

Chapter 1

Automatic Refinement Checking for Formal System Models

Julia Seiter, Robert Wille, Ulrich Kühne, and Rolf Drechsler

Abstract For the design of complex systems, formal modelling languages such as UML or SysML find significant attention. The typical model-driven design flow assumes thereby an initial (abstract) model which is iteratively refined to a more precise description. During this process, new errors and inconsistencies might be introduced. In this chapter, we propose an automatic method for verifying the consistency of refinements in UML or SysML. For this purpose, a theoretical foundation is considered from which the corresponding proof obligations are determined. Afterwards, they are encoded as an instance of *satisfiability modulo theories (SMT)* and solved using proper solving engines. The practical use of the proposed method is demonstrated and compared to a previously proposed approach.

1.1 Introduction

Due to the increasing complexity of today's systems and, caused by this, the steady strive of designers and researchers towards higher levels of abstractions, modelling languages such as the *unified modeling language (UML)* [28] or the *Systems Modeling Language (SysML)* [33] together with textual constraints, e.g., provided by the *object constraint language (OCL)* [32] received significant attention in computer-aided design. They allow for a *formal* specification of a system prior to the implementation. Such an initial blueprint can be iteratively refined to a final model to be implemented. The actual implementation is then carried out manually, by using automatic code generation, or a mix of both.

An advantage of using formal descriptions like UML or SysML is that the initial system models can already be subject to (automatic) correctness and plausibility

J. Seiter (✉) • U. Kühne
Institute of Computer Science, University of Bremen, 28359 Bremen, Germany
e-mail: jseiter@informatik.uni-bremen.de; ulrichk@informatik.uni-bremen.de

R. Wille • R. Drechsler
Institute of Computer Science, University of Bremen and Cyber-Physical Systems,
DFKI GmbH, 28359 Bremen, Germany
e-mail: rwille@informatik.uni-bremen.de; drechsle@informatik.uni-bremen.de

checks. By this, inconsistencies and/or errors in the specification can be detected even before a single line of code has been written. For this purpose, several approaches have been introduced [3, 11, 12, 30, 31]. They tackle verification questions such as “Does the conjunction of all constraints still allow the instantiation of a legal system state?” or “Is it possible to reach certain bad states, good states, or deadlocks?”. These verification tasks are typically categorized in terms such as consistency, reachability, or independence [17].

However, these verification techniques are usually carried out on a single model and, hence, are not sufficient for the typical model-driven design flow in which an abstract model is generated first and iteratively refined to a more precise description. Indeed, they enable the detection of errors and inconsistencies in one iteration, but they need to be re-applied in the succeeding iteration even for minor changes. Instead, it is desirable to check whether a refined model is still consistent to the original abstract model. In this way, verification results from abstract models will also be valid for later refined models.

For the creation of software systems, such a refinement process has already been established. Here, frameworks such as the *B-method* [1], *Event-B* [2], and *Z* [34] exist. These methods rely on a rigorous modelling using first-order logic. Extensions, e.g., of *Event-B* to the UML-like *UML-B* or translations of UML to *B* models, are available in [5, 29], respectively. But since the proof obligations for a correct refinement in these frameworks are undecidable in general, usually manual or interactive proofs must be conducted—a time-consuming process which additionally requires a thorough mathematical background.

Hence, *automatic* proof techniques are desired. For this purpose, existing solutions proposed in the context of hardware verification and the design and modelling of reactive systems may be applied. Here, the relation of an implementation and its specification (comparable to a refined and an abstract system) is traditionally described by *simulation relations* on finite state systems (see, e.g., [7, 16, 21, 23]). There exist algorithms for computing such relations [10, 25]. However, since these algorithms operate on explicit state graphs, they do require the consideration of all possible system states and operation calls—*infeasible* for larger designs. A similar difficulty occurs when attempting to automatize the verification process proposed by the *B-method*. In [20], an extension to the tool ProB has been proposed which automatically solves all proof obligations created in the refinement process. However, according to their evaluation, the run-time for the verification grows exponentially.

As a consequence, an alternative solution is proposed in this chapter which exploits the recent accomplishments in the domain of model-based verification (i.e. approaches like [3, 11, 12, 30, 31]) using a symbolic state representation. Based on a theoretical foundation of refinement, we can prove the preservation of safety properties from an abstract model to a more detailed model. In contrast to the existing approaches like in [24], this also includes *non-atomic* refinements, where an abstract operation is replaced by a sequence of refined operations. By this, the consistency of a refined model against an original (abstract) model can be checked automatically.

The remainder of the chapter is structured as follows. The chapter starts with a brief review on models and their notation in Sect. 1.2. Section 1.3 describes the addressed problem which, afterwards, is formalized in Sect. 1.4. The proposed solution is introduced in Sect. 1.5 and its usefulness is demonstrated in Sect. 1.6 where it is applied to several examples and compared to the results from [20]. The chapter is concluded in Sect. 1.8.

1.2 Models and Their Notation

Modelling languages provide a description of a system to be realized, i.e. proper description means to *formally* define the structure and the behaviour of a system. At the same time, implementation details which are not of interest in the early design/specification state remain hidden. In the following, we briefly review the respective description by means of UML and OCL. The approaches proposed in this chapter can be applied to similar modelling languages (e.g. such as SysML) as well.

Definition 1. A *model* is a tuple $m = (C, R)$ composed of a set of classes C and a set of associations R . A *class* $c = (A, O, I) \in C$ of a model m is a tuple composed of attributes A , operations O , and invariants I . An n -ary *association* $r \in R$ of a model m is a tuple $r = (r_{\text{ends}}, r_{\text{mult}})$ with *association ends* $r_{\text{ends}} \in C^n$ for a given set of classes C and *multiplicities* $r_{\text{mult}} \in (\mathbb{N}_0 \times \mathbb{N})^n$ that is defined as a range with a *lower bound* and an *upper bound*.

Example 1. Figure 1.1a shows a model composed of the class Phone which itself is composed of the attributes $A = \{\text{credit}\}$, the operations $O = \{\text{charge}\}$, and the invariant $I = \{i1\}$.

Invariants in the model describe additional constraints which have to be satisfied by each instantiation of the model. For this purpose, textual descriptions provided in OCL can be applied. OCL also allows the specification of the behaviour of operations.

