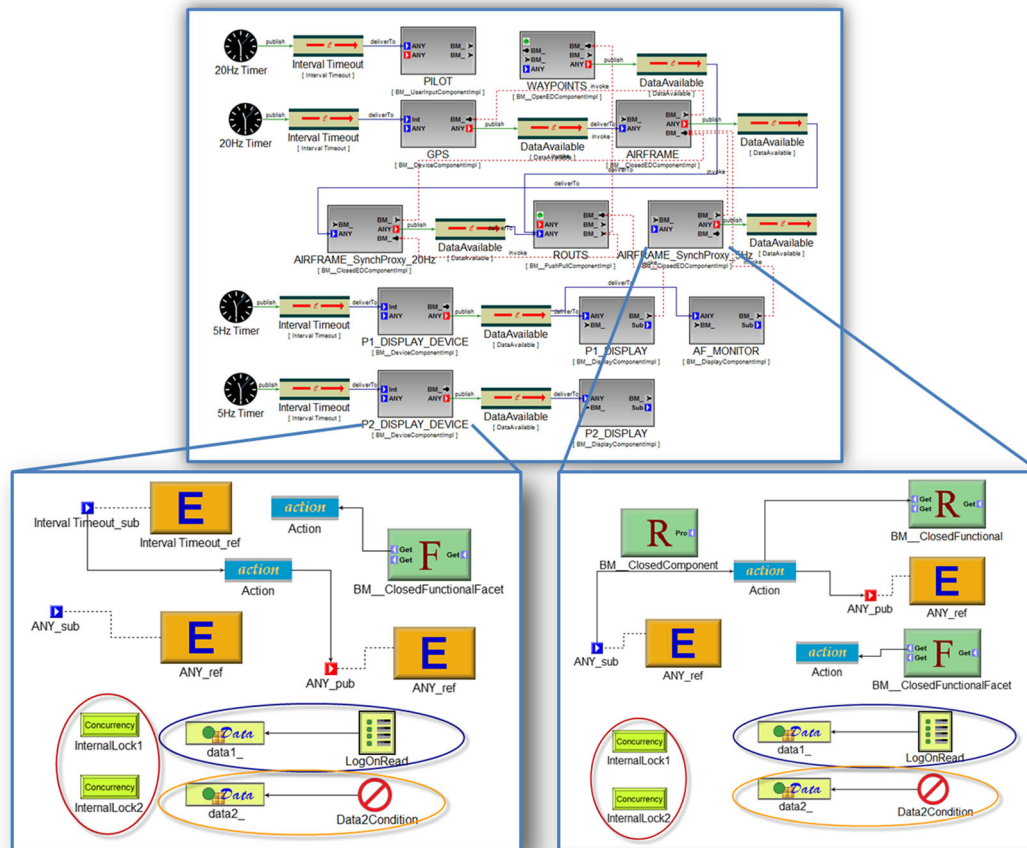


Figure 1. Adding nonfunctional properties in Embedded System Modeling Language models (adapted from [17]).

The tasks of adding data constraints and synchronization policies share a similar process as Example 2.1, and therefore, can be solved using the same approach. A correct process for maintaining these nonfunctional properties involves navigating the model and locating the correct places to insert the nonfunctional properties and then adding the correct nonfunctional elements and specifying the corresponding attributes. Although the maintenance process can be carried out by manual editing, it is tedious, time-consuming, and error-prone, particularly when the base model is large and complex (e.g., the partial system model shown in Figure 1 is only a subset of a system with thousands of components). Thus, an automated process can benefit the maintenance process when evolving model properties.

2.2. Maintaining nonfunctional system properties in QoS Adaption Modeling Language Models

The maintenance and evolution of distributed real-time and embedded systems is often challenging because of the consideration of different QoS constraints that might conflict with each other and must be treated as trade-offs among a series of alternative design decisions [18]. The QoS Adaption Modeling Language (QoSAML) [18] was designed to address the challenges of using an MDE approach, which uses a finite state machine representation extended with hierarchy and concurrency mechanisms to model the QoS adaptive behavior of the system.

One successful usage of QoSAML is to specify the QoS properties for an unmanned aerial vehicle (UAV) [21], which is an aircraft that is capable of surveying dangerous terrain and territories. The UAV continuously sends video streams to a central distributor, so that operators can observe the video and give further commands to the UAV. In order to reach a precise and timely response from operations, a smooth video stream must be guaranteed, which means that

the video must not be stale or be affected by jittering. However, because of the changing conditions in the surveillance environment, the fidelity of the video stream must be maintained by adjusting the QoS parameters. For example, as shown in Figure 2, in good conditions where a reliable network transmission is available, a smooth video stream can be kept using a high video *Size* and a high video *FrameRate*, whereas in a poor environment, both video *Size* and *FrameRate* should be reduced in order to keep the same video transmission latency.

Figure 3 is part of a QoSAML model that specifies the QoS properties for a UAV application. In this model, the latency is a dependent variable input to a hierarchical state machine called *Outer State*. Inside the *OuterState*, there are state machines that describe the adaption of identified independent control parameters, such as *Size State* and *FrameRate State*. In each of the state machines, several *States* are included to represent the options for the corresponding control parameter. A state specifies three *Data* for the option, *Pri* defines the priority of this option; *Max* defines the maximum value that can be used for this parameter; and *Min* defines the minimum value for the parameter. Model

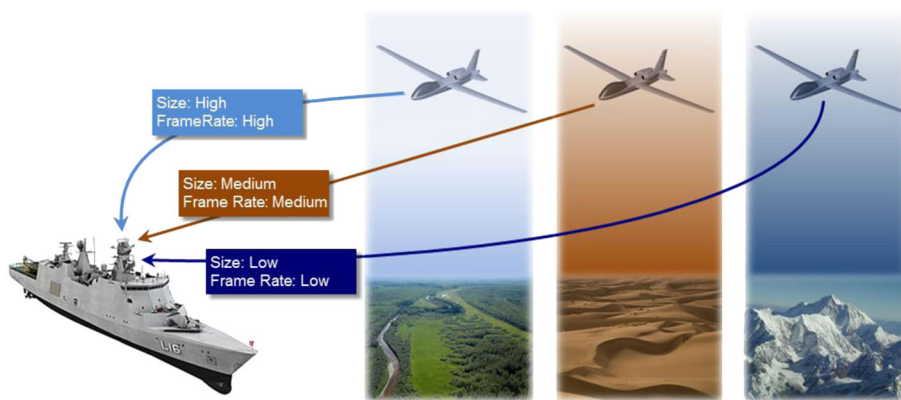


Figure 2. QoS Issue in the UAV Image Transmission Scenario.

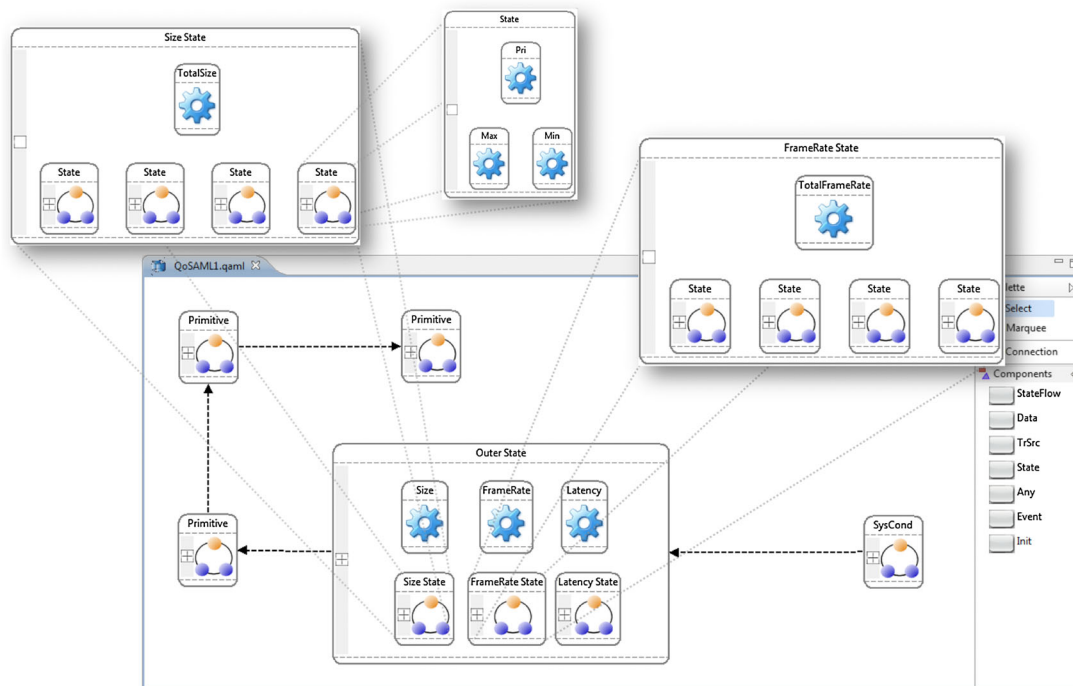


Figure 3. QoSAML Model.

translators have been developed in previous efforts that generate run-time conditions (expressed in the Contract Description Language [21]) from the QoSAML models automatically, which can be integrated into the run-time kernel of the system [18].