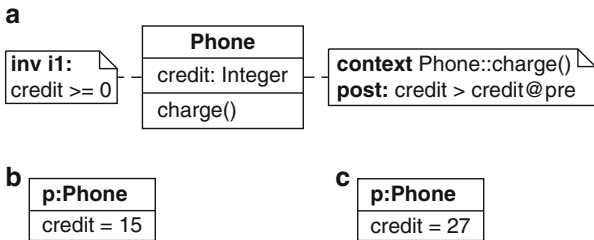


Fig. 1.1 Example of a model and its instantiation. (a) Given model; (b) State σ_0 ; (c) State σ_1

Definition 2. *OCL expressions* Φ are textual constraints over a set of *variables* $V \supseteq A \times R$ composed of the attributes A of the respective classes, but also further (auxiliary) variables. An OCL condition $\varphi \in \Phi$ is defined as a function $\varphi : V \rightarrow \mathbb{B}$. They can be applied to specify the invariants of a class as well as the pre- and post-condition of an operation, i.e. $I \subseteq \Phi$. An *operation* $o \in O$ is defined as a tuple $o = (\triangleleft, \triangleright)$ with pre-condition $\triangleleft \in \Phi$ and post-condition $\triangleright \in \Phi$, respectively. The valid initial assignments of a class are described by a predicate $\text{init} \in \Phi$.

Example 2. In the model from Fig. 1.1a, the invariant $i1$ states that `credit` always has to be greater or equal to 0. The post-condition of the operation `charge` ensures that after invoking the operation, `credit` is increased.

Any instance of a model is called a *system state* and is visualized by an object diagram.

Definition 3. *Object diagrams* represent precise system states in a model. A *system state* is denoted by σ and is composed of *objects*, i.e. instantiations of classes. The attributes of the objects are derived from the classes and assigned precise values. Associations are instantiated as precise *links* between objects.

In order to evaluate a model, it is crucial to particularly consider whether system states are valid or sequences of system states represent valid behaviour. This requires the evaluation of the given OCL expressions.

Definition 4. For a system state σ and an OCL expression φ , the evaluation of φ in σ is denoted by $\varphi(\sigma)$. A system state σ for a model $m = (C, R)$ is called *valid* iff it satisfies all invariants of m , i.e. iff $\bigwedge_{c \in C} I_c(\sigma)$. An operation call is valid iff it transforms a system state σ_t satisfying the pre-condition to a succeeding system state σ_{t+1} satisfying the post-condition,¹ i.e. iff $\triangleleft(\sigma_t)$ and $\triangleright(\sigma_t, \sigma_{t+1})$. A sequence of system states is called *valid*, if all operation calls are valid.

Example 3. Figures 1.1b and c show two valid system states (in terms of object diagrams) for the model from Fig. 1.1a. This is a valid sequence of system states which can be created by calling the operation `charge`.

1.3 Refinement of Models

Using the description means reviewed in the previous section allows for a *formal* specification of a system to be implemented. By this, precise blueprints are available already in the early stages of the design. A rough initial model is thereby created

¹The post-condition is a binary predicate, since it can also depend on the source state, which is expressed using `@pre` in OCL.

first which covers the most important core functionality. Afterwards, a *refinement* process is conducted in which a more precise model of the respective components and operations is created. This refinement process may include

- the addition of new components and relations (i.e. classes and their associations),
- the extension of classes by new attributes,
- the extension of the behavioural description (i.e. the addition of new operations as well as pre- and post-conditions and the strengthening of existing pre- and post-conditions), and
- the extension of the constraints (i.e. the addition of new and the strengthening of existing invariants).

Example 4. Consider the model from Fig. 1.2a representing a simple phone application. It consists of a phone with a credit which can be charged by a corresponding operation. A possible refinement of this model is depicted in Fig. 1.2b. Here, the post-condition of the operation *charge* has been rendered more precise, i.e. a parameter defining the amount of credits to be charged has been added.

Remark. Up to this point, we do not consider the refinement of associations and operation parameters. This includes the type of association and the multiplicities of the associations ends. However, this is not due to a technical limitation of our approach which can easily be extended to further description means. Here, we decided to focus on the refinement of attributes and operations, considering in particular non-atomic refinement, as these are the most important modelling elements in formal system specifications. Other kinds of refinement, e.g. for operation parameters, can be conducted analogously.

In the following, we denote the *abstract model* by m^a and the *refined model* by m^r . A refinement is described by a refinement relation defined as follows:

Definition 5. A refinement relation is a pair $\text{Ref} = (\text{Ref}_\Sigma, \text{Ref}_\Omega)$ with

- Ref_Σ describing the refinement of the states, i.e. Ref_Σ^{-1} is a function mapping a refined state σ^r to its corresponding abstract state σ^a , and
- Ref_Ω describing the refinement of operations, i.e. Ref_Ω is a function mapping an abstract operation o^a to a sequence $o_1^r \cdot o_2^r \cdot \dots \cdot o_k^r \in (O^r)^+$ of refined operations.

Example 5. The refinement from the model in Fig. 1.2a to the model in Fig. 1.2b is described by the relation $\text{Ref}=(\text{Ref}_\Sigma, \text{Ref}_\Omega)$. That is, each state σ^r in the

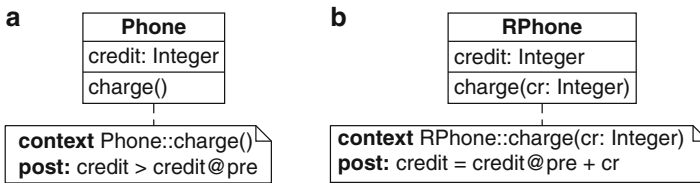


Fig. 1.2 Refinement step. (a) Abstract model; (b) Refined model

refined model (composed of objects from class RPhone) has one corresponding state $\text{Ref}_{\Sigma}^{-1}(\sigma^r) = \sigma^a$ in the abstract model (composed of objects from class Phone) such that $\text{RPhone.credit} = \text{Phone.credit}$. Furthermore, the operation $\text{Phone}::\text{charge}()$ is refined so that $\text{Ref}_{\Omega}(\text{Phone}::\text{charge}()) = \text{RPhone}::\text{charge}(cr)$, i.e. a corresponding operation with an additional parameter.

Adding details step by step—like in the above example—is common practice in model-driven design using UML or SysML. Nevertheless, during this manual process, new errors might be introduced, leading to a refined model that is not consistent with the abstract model any more. In fact, the refinement sketched above contains a serious flaw.

Example 6. The refined model in Fig. 1.2b allows for a behaviour that is not specified by the abstract model. It is possible to assign a value equal to or less than 0 to the parameter cr , so that after calling the operation charge , the value of credit does not change at all or even decreases. This contradicts the behaviour described in the abstract model which only allows for a strict increase of that attribute. As a possible repair of this inconsistency, the precondition $\text{pre}: cr > 0$ could be added to the operation $\text{RPhone}::\text{charge}(cr)$.

In order to identify and fix inconsistencies of the refinement, designers have to intensely check the refined model against the abstract original—often a complicated and cumbersome task which results in a manual and time-consuming procedure. In the worst case, all components, constraints, and possible executions have to be inspected. While this might be feasible for the simple model discussed above, it becomes highly inefficient for larger models. Hence, in the remainder of this chapter we consider the question “How to automatically check whether a refined model m^r is consistent with respect to the originally given abstract model m^a ?”

1.4 Theoretical Foundation

This section formalizes the problem sketched above. For this purpose, we exploit the theoretical foundation of Kripke structures and their concepts of simulation relations. We show how these concepts can be applied for the refinement of system models provided e.g. in UML or SysML. This provides the basis for the proposed solution which is described afterwards in Sect. 1.5.

Since we are considering models mostly in the context of software and hardware systems, we assume bounded data types and a bounded number of instances in the following.² Based on these assumptions, the behaviour of a model can be described as a finite state machine, e.g. a *Kripke structure*.

²This restriction is common in many approaches (e.g. [3, 11, 12, 30, 31]) and also justified by the fact that, eventually, the implemented system will be realized by bounded physical devices anyway.

Definition 6. A Kripke structure is a tuple $\mathcal{K} = (S, S_0, AP, \mathcal{L}, \rightarrow)$ with a finite set of states S , initial states $S_0 \subset S$, a set of atomic propositions AP , a labelling function $\mathcal{L} : S \rightarrow 2^{|AP|}$, and a (left-total) transition relation $\rightarrow \subseteq S \times S$.

Using this formalism, we can define the behaviour of a UML or SysML model and its operations as follows:

Definition 7. A model $m = (C, R)$ induces a Kripke structure $\mathcal{K}_m = (S, S_0, AP, \mathcal{L}, \rightarrow)$ with

- S being the set of all valid system states of $m = (C, R)$,
- S_0 being the set of initial states defined by the predicate `init` (cf. Definition 2), i.e. $S_0 = \{\sigma \in S \mid \text{init}(\sigma)\}$,
- \rightarrow being the transition relation including the identity (i.e. $\sigma \rightarrow \sigma$) as well as all transitions caused by executing operations $o = (\triangleleft, \triangleright) \in O$ of the model (i.e. $\sigma_1 \rightarrow \sigma_2$ with $\triangleleft(\sigma_1)$ and $\triangleright(\sigma_1, \sigma_2)$), and
- AP and \mathcal{L} are defined s.t. \mathcal{L} can be used to retrieve the values of the attributes of σ in the usual bit-vector encoding.

We will write $\sigma_1 \xrightarrow{o} \sigma_2$ to make clear that an operation o transforms a state σ_1 to a state σ_2 .

With this formalization, we can make use of known results for finite and reactive systems. To describe refinements in this domain, *simulation relations* are usually applied for this purpose (see, e.g., [7, 10, 16, 21, 23, 25]). In this chapter, we adapt this concept for the considered formal models. This leads to the following definition of a simulation relation.

Definition 8. Let $\mathcal{A} = (S_{\mathcal{A}}, S_{\mathcal{A}0}, AP_{\mathcal{A}}, \mathcal{L}_{\mathcal{A}}, \rightarrow_{\mathcal{A}})$ be a Kripke structure of an abstract model and $\mathcal{R} = (S_{\mathcal{R}}, S_{\mathcal{R}0}, AP_{\mathcal{R}}, \mathcal{L}_{\mathcal{R}}, \rightarrow_{\mathcal{R}})$ be a Kripke structure of a refined model with $AP_{\mathcal{R}} \supseteq AP_{\mathcal{A}}$. Then, a relation $H \subseteq S_{\mathcal{R}} \times S_{\mathcal{A}}$ is a simulation relation iff

1. all initial states in the refined model have a corresponding initial state in the abstract model, i.e. $\forall s_0 \in S_{\mathcal{R}0} \exists s'_0 \in S_{\mathcal{A}0}$ with $H(s_0, s'_0)$,
2. all states in the refined model are constrained by at least the same propositions as their corresponding abstract state, i.e. $\forall s, s' : H(s, s') \Rightarrow \mathcal{L}_{\mathcal{R}}(s) \cap AP_{\mathcal{A}} = \mathcal{L}_{\mathcal{A}}(s')$, and
3. all possible transitions in the refined model have a corresponding transition in the abstract model leading to a corresponding succeeding state, i.e. $\forall s, s' : H(s, s') \Rightarrow s \rightarrow_{\mathcal{R}} t \Rightarrow \exists t' \in S_{\mathcal{A}}$ s.t. $s' \rightarrow_{\mathcal{A}} t'$ and $H(t, t')$.

We say that \mathcal{R} is simulated by \mathcal{A} (written as $\mathcal{R} \preceq \mathcal{A}$), if there exists a simulation relation.

Example 7. As an illustration of the above definition, Fig. 1.3a shows the general scheme of a transition between states from a refined model (denoted by s and t) and a corresponding transition in an abstract model (from s' to t'). The simulation relation H is indicated by dashed lines. Figure 1.3b on the right shows an example

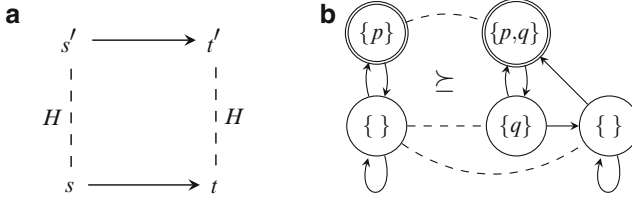


Fig. 1.3 Simulation relation. (a) Correspondence of states; (b) Example for simulation

for two Kripke structures. The abstract model is the one on the left-hand side and simulates the refined model on the right-hand side. Initial states are marked by a double outline. While all corresponding states agree on the atomic proposition p , the refined model has an additional proposition q . It can easily be checked that for each refined transition, there is a corresponding abstract one.

The simulation relation ensures that a refined model is consistent to an abstract system, i.e. whatever the refined system does must be allowed by the abstract system. Besides that, there might be more behaviour allowed in the abstract system than implemented. If we have $\mathcal{R} \leq \mathcal{A}$, then the traces of \mathcal{R} are contained in those of \mathcal{A} . This also means that globally valid properties of \mathcal{A} carry over to \mathcal{R} , as, for example, the non-reachability of bad states. Hence, by proving that the applied refinement Ref (cf. Definition 5) satisfies the properties of a simulation relation H , the consistency of a refined model can be verified.