2.2.1. Maintaining nonfunctional properties in QoS Adaption Modeling Language. The model in Figure 3 is not complete, because the transitions between different states have not been specified. A transition connects from a source state to a target state, representing how a control parameter can be changed and adapted. To specify the transitions, there are two different strategies that can be used, as illustrated in Figure 4. The left side of the figure specifies a protocol that exhausts the effect of one independent parameter (*Size*) before attempting to adjust another independent parameter (*FrameRate*). In other words, the *FrameRate* parameter has a higher priority so that it is not reduced or increased until there is no further reduction or increment possible to the *Size*. By contrast, the strategy in the right side of the figure is more equitable, with a zig-zag pattern suggesting that the reduction or the increment of one parameter is staggered with the reduction or the increment of another [18]. From this scenario, it can be seen that weaving the strategy protocols becomes a challenging task when more control parameters are involved, or a large number of intermediate states are included in the state.

As the second example for this paper, we choose the task of weaving the priority exhaustive protocol to the QoSAML model, which is defined as follows:

Example 2.2: In a given state machine, for each pair of two states included in the state machine, if their priority data are less than 5, and if the priority data of one state is one less than the other, add a transition between the two states from the state with the lower priority (SourceState) to the one with the higher priority (TargetState). In addition, set up the attributes for the transition: the Guard of the transition should be given from the users input, and the Action of the transition should be in the format of ‘ControlParameterName = (SourceState.Max + TargetState.Max) / 2.’

For example, Figure 5 shows the model after applying the priority exhaustive protocol to *Size State* and *FrameRate State*. The challenges of this task result from locating all pairs of the states that consist of the qualified priority data (i.e., $SourceState.Priority = TargetState.Priority - 1$, $SourceState.Priority < 5$, $TargetState.Priority < 5$), performing the repeated computation to obtain the average value from the two Max data, as well as enabling user input.

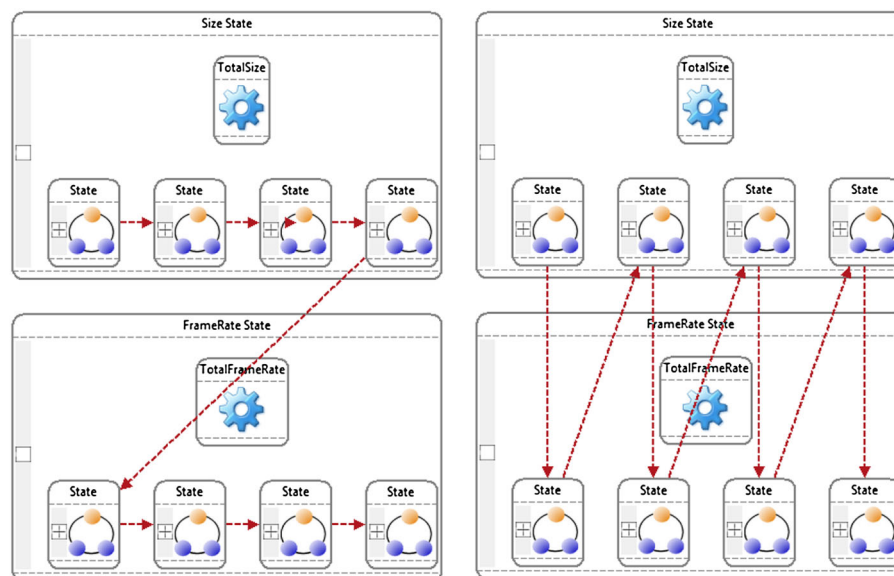


Figure 4. Two state transition protocols to adapt to environment—Priority Exhaustive (left) and Zig-zag (right).

3. DEMONSTRATION-BASED NONFUNCTIONAL PROPERTY MAINTENANCE USING MODEL TRANSFORMATION BY DEMONSTRATION

Model Transformation By Demonstration (MTBD) is our solution that enables general end users to evolve and maintain nonfunctional properties of models, without knowing programming languages or metamodel definitions. MTBD is a new model transformation approach that was initially designed to handle model refactoring problems [31]. The idea of MTBD is to offer an alternative to writing model transformation rules manually, such that users are asked to demonstrate how the model transformation should be performed by directly editing (e.g., add, delete, connect, and update) the model instance to simulate the model transformation process step-by-step. The user evolves and maintains a source model to the target model during the demonstration process. A recording and inference engine captures all user operations and infers a user's intention in a model transformation task. A transformation pattern is generated from the inference, which can be reused and executed in any model instance at any time to carry out the same model transformation process.

We have extended the MTBD concept to nonfunctional property maintenance task, so that it can be used to demonstrate the specific maintenance process and generate a reusable pattern to automatically realize the process to the rest of the model, as well as other model instances.

Figure 6 is an overview of the approach. In any nonfunctional property maintenance task, two fundamental specifications are required to precisely define a task. First, we need to specify the 'where' (i.e., the locations in the models where the nonfunctional property or behavior emerges). The desired locations can be specified through constraints on its structure (i.e., the presence of required model elements and connections) and attributes (i.e., the model elements and connections with qualified attribute values). Second, we need to specify the 'what' (i.e., the specific composition of the nonfunctional properties or behavior). As shown in the two examples, a nonfunctional property can contain various types of actions ranging from adding new elements or connections (removing and modifying elements and connections also happens in some tasks) to conducting computations, asking for user input, and initializing attributes. Using the MTBD approach, users can simply select one of the qualified locations, and then perform the necessary nonfunctional property maintenance tasks on the specific location. At the same time, the recording engine captures all the information about the demonstration, and later generalizes an initial pattern from the user's

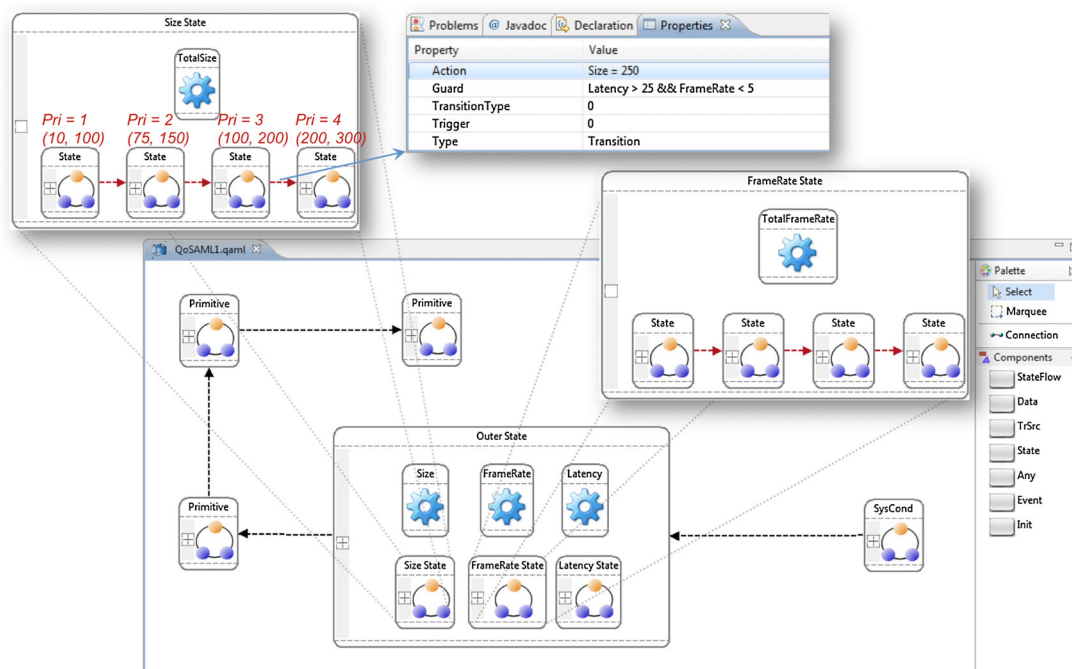


Figure 5. A QoSAML model after applying the priority exhaustive protocol.

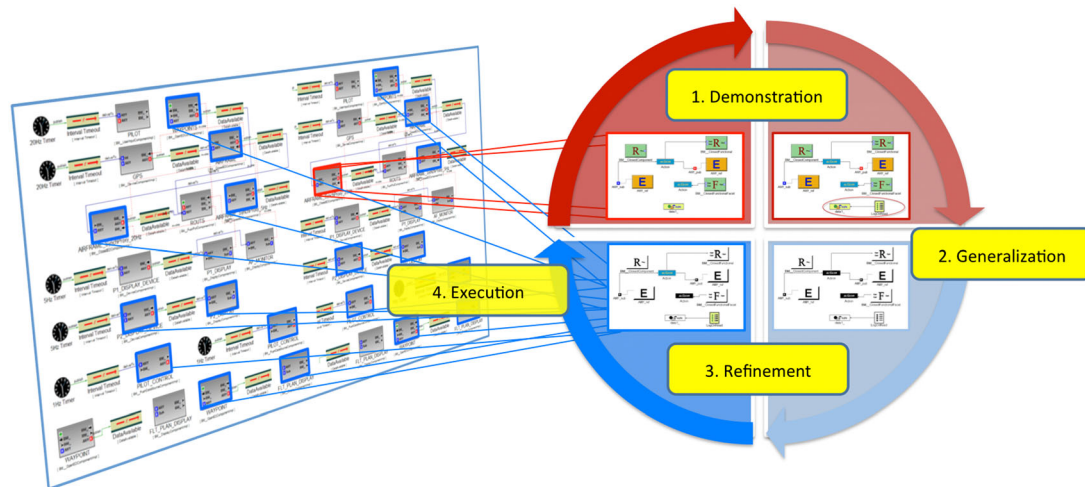


Figure 6. High-level overview of using Model Transformation By Demonstration to address nonfunctional property maintenance.

demonstrated behavior. Because the initial pattern may not reflect the precise specification about the locations, refinements are enabled in the extended MTBD to allow a user to provide more detailed feedback in order to restrict the pattern in terms of the structure and attributes. The pattern is finalized after the refinement, and can be executed and applied to the whole model instance to realize the complete weaving process.

In Section 3.1, we explain the approach with more implementation details about each step, along with solving the motivating Example 2.1. Section 3.2 continues to illustrate the approach by presenting the solution to motivating Example 2.2.

3.1. MTBD in action: specifying nonfunctional concerns in the ESML using MTBD

The main steps of the extended MTBD approach for nonfunctional property maintenance are shown in Figure 7. To better describe the idea, we use Example 2.1 to explain the specific processes.

3.1.1. Step 1—User demonstration and recording. Users first give a demonstration by locating one of the correct places in the model where nonfunctional properties are needed, and directly editing a model instance (e.g., add a new model element or connection, modify the attribute of a model element, and connect two model elements) to simulate the maintenance task. During the demonstration, users are expected to perform operations not only on model elements and connections but also on their

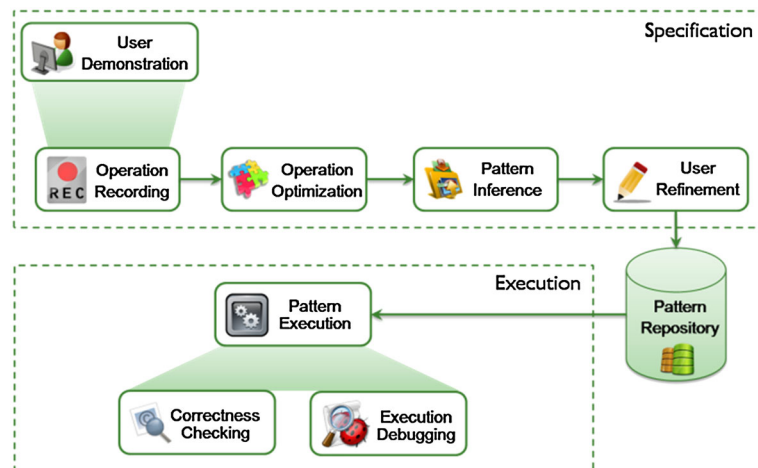


Figure 7. The implementation of Model Transformation By Demonstration (initially described in [31]).

attributes, so that the attribute composition can be supported. An attribute refactoring editor has been developed to enable users to access all the attributes in the current model editor and specify the desired computation (e.g., string and arithmetic computation). It also provides a mechanism for users to create a temporary data pair (i.e., a variable name plus its value), and use this data to simulate user input data during the whole demonstration. At the same time, an event listener has been developed to monitor all the operations occurring in the model editor and collect the information for each operation in sequence.

In Example 2.1, the task is to add a *Logging* mechanism to each *Data* in the *Component* whose name ends with ‘*impl*’; there also must be an *Action* in the *Component*. Therefore, the demonstration is performed on a selected *Component* (*BM_DeviceComponentImpl*) that qualifies these conditions as shown in the left of Figure 8.

With a desired location, the demonstration is straightforward. A user can add a new *Logging* element in the *Component*, and set up its type attribute as ‘*LogOnRead*’, after which a connection can be made in the demonstration between the new *Logging* element to a *Data* element. List 1 shows all the operations performed during the demonstration, and the model after the demonstration is presented in the right of Figure 8. The demonstration only needs to be carried out for one of the *Data* elements, because the execution engine can automatically match all the other locations.

List 1. Operations performed for Example 2.1 in the demonstration

Sequence	Operation Performed
1	Add a <i>Logging</i> in <i>ESMLRoot.BM_DeviceComponentImpl</i>
2	Set <i>Logging.LogType</i> = ‘ <i>LogOnRead</i> ’
3	Connect <i>Logging</i> to <i>data1_</i>

3.1.2. Step 2—Operation optimization. The list of recorded operations indicates how a nonfunctional property should be composed in the base model. However, there is no guarantee that all operations in the demonstration are meaningful. Users may perform useless or inefficient operations during the demonstration. For instance, without a careful design, it is possible that a user first adds a new element and modifies its attributes, and then deletes it in another operation later, with the result being that all the operations regarding this element contradict each other and do not take effect in the transformation process. Thus, the operations are meaningless, even though they do not lead to an incorrect demonstration. Thus, after the demonstration, the engine optimizes the recorded operations to eliminate any meaningless actions. The optimization algorithm is specified in [31]. The operations in List 1 are all meaningful and are unchanged after this step.

3.1.3. Step 3—Pattern inference. With an optimized list of recorded operations, the transformation can be inferred by generalizing the behavior in the demonstration. Because the MTBD approach does not rely on any model transformation language, it is not necessary to generate specific transformation rules, although that is possible. Instead, we generate a transformation pattern, which summarizes the

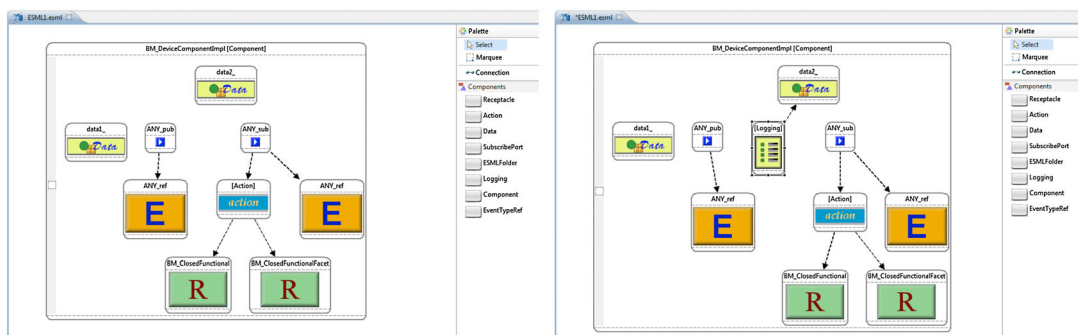


Figure 8. A selected Component before (left) and after (right) the demonstration of adding a logging mechanism.

precondition of a transformation (i.e., where a nonfunctional property emerges) and the actions needed in a transformation (i.e., how a nonfunctional property should be maintained in this location).