However, determining a simulation relation requires a strict step-wise correspondence between the transition in the refined model and in the abstract one. But refinements of UML or SysML models often include the replacement of a single abstract operation by a sequence of refined operations (also known as *non-atomic refinement* [6]). In order to formalize this, we need a more flexible relation. This is provided by the notion of *divergence-blind stuttering simulation* (dbs-simulation).

Definition 9. Given two Kripke structures \mathcal{R} and \mathcal{A} with $AP_{\mathcal{R}} \supseteq AP_{\mathcal{A}}$, a relation $H \subseteq S_{\mathcal{R}} \times S_{\mathcal{A}}$ is a divergence blind stuttering simulation (dbs-simulation) iff

1. $\forall s_0 \in S_{\mathcal{R}0} \exists s'_0 \in S_{\mathcal{A}0}$ with $H(s_0, s'_0)$,
2. $\forall s, s' : H(s, s') \Rightarrow \mathcal{L}_{\mathcal{R}}(s) \cap AP_{\mathcal{A}} = \mathcal{L}_{\mathcal{A}}(s')$, and
3. each possible transition in the refined model corresponds to a sequence of 0 or more abstract transitions, i.e. $\forall s, s' : H(s, s')$ and $s \rightarrow_{\mathcal{R}} t$, then there exist $t'_0, t'_1 \dots t'_n$ ($n \geq 0$) such that $s' = t'_0$ and $\forall i < n : t'_i \rightarrow_{\mathcal{A}} t'_{i+1} \wedge H(s, t'_i)$ and $H(s', t'_n)$.

We say that \mathcal{R} is dbs-simulated by \mathcal{A} , written as $\mathcal{R} \leq_{\text{dbs}} \mathcal{A}$, if there exists a dbs-simulation.

Compared to the original simulation relation, this definition is less precise with respect to the *duration* of specific operations. But, it still guarantees that the functional behaviour of the refined model is consistent with the behaviour of the

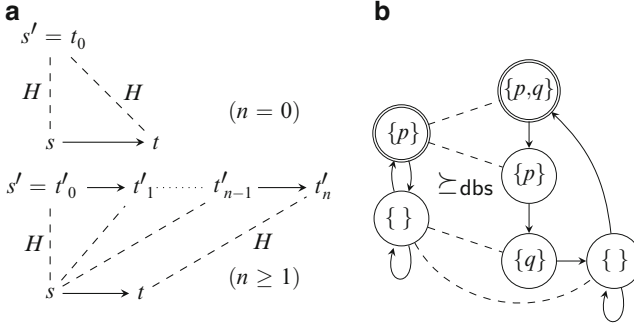


Fig. 1.4 dbs —simulation relation. (a) Correspondence of states; (b) Example for dbs -simulation

abstract model—even in the absence of a (step-wise) one-to-one correspondence of the transitions. In particular, if the properties of the dbs -simulation are satisfied, a bad state unreachable in \mathcal{A} is also unreachable in \mathcal{R} .

Example 8. In Fig. 1.4, the dbs -simulation relation is illustrated. The general scheme of corresponding states and transitions is shown in Fig. 1.4a. In Fig. 1.4b, the abstract model on the left-hand side dbs -simulates the refined model on the right-hand side. Note that the transition from the initial state of the refined model is a *stuttering* transition, since it corresponds to an empty sequence of transitions in the abstract model.

The above definitions provide the formal foundation for consistency checks of refinements. By referring to dbs -simulation, we can preserve safety properties from an abstract model to a refined model. Hence, by proving that the actually applied refinement Ref (c.f. Definition 5) indeed satisfies the properties of a dbs -simulation H (cf. Definition 9), the consistency of the refinement is shown. In the next section, we describe how the refinement for UML or SysML models can efficiently be checked.

1.5 Proposed Solution

In this section, we present the proposed solution to automatically check the refinement of the model m^a to the model m^r . As outlined above, we particularly require that the applied refinement Ref satisfies the properties of a dbs -simulation H . For this purpose, *all* (valid) system states as well as *all* possible operation calls in those states need to be considered. Naive schemes, e.g., relying on enumerating all possible scenarios are clearly infeasible for this purpose. Hence, we propose an approach that maps the problem to an instance of *satisfiability modulo theories* (SMT) and, afterwards, exploits the efficiency of corresponding solving techniques (such as [9]).

To this end, we represent arbitrary system states and transitions for the abstract model m^a as well as the refined model m^r together with their invariants and the refinement relation Ref in terms of bit-vectors and bit-vector constraints. In the same way, the verification objectives proving that the applied refinement Ref indeed ensures **dbs**-simulation are encoded and checked automatically. In the following, the resulting verification objectives are briefly sketched. Then, we illustrate how to encode these in SMT.

1.5.1 Verification Objectives

As motivated in Sect. 1.3, we are interested in the relation between abstract operations and their possibly non-atomic refinements. These operation refinements are given as operation sequences according to Definition 5. By this, the refinement check is reduced to the question of whether there is a sequence of operation calls in the refined model that corresponds to a single call in the abstract model (according to the given refinement relation), but violates the requirements of the abstract operation. Unsatisfiability of such an instance shows that no such sequence exists and, hence, the refinement is correct. Otherwise, a counterexample showing the inconsistency is provided.

Based on this intuitive notion of refinement, we derive three verification objectives that prove the correspondence of an abstract operation and its refined operations and are sufficient to prove **dbs**-simulation. By this, the preservation of safety properties is guaranteed and the refinement is proven consistent. The three objectives read as follows:

1. Check whether all initial states in the refined model indeed correspond to the respective initial states in the abstract model, i.e.

$$\forall \sigma_0^r : \text{init}(\sigma_0^r) \Rightarrow \text{init}(\text{Ref}_\Sigma^{-1}(\sigma_0^r)).$$

This check is illustrated in Fig. 1.5a.