As the initial generalized pattern, the precondition in this pattern is called the minimum precondition, which means that it only ensures that all the operations can be executed correctly with the minimum required model elements and connections. As shown in Figure 9, the precondition in this pattern defines that there must be a *Component* available, and in this *Component*, a *Data* element must exist. Obviously, this precondition provides sufficient operands for the operation actions to be executed correctly. However, in many cases, the generalized precondition is not restrictive and precise enough because of the limited expressiveness of the user demonstration. For example, although an *Action* is included in the *Component* in the demonstration, it is not incorporated in the generalized precondition, because this element was never used in any of the user operations. Similarly, we intentionally chose the *Component* called '*BM_DeviceComponentImpl*' as the desired location of our demonstration. The constraint on the name of this component is not inferred. Therefore, to provide further constraints on the precondition, more feedback is needed from users.

3.1.4. Step 4—User refinement. The inaccuracy of the initial pattern resulting from the limitation on the expressiveness of the user demonstration is a main issue undermining the applicability of MTBD to the maintenance tasks of nonfunctional properties, because it is often essential to accurately specify the exact locations of applying a nonfunctional property based on precise constraints. In order to address this problem, we extended the original version of MTBD to permit users to refine the inferred transformation by providing more feedback for the precondition of the desired transformation scenario from two perspectives—structure and attributes. For instance, users can restrict the precondition by selecting and confirming extra model elements or connections in the model editors that must be included in the pattern. A new type of operation—*Confirm Containment*—is implemented in the editor for this purpose, which allows users to right-click on any model element in the editor and confirm to the engine that it should be contained in the final precondition. In addition, the precondition related with the implicit structure is also supported, such as replacing element *A* only if *A* has no incoming or outgoing connections. The refinement on the attributes can be realized by choosing the element in the demonstration and typing the specific conditions (e.g., add new element *B* in *C* only when the attribute value of *C* is greater than 200). All the user refinements are still performed at the model instance level without being aware of the metamodel definitions, after which a transformation pattern will be finalized and stored in the pattern repository for future use.

Reconsidering Example 2.1, two more user refinement operations are needed in this step, as shown in List 2.

List 2. Refinement operations performed for Example 2.1 in the demonstration

Sequence	Operation Performed
4	Confirm the containment of <i>ESMLRoot.BM_DeviceComponentImpl.Action</i>
5	Specify precondition: <i>ESMLRoot.BM_DeviceComponentImpl.name.endsWith('Impl')</i>

Figure 10 shows the final transformation pattern after user refinement (the bold rectangle outline represents the attachment of preconditions on its attributes). The generated pattern will be stored in the repository for future reuse.

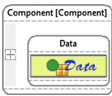
Precondition		Actions
elem1.elem2		1. Add Logging in elem2
elem1.elem2.elem3		2. Set up Logging.LogType = "LogOnRead"
elem1: ESMLRoot		3. Connect Logging to elem3
elem2: Component		
elem3: Data		

Figure 9. The initial generalized pattern.

3.1.5. Step 5—Pattern execution. The final generated patterns can be executed on any model instances. Users can select a specific part of the model to apply the pattern, or by default apply the pattern on the whole model. Because a pattern consists of the precondition and the transformation actions, the execution starts with matching the precondition in the new model instance and then carrying out the transformation actions on the matched locations of the model. The MTBD engine has also been updated to validate the preconditions refined by users in the previous step. The execution can be carried out in any selected locations in the model or to the whole model by default. The execution engine applies a backtracking algorithm that works similar to graph matching, which can be found in [31]. The selected existing model elements in a model instance serve as a candidate pool, and the precondition is matched using subsets of the pool to find the satisfied locations based on both the structural and attribute constraints.

Users can select the pattern in the execution controller dialog to execute an inferred transformation pattern from the repository. Users can select multiple patterns to execute in sequence, which is particularly useful when a model transformation task is divided by subtasks and specified by different demonstrations. In addition, the total times for executing the selected pattern(s) can be specified, because in some use cases (e.g., model scalability), a transformation pattern(s) needs to be executed multiple times to transform the model to a specific state and configuration.

Executing the generated pattern from Step 4 leads to the *Logging* element added to all the locations that qualify the conditions.

3.1.6. Step 6—Correctness checking and debugging. Although the location matching the precondition guarantees that all transformation actions can be executed with the necessary operands, it does not ensure that executing them will not violate the syntax, semantics definitions, or external constraints. Therefore, the execution of each transformation action will be logged and model instance correctness checking on syntax, and static semantics is performed after every execution. If a certain action violates the metamodel definition, all executed actions are undone and the whole transformation is canceled. An execution control component has been developed as part of MTBD to control the number of execution times, and enable the execution of multiple patterns together in sequence. A debugger has been created to enable end users to track the execution of the transformation pattern without being exposed to low-level execution information [32].

3.2. solving nonfunctional property maintenance tasks in QoS Adaption Modeling Language

The demonstration of adding a QoS transition strategy is performed on the selected *Size State*, as shown in Figure 11. Inside the *Size State*, we locate the two *States* with the proper *Pri* values, and perform the operations in List 3. The *Action* attribute configuration by operation 2 is conducted through the attribute refactoring editor, which allows users to access all the model elements and connections in the current model editor and load the attribute values to specify the computation process and evaluate the concrete value. The attribute refactoring editor also provides a mechanism to let users create a temporary data pair, with a given name and a value. The creation of the temporary data is actually used to simulate the user input process, and the data can be used in any attribute configuration and computation process through the entire demonstration. The creation of the temporary data will be generalized as a user input action and will display an input box when the final pattern is executed.

List 3. Operations performed for *Example 2.2* in the demonstration

Sequence	Operation Performed
1	Add a <i>Transition</i> between <i>QoSAMLRoot.OuterState.SizeState.State1</i> and <i>QoSAMLRoot.OuterState.SizeState.State2</i>
2	Set <i>Transition.Action</i> = <i>QoSAMLRoot.OuterState.SizeState.name</i> + ' = ' + (<i>QoSAMLRoot.OuterState.SizeState.State1.Max.value</i> + <i>QoSAMLRoot.OuterState.SizeState.State2.Max.value</i>) / 2 = ' <i>Size = 125</i> '
3	Create a temporary data pair (Name: <i>guard</i> , Value: ' <i>Latency > 25 && FrameRate < 5</i> ') Set <i>Transition.Guard</i> = <i>guard.value</i> = ' <i>Latency > 25 && FrameRate < 5</i> '
4	

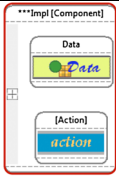
Precondition		Actions
elem1.elem2 (name.endsWith("Impl")) elem1.elem2.elem3 elem1.elem2.elem4		1. Add Logging in elem2 2. Set up Logging.LogType = "LogOnRead" 3. Connect Logging to elem3
elem1: ESMLRoot elem2: Component elem3: Data elem4: Action		

Figure 10. The final generated pattern after user refinement.

The initial pattern generalized from the demonstration is shown in Figure 12. Similar to Example 2.1, the precondition is not accurate enough, because the relationship between the two *Pri* values are not reflected in the demonstration. Moreover, although the two *Max* elements are included in the pattern, they share the same type as *Min* and *Pri* (their meta type are all *Data*), the consequence being that it is possible that the execution engine incorrectly uses *Min* or *Pri* to calculate the average value, or use *Max* or *Min* to compare the *Pri* relationship. Thus, it is necessary to further restrict the *Data* involved in the pattern with their names. The following operations in List 4 are performed in the user refinement step. Operations 5–10 confirm the required *Pri* data elements and their relationship, whereas operations 11–12 ensure that *Max* data elements exist in the two *States*. Figure 13 shows the final generated transformation pattern.

List 4. Refinement operations performed for Example 2.2 in the demonstration

Sequence	Operation Performed
5	Confirm the containment of <i>QoSAMLRoot.OuterState.SizeState.State1.Pri</i>
6	Confirm the containment of <i>QoSAMLRoot.OuterState.SizeState.State2.Pri</i>
7	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State1.Pri.name</i> = 'Pri'
8	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State2.Pri.name</i> = 'Pri'
9	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State1.Pri.value</i> = <i>QoSAMLRoot.OuterState.SizeState.State2.Pri.value</i> —1
10	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State2.Pri.value</i> < 5
11	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State1.Max.name</i> == 'Max'
12	Specify precondition <i>QoSAMLRoot.OuterState.SizeState.State2.Max.name</i> == 'Max'

Executing the pattern on any selected *States* will make the execution engine automatically traverse the state and locate all the pairs of included *States* that satisfy the *Pri* relationship constraint and contains the needed *Max* elements. This allows the *Transition* to be added correctly and combined with a user input *Guard* value.