2. For each step o_j^r of the refined operation which transforms a refined state σ_1^r , check whether this step does not lead to a succeeding state σ_2^r which is inconsistent to its corresponding abstract states. In fact, the succeeding state σ_2^r either has to correspond to the unchanged abstract state or to its abstract state which results after applying the corresponding abstract operation o^a , i.e. for each step o_j^r

$$\begin{aligned} \forall \sigma^a, \sigma_1^r, \sigma_2^r : & \text{Ref}_\Sigma(\sigma^a, \sigma_1^r) \wedge \sigma_1^r \xrightarrow{o_j^r} \sigma_2^r \\ & \Rightarrow (\text{Ref}_\Sigma^{-1}(\sigma_2^r) = \sigma^a \vee (\triangleleft_{o^a}(\sigma^a) \wedge \triangleright_{o^a}(\sigma^a, \text{Ref}_\Sigma^{-1}(\sigma_2^r)))) \end{aligned}$$

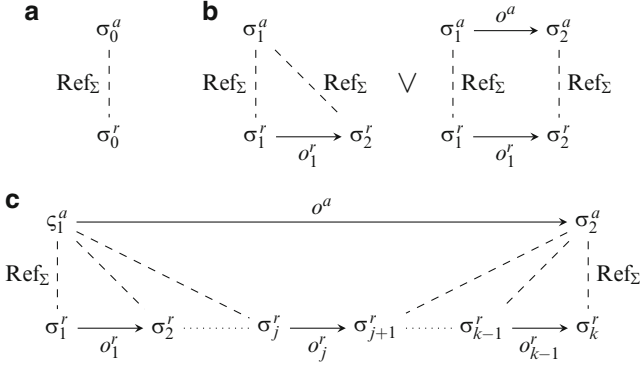


Fig. 1.5 Verification objectives. (a) Initialization; (b) Single step correspondence; (c) Chaining of refined steps

is checked. This check is illustrated in Fig. 1.5b. These two objectives are already sufficient to prove **dbS**-simulation. Nevertheless, a third objective is additionally checked.

3. Check whether the joint effect of the refined operation sequence adheres exactly to the specification of the abstract operation. That is, for each operation o^a and its refinement $o_1^r \dots o_k^r$

$$\begin{aligned} \forall \sigma_1^a, \sigma_1^r, \sigma_2^r \dots \sigma_{k+1}^r : & \text{Ref}_\Sigma(\sigma_1^a, \sigma_1^r) \wedge \sigma_1^r \xrightarrow{o_1^r} \sigma_2^r \dots \sigma_k^r \xrightarrow{o_k^r} \sigma_{k+1}^r \\ & \Rightarrow (\triangleleft_{o^a}(\sigma_1^a) \wedge \triangleright_{o^a}(\sigma_1^a, \text{Ref}_\Sigma^{-1}(\sigma_{k+1}^r))) \end{aligned}$$

is checked. This check is illustrated in Fig. 1.5c. This check particularly considers the common UML or SysML refinement which often refines a single abstract operation into a sequence of refined operations.

Together, these three objectives represent the verification tasks to be solved by the respective solving engine. Next, we illustrate how they are encoded as an SMT instance.

1.5.2 Basic Encoding

In order to represent arbitrary system states and transitions in an SMT instance, we use an encoding similar to the ones previously presented, e.g., in [3, 11, 12, 30] and particularly in [31]. Here, systems states (basically defined by the values of their attributes) and links are represented by corresponding bit-vector variables. Invariants are represented by corresponding SMT constraints. By this, it is ensured

that the solving engine only considers systems states σ composed of objects satisfying all invariants of the underlying class, i.e. $I_c(\sigma)$.

In order to encode transitions caused by operation calls, bit vectors $\omega_i \in \mathbb{B}^{\lceil \text{id}(o) \rceil}$ are created for each step i in the refined model. Depending on the assignment to ω , the respective pre-conditions and post-conditions have to be enforced. This can be realized by a constraint

$$\omega_i = \text{enc}(o) \Rightarrow \triangleleft_o(\sigma_i^r) \wedge \triangleright_o(\sigma_i^r, \sigma_{i+1}^r),$$

where $\text{enc}(o)$ represents a unique binary representation of the operation o , i.e. a number from 0 to $|O_r|$ with $\text{enc}(\text{id}) = 0$. Furthermore, to ensure that only legal values can be assigned to a vector ω , we use a constraint $\omega \leq |O_r|$.

We further introduce auxiliary predicates that reflect the relationship between an abstract operation and its refined steps. For this purpose, the operation refinement Ref_Ω is utilized:

$$\begin{aligned} \text{step}_i(o^a, o_j^r) &\Leftrightarrow \text{Ref}_\Omega(o^a) = o_1 \cdot o_2 \dots o_k \wedge o_i = o_j \\ \text{step}(o^a, o_j^r) &\Leftrightarrow \bigvee_{i=1}^{|\text{Ref}_\Omega(o^a)|} \text{step}_i(o^a, o_j^r). \end{aligned}$$

Here, $\text{step}_i(o^a, o_j^r)$ evaluates to true iff the refined operation o_j^r is the i th step in the refinement of o^a , while $\text{step}(o^a, o_j^r)$ reflects that o_j^r occurs in any position in the refinement of o^a .

In order to encode the chaining of the refined operation steps according to the scheme in Fig. 1.5c, we define the predicate chain:

$$\text{chain}(o^a) \Leftrightarrow \bigwedge_{i=1}^l (\text{step}_i(o^a, o_i^r) \wedge \omega_i = \text{enc}(o_i^r) \vee i > |\text{Ref}_\Omega(o^a)| \wedge \omega_i = \text{enc}(\text{id})).$$

In the above formula, in order to cover all abstract operations in one instance, the refined operation sequences are brought to the same maximal length l by filling up the sequence with the identity function for operations where $|\text{Ref}_\Omega(o)| < l$. We thereby make use of the maximum number of steps according to Ref_Ω , i.e. $l = \max\{|\text{Ref}_\Omega(o^a)| \mid o^a \in O^a\}$. Next, the above ‘‘ingredients’’ are put together in order to encode the verification objectives of a refinement.

1.5.3 Encoding the Verification Objectives

While the encodings from above ensure a proper representation of the models, system states, and execution of operations in an SMT instance, finally the verification objectives from Sect. 1.5.1 are encoded. In order to prove (1), we encode its negation

and check for unsatisfiability, i.e.

$$\exists \sigma_0^r, \sigma_0^a : \text{Ref}_\Sigma(\sigma_0^a, \sigma_0^r) \wedge \text{init}(\sigma_0^r) \wedge \neg \text{init}(\sigma_0^a). \quad (1.1)$$

To check (2), we try to determine a refined operation call that cannot be matched with one of the schemes in Fig. 1.5b. Hence, instead of encoding (2) for each individual refined operation, we let the solving engine choose a refined step that violates the requirements, i.e.

$$\begin{aligned} \exists \sigma_1^a, \sigma_2^a, \sigma_1^r, \sigma_2^r, o^a, o^r : & \text{Ref}_\Sigma(\sigma_1^a, \sigma_1^r) \wedge \omega_1 = \text{enc}(o^r) \\ & \wedge \text{step}(o^a, o^r) \wedge \text{Ref}_\Sigma(\sigma_2^a, \sigma_2^r) \\ & \wedge \neg (\sigma_1^a = \sigma_2^a \vee \triangleleft_{o^a}(\sigma_1^a) \wedge \triangleright_{o^a}(\sigma_1^a, \sigma_2^a)). \end{aligned} \quad (1.2)$$

That is, we check that, given a pair of corresponding states and an operation call in the refined state, whether it is possible that the reached refined state neither corresponds to the original abstract state nor does it satisfy the specification of the abstract operation. In case this instance is unsatisfiable, objective (2) has been proven.

Finally, for (3) we need to check whether we can determine an instantiated sequence of refined operation calls, such that their joint effect does not adhere to the specification of the respective abstract operation. For this purpose, we use the chain predicate as defined in the previous section to construct the *unrolled* operation sequence, i.e.

$$\begin{aligned} \exists \sigma_1^a, \sigma_2^a, \sigma_1^r \dots \sigma_{l+1}^r, o^a, o_1^r \dots o_l^r : & \text{Ref}_\Sigma(\sigma_1^a, \sigma_1^r) \wedge \text{chain}(o^a) \wedge \text{Ref}_\Sigma(\sigma_1^a, \sigma_{l+1}^r) \\ & \wedge \neg (\triangleleft_{o^a}(\sigma_1^a) \wedge \triangleright_{o^a}(\sigma_1^a, \sigma_{l+1}^a)). \end{aligned} \quad (1.3)$$

That is, we are searching for a chain of $l + 1$ refined states and connected by l operation calls such that there are no corresponding abstract states which satisfy the pre- and post-conditions of the respective abstract operation. Unsatisfiability proves that no such chain exists and, hence, objective (3) holds.

1.6 Evaluation

The approach presented in this chapter has been implemented in Java, using the SMT solver *Boolector* [9] as underlying solving engine. In order to evaluate the applicability and scalability of our approach, we have applied it to two systems based on examples presented in [1]. For the sake of comparison, these examples have additionally been verified using the previously proposed B method following a manual as well as an automatic scheme [20].

The first example describes an access control system (AC) which is employed to grant access to a building when presented with an authorized ID by a user. Two refinement steps have been modelled, a correct and an erroneous one, which are depicted in Fig. 1.6 together with the abstract model. All types of refinement

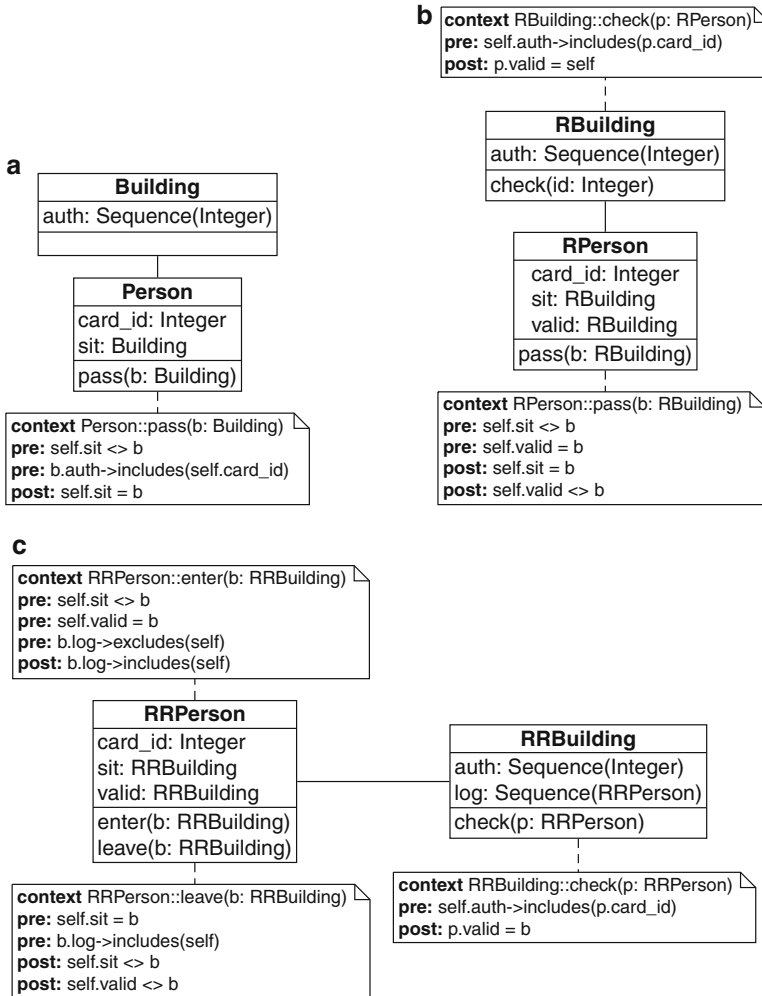


Fig. 1.6 Access control system. (a) Abstract model; (b) First refinement; (c) Second refinement

presented in this chapter have been applied to this model, i.e. attribute refinement as well as atomic and non-atomic operation refinement.

Table 1.1 provides the sizes of the three models (denoted by AC0, AC1, and AC2), i.e. the number of classes, attributes, operations, and OCL constraints are listed. As can be seen, the abstract model (AC0) and the two refined models (AC1, AC2) are relatively small regarding the number of UML elements. Only the number of OCL constraints increases slightly as the added and refined operations are extended.

In order to compare our work to the traditional B approach, we re-modelled this example in B and verified the refinement manually, using the event-B tool

Table 1.1 Size of examples

| Model | #Classes | #Attributes | #Operations | #Constraints |
|-------|----------|-------------|-------------|--------------|
| AC0 | 2 | 3 | 1 | 3 |
| AC1 | 2 | 4 | 2 | 6 |
| AC2 | 2 | 5 | 3 | 10 |
| MPC0 | 2 | 2 | 4 | 12 |
| MPC1 | 2 | 6 | 12 | 40 |
| MPC2 | 2 | 8 | 16 | 52 |
| MPC3 | 2 | 8 | 16 | 60 |

Rodin. The first refinement step led to a total of 14 proof obligations that had to be discharged. While five of them could be proven fully automatically and some further proofs needed only minor effort, the remaining ones required rather complex interactions like manually entered hypotheses or case splitting. Furthermore, the event-B model required additional invariants. This became particularly crucial for the second refinement which, due to the non-atomic nature of the refinement conducted here, could not be modelled in a straight-forward fashion in event-B.