4. COMPARISON BETWEEN MTBD AND TRADITIONAL MODEL TRANSFORMATION LANGUAGES

The MTBD implementation is a plug-in called MT-Scribe that is hosted within the GEMS [15] modeling tool. MT-Scribe is triggered by the end user by pressing a 'record' button in the model editor. Thus, any modeling language defined in GEMS that can be edited in the model editor can

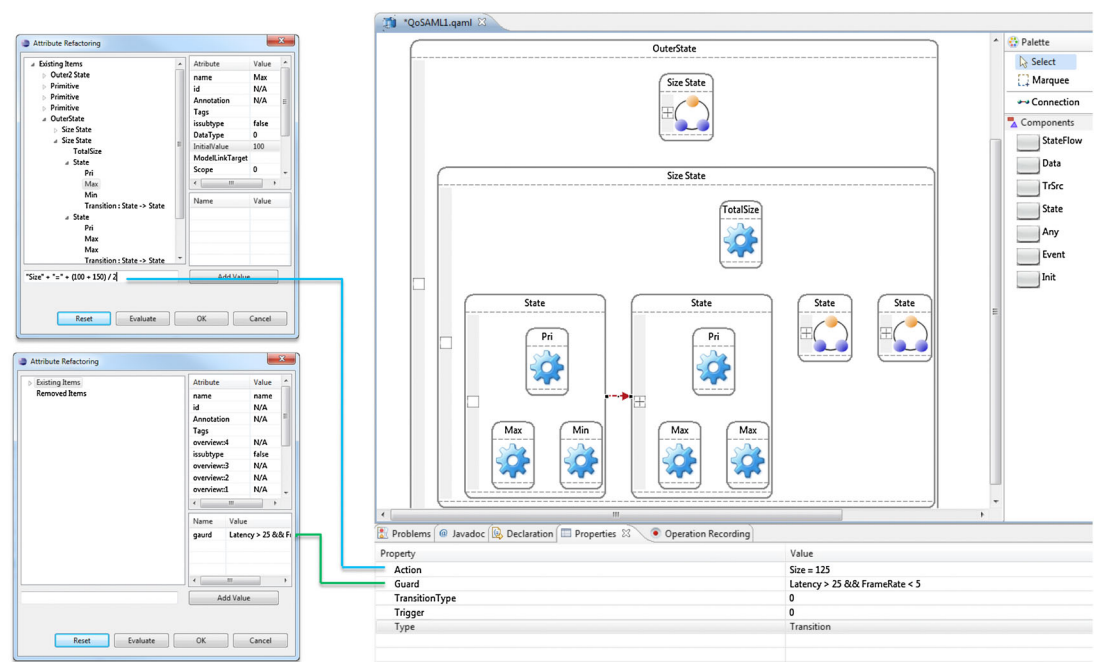


Figure 11. Demonstration of adding a transition and setting up the attributes for the new transition.

apply MTBD to address the maintenance of nonfunctional properties, which means MTBD is a general solution that can be applied across multiple DSMLs.

Using MTBD, users are only involved in editing model instances to demonstrate the nonfunctional property maintenance process and giving feedback after the demonstration. In other words, users only participate in Step 1 and Step 4 (Section 3.1). All of the other procedures (i.e., optimization, inference, generation, execution, and correctness checking) are fully automated. In both Step 1 and Step 4 where users are involved, all of the information exposed to users is at the model instance level, rather than the metamodel level. For instance, the demonstration is carried out using the basic editing operations in the GEMS model editor; the attribute configuration is accomplished using the attribute refactoring editor, which contains all the concrete attribute values from all the available elements and connections. The containment confirmation is simply realized by a one-click operation on the desired model element or connection; and the extra precondition is given using the dialog where users can access all the elements touched in the demonstration and type the constraints directly. The generated patterns are invisible to users (Figures 9, 10, 12, and 13 are presented for the sake of explanation; the generated patterns and actions are not visible to users). Therefore, apart from the basic knowledge about the

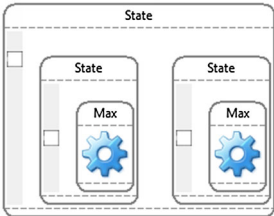
Precondition	
elem1.elem2.elem3.elem4 elem1.elem2.elem3.elem5 elem1.elem2.elem3.elem4.elem6 elem1.elem2.elem3.elem5.elem7	
elem1: QoSAMLRoot	
elem2: State	
elem3: State	
elem4: State	
elem5: State	
elem6: Data elem7: Data	
Actions	
1. Add a <i>Transition</i> between <i>elem4</i> and <i>elem5</i>	
2. Set <i>Transition.Action</i> = <i>elem3.name</i> + "=" + (<i>elem6.value</i> + <i>elem7.value</i>) / 2	
3. Create a data pair (<i>guard</i> = "Latency > 25 && FrameRate < 5")	
4. Set <i>Transition.Guard</i> = <i>guard.value</i>	

Figure 12. The initial generalized pattern.

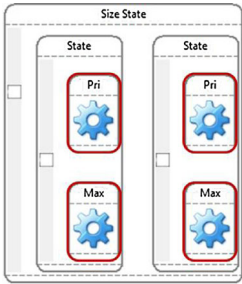
Precondition elem1.elem2.elem3.elem4 elem1.elem2.elem3.elem5 elem1.elem2.elem3.elem4.elem6 (elem6.name == "Max") elem1.elem2.elem3.elem5.elem7 (elem7.name == "Max") elem1.elem2.elem3.elem4.elem8 (elem8.name == "Pri") elem1.elem2.elem3.elem5.elem9 (elem9.name == "Pri") (elem8.value == elem9.value - 1) (elem9.value < 5) <hr/> elem1: QoSAMLRoot elem2: State elem3: State elem4: State elem5: State elem6: Data elem7: Data elem8: Data elem9: Data	
Actions 5. Add a Transition between elem4 and elem5 6. Set Transition.Action = elem3.name + "=" + (elem6.value + elem7.value) / 2 7. Create a data pair (guard = "Latency > 25 && FrameRate < 5") 8. Set Transition.Guard = guard.value	

Figure 13. The final generated pattern after user refinement.

domain, users are isolated from the formal and abstract metamodel definitions and implementation details. Furthermore, no model transformation languages and tools are used in the implementation of MTBD. Thus, users are completely isolated from knowing model transformation languages or programming language concepts.

To compare MTBD to traditional usage of transformation languages, we considered MTBD with the actual transformation rules used to specify the same transformation tasks for the two motivating examples [17, 18]. These transformation rules were used by C-SAW (Constraint-Specification Aspect Weaver) [17, 25], as shown in Figures 14 and 15. We compare the two approaches by analyzing the specific implementation differences, as well as some general development summaries in Section 4.1 and 4.2. As a special type of MTLs, we discuss the differences and advantages of MTBD over graphical MTLs in Section 4.3.

4.1. Comparison on specific implementation differences between Model Transformation By Demonstration and Constraint-Specification Aspect Weaver

Identifying the desired locations of the nonfunctional properties that must be maintained is simplified using MTBD (i.e., a user just selects an example of the specific location using an instance model). In comparison, with the formally specified transformation rules sent to C-SAW, the location of the nonfunctional properties is defined by extended OCL constraints (e.g., `forAll()`, `select()`) together with APIs provided in the transformation language (e.g., `models("State")`, `atoms()`). The process using a traditional transformation language becomes more complex when the different APIs are called and used together in a single statement. By contrast, the main location specification in MTBD is automatically handled in the demonstration process. It is the MT-Scribe recording engine that detects the location of the operation occurring and generalizes the location context information, so that users focus on selecting a desired location without being aware of a generalized locating process.

From our experience, the specific constraints on the preconditions using MTBD are more intuitive and direct than writing formal transformation rules. By selecting and clicking on the desired model elements or connections, constraints on the structure can be specified. The location and selection of the attribute in MTBD is realized by clicking on the element in the precondition specification dialog and providing simple expressions based on the instance model. However, in C-SAW, OCL expressions and condition statements need to be applied (e.g., `if ...select(m | m.kindOf() ==`

```

defines Start, FindData1, AddLog;

strategy FindData1()
{
    Atoms()->select(a | a.kindOf() == "Data")->AddLog();
}

strategy AddLog()
{
    declare parentModel : model;
    declare dataAtom, logAtom : atom;

    dataAtom := self;
    parentModel := parent();

    if (parentModel.atoms("")->select(m | m.kindOf() == "Action")->size() >= 1)
    then
        logAtom := parentModel.addAtom("Log", "LogOnRead");
        logAtom.setAttribute("Kind", "On Read");
        parentModel.addConnection("AddLog", logAtom, dataAtom);
    endif;
}

strategy Start()
{
    rootFolder().findFolder("ComponentTypes").models().select(m | m.name()
        .endsWith("impl"))->FindData1();
}

```

Figure 14. The Enterprise Control Language code to implement Example 2.1 (adapted from [25]).

"Action")->size() >= 1). When it comes to defining the precondition on attribute values, we believe that MTBD is simpler than using conditional statements that access the modeling tool APIs. For example, the following is an expression that would be needed in a typical model transformation rule using the traditional approach: `findAtom("Priority").`

`findAttributeNode("InitialValue").getInt(pri);`

Nonfunctional property maintenance and composition are implemented using model manipulation APIs in C-SAW (e.g., `parent().addConnection("Transition", "Transition", "Transition", endID, prevID), aConnection.addAttribute("Guard", guard)`), whereas using MTBD, the composition process is demonstrated using the basic operations in the model editor (i.e., add, delete, and update attributes).

4.2. Comparison on development efforts between Model Transformation By Demonstration and Constraint-Specification Aspect Weaver

We have not performed a comprehensive formal user study on the comparison between MTBD and the use of formal transformation languages. However, we have made several general summary observations related to development using the two approaches. Table I lists the effort required to create the deliverable transformation artifact, indicating that with the demonstration approach, only a small number of operations are needed using MTBD to accomplish the exact same tasks that were carried out by writing dozens of lines of transformation code. To better measure the effort, we asked two users to accomplish the two tasks using these two approaches (one user applied MTBD, whereas the other one wrote model transformation rules in C-SAW) so that we could measure the time spent using each approach. The user who used MTBD was a trained MTBD user, whereas the other user who wrote the transformation rules is one of the C-SAW co-authors. The results in Table II suggest that using MTBD to solve these two examples was much faster than writing C-SAW rules. We believe that MTBD and the associated demonstrations enable the users to test the patterns and make modifications more quickly than using equivalent formal model transformations. For instance, in order to determine if a generated pattern is too general or not, users just need to execute the pattern and test the result—all within the context of the model instances. In some cases,

```

defines AddTransition, FindConnectingState, ApplyTransitions;

strategy AddTransition(stateName, prevID, guard : string; prevPri : integer)
{
    declare pri, minVal, maxVal, avgVal : integer;
    declare endID : string;
    declare aConnection : node;
    findAtom("Priority").findAttributeNode("InitialValue").getInt(pri);
    if(pri == prevPri + 1)
    then
        getID(endID);
        findAtom("Min").findAttributeNode("InitialValue").getInt(minVal);
        findAtom("Max").findAttributeNode("InitialValue").getInt(maxVal);
        avgVal := (minVal + maxVal) / 2;
        <<CComBSTR action(stateName);    // exit statement to C++
        action.Append("="+XMLParser::itos(avgVal)); >>
        aConnection :=
            parent().addConnection("Transition", "Transition", "Transition",
                                   endID, prevID);
        aConnection.addAttribute("Guard", guard);
        aConnection.addAttribute("Action", action);
    endif;
}

strategy FindConnectingState(stateName, guard : string)
{
    declare pri : integer;
    declare startID : string;
    findAtom("Priority").findAttributeNode("InitialValue").getInt(pri);
    getID(startID);
    if(pri < 5)
    then
        parent().models("State")->
            forAll(AddTransition(stateName, startID, guard, pri));
    endif;
}

strategy ApplyTransitions(stateName, guard : string)
{
    declare theModel : node;
    theModel := findModel(stateName);
    theModel.models("State")->forAll(FindConnectingState(stateName, guard));
}

```

Figure 15. The Enterprise Control Language code to implement Example 2.2 (adapted from [18]).

modifications have to be made to fix demonstration errors followed by re-executing the new pattern. We have observed that MTBD shortens these development iterations.

Table III describes the prerequisite knowledge for each approach. For MTBD, the development environment is integrated in the modeling tool itself, so users only need to know a few additional buttons to use MTBD. The demonstration is performed using the same model editing operations, so the

Table I. The comparison of development effort to solve the two motivating examples.

Example	MTBD	C-SAW Rules
Example 2.1	three editing operations two precondition refinement (one mouse-click on one element, one string expression constraint on an attribute)	23 SLOC
Example 2.2	four editing operations eight precondition refinement (two mouse-click on two elements, four string expression constraints on four attributes, two numerical constraints on two attributes)	40 SLOC

Table II. The comparison of development time to solve the two motivating examples.

Example	MTBD	C-SAW Rules
Example 2.1	30 seconds for demonstration 1.5 minute for refinement	20 minutes
Example 2.2	40 seconds for demonstration 4 minutes for refinement	35 minutes

only new things to learn are the three types of UIs for inputting pattern refinements. For C-SAW, the development environment is a plug-in to Generic Modeling Environment, and it requires a regular textual editor to specify the transformation rules. Using a formal transformation language like C-SAW requires learning the new syntax and semantics of the language, including a number of new keywords, program structures, and functions. Finally, we identified the possible errors that could happen in the development process using the two approaches. For MTBD, users might perform incorrect editing operations or give invalid expressions in the precondition specification and refinement. The base pattern and minimum preconditions are automatically generated in the background, such that user mistakes are minimized in MTBD. For C-SAW or most other textual or visual transformation programming languages, the transformation rules could suffer from general programming issues, such as syntax and logic errors stemming from a misunderstanding of the transformation language or its application (Table IV).

4.3. Comparison with graphical Model Transformation Languages

Besides the textual C-SAW transformation comparison, MTBD also shows advantages over general graphical MTLs (e.g., Graph Rewriting and Transformation Language [2]). Graphical MTLs specify the transformation rules by configuring the source model in the left side and the target model in the right side. The transformation engine matches the source model and replaces it with the target model. Compared with the textual MTLs, graphical MTLs provide a more visual and straightforward approach to specify the rules. However, in order to specify a transformation rule using graphical MTLs, users must map the concrete model elements with the concrete syntax of a certain domain to the syntax of the graphical MTL they choose. Depending on the specific domain and MTLs, the syntax and semantics gap of converting the desired model transformation scenario into a rule varies. Additionally, the rules are still defined using metamodel definitions even though they are represented graphically. Similar to graphical MTLs, the execution engine of MTBD applies a similar semantics to match a precondition and carry out the execution process, but MTBD hides the left side and right side rules (the right side rules in MTBD are in the format of an operational change list as described in Section 3) and avoids the gap by the demonstration process. Users

Table III. The comparison of prerequisite knowledge to solve the two motivating examples.

MTBD	C-SAW
1. MTBD development environment 2. How to demonstrate by model editing 3. How to perform refinement (i.e., the usage of three types of refinement UIs)	1. C-SAW syntax and semantics including 12 keywords, five main program structures, and four functions. 2. C-SAW development environment

Table IV. The comparison of possible errors to solve the two motivating examples.

MTBD	C-SAW
1. Demonstrate with the wrong editing operation 2. The incorrect specification on the attribute precondition expressions	1. General programming language syntax error. 2. Semantics error or bugs on the transformation logic

specify the transformation on concrete model instances by directly editing, without having to convert the idea into another syntax format.

Apart from the general graphical MTLs, some innovative new graphical MTLs (e.g., Modeling Aspects Using a Transformation Approach (MATA) [36]) provide an improved user experience and ease of use by applying concrete syntax to specify the rules without the knowledge of metamodel definitions, as well as combining the left side rule and right side rule together into a single view. This type of graphical MTLs shows similar user experience as MTBD. However, there are still some fundamental differences between the two approaches, as highlighted in the remaining paragraphs of this section.

First, MTBD is not designed as a language or programming language but as a new user experience and interface integrated with the model editing environment to perform model transformations. Thus, programming concepts are avoided in the MTBD process. Users demonstrate transformations rather than writing programmatic rules to specify them. The underlying tool takes care of deriving the programmatic expressions to generalize the demonstrated transformation. Some of the expressions used in the precondition specification are programmatic in nature, but have no conditional logic, loops, or complex expressions, and their usage is minimized through intuitive UIs (e.g., dialog showing the related model element and connections) to assist users' specification. By comparison, graphical MTLs such as MATA are still programming languages based on a grammar or specific metamodel (e.g., UML), with its own syntax and semantics; that is, users still need to have the basic knowledge about a programming language or graph transformations to correctly utilize the language. For many domains, users do not have sufficient experience with programming language or graph transformation concepts to write transformation rules.