In contrast, both steps could be automatically verified in negligible time by the approach proposed in this chapter. The non-atomic refinement did not lead to an increased run-time in this case.

The second example is a mechanical press controller (MPC), which has also been used to evaluate the automatic verification approach in [20] with the tool ProB. It describes a mechanical press with a motor, a clutch, and a door which interact in such a way as to guarantee a safe use. As in [20], we have modelled the first three refinement steps in UML and verified them with the proposed SMT-based approach. Here, the refinement contains the introduction of new attributes and constraints as well as atomic operation refinement. All three refinement steps have been proven to be correct.

Again, the size of the abstract model and its refinements is shown in Table 1.1 (denoted by MPC0, MPC1, MPC2, and MPC3). In contrast to the first example, the amount of attributes and operations as well as the number of OCL constraints increases. Especially the growing number of operations is important, since, for the SMT-based approach, each of these operations has to be verified according to the criteria presented earlier.

Table 1.2 shows the run-times of our experiments compared to those of ProB. The first two columns indicate which models have been verified against each other. The third and fourth columns contain the run-times of ProB without and with XSB Prolog taken from [20]. In [20], the ProB tool has already been compared to an automatic refinement verification approach based on CSP (namely [18]) which was clearly outperformed by ProB.

Again, the proposed approach proved the correctness of all three refinement steps in negligible time whereas the run-time of ProB was much larger. Also, with and without XSB Prolog, ProB's run-time increased drastically with every step in the

Table 1.2 Experimental results

| Abstract | Refined | Run-time | | |
|----------|---------|----------|--------------|---------------|
| | | ProB (s) | ProB+XSB (s) | SMT-based (s) |
| AC0 | AC1 | - | - | < 0.01 |
| AC1 | AC2 | - | - | < 0.01 |
| MPC0 | MPC1 | 6.28 | 2.85 | < 0.01 |
| MPC1 | MPC2 | 70.57 | 26.66 | < 0.01 |
| MPC2 | MPC3 | 333.85 | 136.12 | < 0.01 |

refinement process. A similar development has not been observed for the SMT-based method so far.

These experiments confirm that our approach is robust in such a way that it is applicable to various types of models and refinements. Neither errors in the refinement process nor the type of operation refinement—atomic or non-atomic—have a significant influence on the run-time.

1.7 Discussion: Extraction of a Refinement Relation

While the approach presented in Sect. 1.5 serves very well to prove the consistency of a given refinement relation, it may not always be applicable. In order to verify a refinement step, a refinement relation is necessary; otherwise, none of the verification objectives can be checked. However, such a relation may not be present in case that several designers are involved in the modelling process or the refinement process has not been documented.

In this case, methods to extract a refinement relation from the given models are required. The goal of such an extraction is not to obtain any arbitrary relation, but a correct relation based on the criteria presented in Sect. 1.5.1. In the following, we will discuss some related approaches from the literature before sketching some ideas how such an extraction could be realized in the setting of this chapter.

1.7.1 Existing Approaches

In the past, different approaches to retrieve traceability or refinement information have been proposed. Several works focus on information retrieval techniques [4, 19, 22]. Here, the basic idea is to identify textual similarities which may refer to the same concepts. Some of these works focus on relations between different levels of abstraction, e.g. between code and documentation. However, since information retrieval relies on textual similarities, re-naming model elements is a huge problem which might well occur during refinement. Egyed presents a structural analysis

to determine traceability links in [15]. He uses abstraction rules to map classes, attributes, and association. This method works on UML only without considering OCL constraints specifying the operations' behaviour. Briand et al. discuss the use of information gathered by monitoring the designer's modifications as a means to retrieve traceability links in [8]. Like in the approaches mentioned so far, the model's behaviour is not considered in particular.

The authors of [14] propose an adaptation of the algorithm from [26] by Robinson. They apply Robinson's approach to Z refinements and encode it in a model checker. Another extension of the same algorithm can be found in [27], relying on the same mechanism. A relation R containing all potential mappings is step-wise reduced by incorrect mappings until either a correct relation is determined or the all mappings have been removed. Although these approaches do in fact consider the specified behaviour, depending on the variation of the algorithm, the whole system has to be simulated. Since in the beginning R contains all pairs of states, the method does not scale to larger systems.

1.7.2 SMT-Based Relation Extraction

The related approaches discussed above are either of heuristic nature—and therefore incomplete—or they try to solve the problem in an exact way. In the latter case, representing the refinement relation explicitly is infeasible for larger models. Since the encoding of the refinement verification in SMT has proven very successful in terms of scalability, the question is if this approach could also be used to extract a correct refinement relation.

To understand the complexity of the problem, it is useful to view it in the context of automatic synthesis. Verifying a model wrt. some specification is conceptually easier than synthesizing a model which satisfies this specification. In our case, this is reflected in the complexity of the SMT encoding needed to solve the respective problem. The verification of a given refinement relation can be encoded in a purely existentially quantified formula, that checks or falsifies the existence of some pairs of states which violate the refinement relation. Intuitively, the extraction of a correct relation demands one quantifier more: Does there exist some relation such that for all pairs of states it verifies the refinement of our models? Thus, the problem cannot be solved in a complete manner using a quantifier-free SMT encoding.

While there are some solvers that support quantified formulae—such as Z3 [13]—the run-time and memory foot-print increase significantly with each additional quantifier alternation. Alternatively, a two-stage approach can be used that relies solely on quantifier-free encodings:

1. Find some pair of states and a relation which proves their refinement
2. Check if the found relation is a correct refinement relation
3. If yes, we are done. Otherwise continue with step 1

In the sketched algorithm, possible relations are enumerated by the underlying solver until a correct refinement is found. The verification in step 2 has already been solved in this chapter. The algorithm terminates in case of success or if no more relation can be found in step 1. In the latter case, we can be sure that the two models do not represent a correct refinement.

The effectiveness of the sketched approach critically depends on the first step. In the worst case, the algorithm will enumerate a huge number of incorrect relations that will be rejected by the second step. Additional constraints can help to reduce the number of iterations, but will in the same time increase the complexity of the first step. As a promising direction, *sequence diagrams*—representing test cases of the refined model—can be used to narrow down the set of candidate relations. If a correct refinement exists, it must also be applicable on a feasible run of the two models. This approach is subject to ongoing and future research.