Furthermore, MTBD is designed to be at a higher level of abstraction than the graphical MTLs. This results in the fact that a lot of low-level implementation details are hidden (e.g., patterns are generated and executed based on graph theories in MTBD, but are invisible to end users). On the other hand, the higher level of abstraction actually means that not all possible transformations can be expressed. We did not design MTBD to support all types of model transformation and all possible nonfunctional aspect modeling activities. A few of the tasks that can be specified using regular MTLs cannot be directly demonstrated using MTBD (e.g., using max/min in precondition specification, nonfunctional properties weaving based on multiple metamodels). However, we target the most practical and commonly used transformation scenarios that we have discovered across our general work in model transformation across different domains.

Finally, we would also like to highlight the fact that the target user group of MTBD is a little different from the general MTLs. We expect the main users of MTBD to be domain experts using DSMLs to build models, who do not necessarily have any programming or graph transformation background. For computer scientists or programmers, we believe that they may find it more efficient and comfortable to write specific rules using MTLs to accomplish certain tasks rather than using a demonstration-based approach. However, for users without a programming background, MTBD offers a feasible alternative for them to specify model transformation tasks in a potentially simpler way without a steep learning curve. With many DSMLs emerging across diverse domains, this group of end users who are domain experts is becoming increasingly important in the modeling community.

5. RELATED WORK

Nonfunctional system properties in DSMLs have been well-investigated [5–7], ranging from the specification of nonfunctional properties to the measurement and analysis of those nonfunctional properties. A special issue on the topic also appeared recently in the journal of Software and Systems Modeling [8]. Troya *et al.* [33] regard nonfunctional properties as dynamic semantics and specify the properties using domain-specific observers. Boulanger *et al.* [9] propose a multi-view modeling approach to model both functional and nonfunctional viewpoints. Berardinelli *et al.* [4] apply a unifying conceptual framework to model context-aware aspects of mobile software systems. Yrjönen and Merilinnä [38] propose a nonfunctional properties framework to check the satisfaction of specified requirements. Kupfer and Hadar [23] present a methodology to achieve predefined nonfunctional properties by understanding and representing deployment requirements.

Regarding the maintenance of nonfunctional properties, the current main practice is to use MTLs. For example, Paige *et al.* [28] propose to use Epsilon [12] to capture failure behavior. Of course, the traditional model transformation languages and tools are also applicable to support automating the maintenance of nonfunctional properties in DSMLs. The most direct way is to use a general-purpose programming language, such as Java or C++, to access and manipulate the internal structure of a model instance using an API provided by a host modeling tool [30], but this approach is not feasible for end users who do not have programming experience. The power of a transformation is often restricted by the supported API within the modeling tool.

Many modeling tools support importing and exporting model instances in the form of XML Metadata Interchange. Therefore, it is also possible to use existing XML tools (e.g., Extensible Stylesheet Language Transformation [35]) to manage the nonfunctional properties outside of a modeling tool using XML Metadata Interchange as an intermediate representation. However, the problem of Extensible Stylesheet Language Transformation is that it is tightly coupled to XML, requiring experience to define the transformations using concepts at a lower level of abstraction than DSMLs.

Model transformation languages are the most popular approach to support nonfunctional property maintenance, because the process of maintaining the properties in the base model can be considered as a model transformation. Textual MTLs specify the transformation rules using high-level expressions and statements. C-SAW [17], ATLAS Transformation Language (ATL) [20], and Query/View/Transformation [27] are representative examples of powerful MTLs that can be used to manage the weaving of models. Moreover, graphical MTLs (e.g., Graph Rewriting and Transformation Language [1], Visual Automated model Transformations [3]) also exist to convert nonfunctional property maintenance tasks into a graph transformation problem by utilizing graph matching and rewriting techniques. Compared with textual MTLs, it is easier to define specific model patterns using graphs, leading to a simplification of the transformation rules in many cases. However, whether a MTL has a high level of abstraction, graphical or textual, its usage may suffer for domain experts based on the challenges mentioned in Section 1 (i.e., the steep learning curve and need to understand the details of the metamodel).

An alternative approach to simplify the maintenance of nonfunctional properties is to use end user model transformation approaches, as proposed in this paper. To address the challenges inherent from using MTLs, Model Transformation By Example [34, 37] was developed so that instead of writing transformation rules manually, users are asked to build a prototypical set of interrelated mappings between the source and target model instances, and then the metamodel-level transformation rules are semiautomatically generated. This approach simplifies model transformation implementation to some extent, but is not appropriate for nonfunctional property maintenance tasks because: (1) it focuses on direct concept mapping between two different domains rather than maintaining property models within the same domain; and (2) support for attribute transformation is limited, preventing the automatic setup of attributes needed in many maintenance tasks.

Another work has been described by Brosch *et al.* [10], which uses an example-based approach to address model refactoring tasks. Because it supports model transformation within the same domain, it also has potential to be applied in nonfunctional property maintenance scenarios. However, the user feedback step may not be at the proper level of abstraction and complex attribute transformation is also more limited.

Eclipse Modeling Framework Refactor [13] is a new open source component that provides tool support for generating and applying refactoring for models based on EMF. The capabilities of EMF Refactor simplify the process of adding new refactoring functions in EMF editors and supporting model evolution activities, which can be used also to address the maintenance of nonfunctional properties. However, the initial definition of the maintenance rules is based on EMF Tiger [14], a graphical MTL, which is based on writing model transformation rules.

6. CONCLUSION AND FUTURE WORK

In this paper, we presented an end-user approach to enable general users (e.g., domain experts or nonprogrammers) to manage the nonfunctional properties in different DSMLs without knowing model transformation languages or metamodel definitions. The MTBD approach is based on

demonstrations made in a modeling tool by the end user, which are then generalized. We have shown through two examples how nonfunctional properties can be introduced using demonstrations provided by a domain expert.

As future work, we will investigate using MTBD to support the maintenance of nonfunctional properties across different domains and metamodels (e.g., separation of common nonfunctional properties in a way that can be reused across multiple modeling languages). For instance, the base model may need to be first transformed to a new model that conforms to a different metamodel (e.g., the model specified in QoSAML may need to be transformed to a new QoS modeling language in order to apply some special features that are not available in QoSAML). After this transformation, the nonfunctional properties may then be applied. The current implementation of MTBD cannot handle this type of maintenance because of the nature of supporting only endogenous model transformation in MTBD. Extending the MTBD idea to exogenous transformation will enlarge its applications in practice.

Because of the higher level of abstraction, which allows domain experts to manage nonfunctional properties, MTBD is less powerful and expressive than a well-defined transformation language. Some functions (e.g., *getMax[]*, *getMin[]*) can be simply realized using function calls or a customized part of code, which can be very challenging to demonstrate. Therefore, in order to enhance the functionality of MTBD, those commonly used functions in the maintenance of nonfunctional properties will be identified and integrated in the demonstration and user refinement steps. Furthermore, there are also some features that can improve the demonstration user experience. For instance, when confirming a containment of an element in the user refinement step, it would be a useful feature to be able to automatically include the basic constraints on not only the element type but also the attributes (e.g., name).

ACKNOWLEDGEMENT

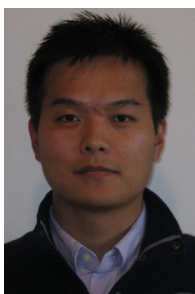
This work is supported by NSF CAREER award CCF-1052616.

REFERENCES

1. Agrawal A, Karsai G, Lédeczi Á. An end-to-end domain-driven software development framework. International Conference on Object-Oriented Programming, Systems, Languages, and Applications - Domain-driven Track, Anaheim, CA, October 2003; 8–15.
2. Balasubramanian D, Narayanan A, Buskirk C, Karsai G The Graph Rewriting and Transformation Language: GreAT. *Electronic Communication of the European Association of Software Science and Technology* 2006; **1**:8.
3. Balogh Z, Varró D. Advanced model transformation language constructs in the VIATRA2 Framework. Symposium on Applied Computing (SAC), Dijon, France, April 2006; 1280–1287.
4. Berardinelli L, Cortellessa V, Marco A. An unified approach to model non-functional properties of mobile context-aware software. The 2nd International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages, Denver, CO, October 2009.
5. Bošković M, Gašević D, Pahl C, Schätz B. The 1st international workshop on non-functional system properties in Domain-Specific Modeling Languages. Models in Software Engineering, Springer-Verlag LNCS 5421, 2008; 227–228.
6. Bošković M, Gašević D, Pahl C, Schätz B. The 2nd international workshop on non-functional system properties in Domain-Specific Modeling Languages. Models in Software Engineering, Springer-Verlag LNCS 6002, 2009; 291–295.
7. Bošković M, Cancila D, Pahl C, Schätz B. The 3rd International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages. Models in Software Engineering, Springer-Verlag LNCS 6627, 2010; 308–311.
8. Bošković M, Gašević D, Pahl C, Schätz B. Guest Editorial to the Theme Issue on Non-Functional System Properties in Domain-Specific Modeling Languages. *Software and System Modeling* 2011; **10**(3):283–286.
9. Boulanger F, Jacquet C, Rouis E, Hardebolle C. Modeling Heterogeneous Points of View with ModHel'X. The 2nd International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages, Denver, CO, October 2009.
10. Brosch P, Langer P, Seidl M, Wieland K, Wimmer M, Kappel G, Retschitzegger W, Schwinger W. An Example is Worth a Thousand Words: Composite Operation Modeling By-Example. International Conference on Model Driven Engineering Languages and Systems (MoDELS), Springer-Verlag LNCS 5795, Denver, CO, October 2009; 271–285.
11. Burnett M, Cook C, Rothermel G. End-user Software Engineering. *Communications of the ACM* 2004; **47**(9):53–58.
12. Eclipse Epsilon. <http://www.eclipse.org/gmt/epsilon/> 2010.
13. EMF Refactor. <http://www.mathematik.uni-marburg.de/~swt/modref/> 2010.
14. EMF Tiger. <http://tfs.cs.tu-berlin.de/emftrans/> 2010.
15. Generic Eclipse Modeling System (GEMS). <http://www.eclipse.org/gmt/gems/> 2010.
16. Graphical Modeling Framework (GMF). <http://www.eclipse.org/modeling/gmf/> 2010.

17. Gray J, Lin Y, Zhang J. Automating Change Evolution in Model-Driven Engineering. *IEEE Computer, Special Issue on Model-Driven Engineering* 2006; **39**(2):51–58.
18. Gray J, Neema S, Zhang J, Lin Y, Bapty T, Gokhale A, Schmidt D. Concern Separation for Adaptive QoS Modeling in Distributed Real-Time Embedded Systems. *Behavioral Modeling for Embedded Systems and Technologies. Applications for Design and Implementation*, Idea Group, Chapter 4, 2009; 85–113.
19. Gray J, Tolvanen J, Kelly J, Gokhale A, Neema S, Sprinkle J. Domain-Specific Modeling. *Handbook of Dynamic System Modeling*. CRC Press, Boca Raton, FL, USA, Chapter 7, 7-1 through 7-20 2007.
20. Jouault F, Allilaire F, Bézivin J, Kurtev I. ATL: A Model Transformation Tool. *Science of Computer Programming* 2008; **72**(1/2):31–39.
21. Karr D, Rodrigues C, Loyall J, Schantz R, Krishnamurthy Y, Pyarali I, Schmidt D. Application of the QuO Quality-of-Service Framework to a Distributed Video Application. *International Symposium on Distributed Objects and Applications*, Rome, Italy, 2001; 299–309.
22. Karsai G, Neema S, Sharp D. Model-Driven Architecture for Embedded Software: A Synopsis and an Example. *Science of Computer Programming* 2008; **73**(1):26–38.
23. Kupfer M, Hadar I. Understanding and Representing Deployment Requirements for Achieving Non-Functional System Properties. *The 1st International Workshop on Non-functional System Properties in Domain-Specific Modeling Languages*, Toulouse, France, October 2008.
24. Ledeczi A, Bakay A, Maroti M, Volgyesi P, Nordstrom G, Sprinkle J, Karsai G. Composing Domain-specific Design Environments. *IEEE Computer* 2001; **34**(11):44–51.
25. Lin Y. A Model Transformation Approach to Automated Model Evolution. Ph.D. Thesis, University of Alabama at Birmingham 2007.
26. Object Management Group. Object Constraint Language Specification. http://www.omg.org/technology/documents/modeling_spec_catalog.htm#OCL. 2010.
27. Object Management Group. Revised Submission for MOF 2.0 Query/View/Transformations RFP (ad/2002-04-10). *OMG Document ad/2005-07-01* 2005.
28. Paige R, Rose L, Kolovos D, Brooke P, Ge X. Automated Safety Analysis for Domain-Specific Modeling Languages. *The 1st International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages*, Toulouse, France, October 2008.
29. Schmidt D. Special Issue on Model-Driven Engineering. *IEEE Computer* 2006; **39**(2):25–32.
30. Sendall S, Kozaczynski W. Model transformation - The Heart and Soul of Model-Driven Software Development. *IEEE Software, Special Issue on Model Driven Software Development* 2003; **20**(5):42–45.
31. Sun Y, White J, Gray J. Model Transformation by Demonstration. *Model Driven Engineering Languages and Systems (MoDELS)*. Springer-Verlag LNCS 5795: Denver, CO, 2009; 712–726.
32. Sun Y, Gray J. Model Transformation By Demonstration Debugger: End-User Support for Debugging Model Transformation Execution. *European Conference on Modeling Foundations and Applications (ECMFA)*, Springer-Verlag LNCS, Montpellier, France, July 2013.
33. Troya J, Rivera J, Vallecillo A. On the Specification of Non-Functional Properties of Systems by Observation. *The 2nd International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages*, Denver, CO, October 2009.
34. Varró D, Balogh Z. Automating Model Transformation by Example using Inductive Logic Programming. *Symposium on Applied Computing (SAC)*, Seoul, Korea, March 2007, 978–984.
35. W3C. XSLT Transformation version 1.0. <http://www.w3.org/TR/xslt> 1999.
36. Whittle J, Jayaraman P, Elkhodary A, Moreira A, Araújo J. MATA: A Unified Approach for Composing UML Aspect Models based on Graph Transformation. *Transactions on Aspect-Oriented Software Development* 2009; **VI**(5560):191–237.
37. Wimmer M, Strommer M, Kargl H, Kramler G. Towards Model Transformation Generation By-Example. *Hawaii International Conference on System Sciences (HICSS)*, Big Island, HI, January 2007, 285.
38. Yrjönen A, Merilinnä J. Extending the NFR Framework with Measurable Non-Functional Requirements. *The 2nd International Workshop on Non-Functional System Properties in Domain-Specific Modeling Languages*, Denver, CO, October 2009.

AUTHORS' BIOGRAPHIES



Dr Yu Sun received his PhD from the University of Alabama at Birmingham in 2011. His research focuses on model-driven engineering, and specifically supporting model evolution process in a user-centric approach. He currently works for a startup company—PAR works, as the director of engineering. The company provides high-precision 3D augmented reality technology for mobile platforms.



Jeff Gray is an Associate Professor in the Department of Computer Science at the University of Alabama. He co-directs the Software Modeling Lab, which focuses on research in areas related to model-driven engineering, DSLs, and software maintenance. He also conducts research and outreach in computer science education, with particular emphasis on grades K-12.



Romain Delamare holds a PhD in Software Engineering delivered by the University of Rennes, France. He carried out post-doctoral studies at the University of Alabama. He worked on software testing and aspect-oriented programming, as well as, model-driven engineering. He currently occupies a research engineer position at Tocea, a start-up company focusing on code and software architecture verification.



Dr Benoit Baudry is a research scientist at INRIA. He received a PhD degree in Computer Science from the University of Rennes, France, in 2003. He first worked at CEA (French center for atomic energy) before joining INRIA in 2004. In 2008, he was a visiting researcher in Colorado State University for 1 year. His research interests include software testing and verification, model-driven engineering, and requirements analysis.

He leads the INRIA TRISKELL team, which investigates model-driven engineering and software product lines from requirements to runtime. He is on the steering committee of the IEEE International Conference on Software Testing, Verification and Validation and on the editorial board of the Journal for Software and Systems Modeling and the Journal of Systems and Software.



Dr Jules White is an Assistant Professor of Computer Science in the Department of Electrical Engineering and Computer Science at Vanderbilt University. His research has won four Best Paper Awards. He has also published over 80 papers. Dr White's research focuses on securing, optimizing, and leveraging data from mobile cyber-physical systems. His mobile cyber-physical systems research spans four key focus areas: (i) mobile security and data collection; (ii) high-precision mobile augmented reality; (iii) mobile device and supporting cloud infrastructure power and configuration optimization; and (iv) applications of mobile cyber-physical systems in multidisciplinary domains, including energy-optimized cloud computing, smart grid systems, healthcare/manufacturing security, next generation construction technologies, and citizen science. His research has been transitioned to industry, where it won an Innovation Award at CES 2013, was a finalist for the Technical Achievement Award at SXSW Interactive, and was a top three for mobile in the Accelerator Awards at SXSW 2013.