1.8 Conclusions

In this chapter, we proposed an automatic approach which proves refinements of UML or SysML class diagrams. By this, we are considering the typical model-driven design flow which usually assumes an initial (abstract) model that is iteratively refined to a more precise representation. Based on a theoretical foundation, we introduced an SMT encoding checking whether the respective refinement relation represents a **dbS**-simulation and, hence, preserves (safety) properties from the abstract model to the refined model. We compared our approach to the tool ProB, which performs automatic refinement verification on B models. An experimental evaluation has shown that the SMT-based technique can verify refinements much faster and scales better than the B-based method.

For future work, we plan to extend our approach in order to support more modelling elements such as refinement of associations or parameters.

Acknowledgements This work was supported by the Graduate School SyDe (funded by the German Excellence Initiative within the University of Bremen’s institutional strategy), the German Federal Ministry of Education and Research (BMBF) within the project SPECiFiC under grant no. 01IW13001, as well as the German Research Foundation (DFG) within the Reinhart Koselleck project under grant no. DR 287/23-1 and a research project under grant no. WI 3401/5-1.

References

1. Abrial, J.R.: The B-book: Assigning Programs to Meanings. Cambridge University Press, New York (1996)
2. Abrial, J.R.: Modeling in Event-B: System and Software Engineering, 1st edn. Cambridge University Press, New York (2010)

3. Anastasakis, K., Bordbar, B., Georg, G., Ray, I.: UML2Alloy: A challenging model transformation. In: International Conference on Model Driven Engineering Languages and Systems, pp. 436–450. Springer, New York (2007)
4. Antoniol, G., Canfora, G., Casazza, G., Lucia, A.D., Merlo, E.: Recovering traceability links between code and documentation. *IEEE Trans. Softw. Eng.* **28**, 970–983 (2002)
5. Ben Ammar, B., Bhiri, M.T., Souquières, J.: Incremental development of UML specifications using operation refinements. *Innov. Syst. Softw. Eng.* **4**(3), 259–266 (2008). doi:10.1007/s11334-008-0056-1
6. Boiten, E.A.: Introducing extra operations in refinement. In: Formal Aspects of Computing, Springer London, pp. 1–13. Springer, London (2012)
7. Braunstein, C., Encrenaz, E.: CTL-property transformations along an incremental design process. *Int. J. Softw. Tools Technol. Transfer* **9**(1), 77–88 (2006). doi:10.1007/s10009-006-0007-9
8. Briand, L.C., Labiche, Y., Yue, T.: Automated traceability analysis for uml model refinements. *Inf. Softw. Technol.* **51**, 512–527 (2009)
9. Brummayer, R., Biere, A.: Boolector: an efficient SMT solver for bit-vectors and arrays. In: Tools and Algorithms for Construction and Analysis of Systems, pp. 174–177. Springer, Berlin (2009)
10. Bulychev, P., Konnov, I.V., Zakharov, V.A.: Computing (bi)simulation relations preserving CTL_X^* for ordinary and fair kripke structures. In: Mathematical Methods and Algorithms, vol. 12, pp. 59–76. Institute for System Programming, Russian Academy of Science (2006)
11. Cabot, J., Clarisó, R., Riera, D.: Verification of UML/OCL class diagrams using constraint programming. In: IEEE International Conference on Software Testing Verification and Validation Workshop, pp. 73–80 (2008)
12. Cadoli, M., Calvanese, D., Giacomo, G.D., Mancini, T.: Finite Model Reasoning on UML Class Diagrams Via Constraint Programming. In: R. Basili, M.T. Piazienza (eds.) AI*IA. Lecture Notes in Computer Science, vol. 4733, pp. 36–47. Springer, Berlin (2007)
13. De Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS’08/ETAPS’08, pp. 337–340. Springer, Berlin/Heidelberg (2008). URL <http://dl.acm.org/citation.cfm?id=1792734.1792766>
14. Derrick, J., Smith, G.: Using model checking to automatically find retrieve relations. *Electron. Notes Theor. Comput. Sci.* **201**, 155–175 (2008)
15. Egyed, A.: Consistent adaptation and evolution of class diagrams during refinement. In: Fundamental Approaches to Software Engineering (2004)
16. Glabbeek, R.: The linear time - branching time spectrum. In: J. Baeten, J. Klop (eds.) CONCUR ’90 Theories of Concurrency: Unification and Extension. Lecture Notes in Computer Science, vol. 458, pp. 278–297. Springer, Berlin/Heidelberg (1990)
17. Gogolla, M., Kuhlmann, M., Hamann, L.: Consistency, independence and consequences in UML and OCL models. In: Tests and Proofs, pp. 90–104. Springer, Berlin (2009)
18. Goldsmith, M., Roscoe, B., Armstrong, P.: Failures-Divergence Refinement - FDR2 User Manual (2005)
19. Hayes, J.H., Dekhtyar, A., Osborne, J.: Improving requirements tracing via information retrieval. In: IEEE International Requirements Engineering Conference (2003)
20. Leuschel, M., Butler, M.: Automatic Refinement Checking for B. In: International Conference on Formal Engineering Methods (2005)
21. Loiseaux, C., Graf, S., Sifakis, J., Bouajjani, A., Bensalem, S., Probst, D.: Property preserving abstractions for the verification of concurrent systems. *Form. Method. Syst. Des.* **6**(1), 11–44 (1995)
22. Natt och Dag, J., Regnell, B., Carlshamre, P., Andersson, M., Karlsson, J.: A feasibility study of automated natural language requirements analysis in market-driven development. *Requir. Eng.* **7**, 20–33 (2002)
23. Nejati, S., Gurfinkel, A., Chechik, M.: Stuttering abstraction for model checking. In: Software Engineering and Formal Methods, pp. 311–320. Springer, Berlin (2005)

24. Pons, C., Garcia, D.: Practical verification strategy for refinement conditions in UML models. In: *Advanced Software Engineering: Expanding the Frontiers of Software Technology*. IFIP International Federation for Information Processing, vol. 219, pp. 47–61. Springer, Berlin (2006)
25. Ranzato, F., Tapparo, F.: Computing stuttering simulations. In: *Concurrency Theory (CONCUR)*. Lecture Notes in Computer Science, vol. 5710, pp. 542–556. Springer, Berlin (2009)
26. Robinson, N.J.: Finding abstraction relations for data refinement. Technical Report, Software Verification Research Center, The University of Queensland (2003)
27. Robinson, N.J.: Incremental derivation of abstraction relations for data refinement. In: *Formal Methods and Software Engineering*. IEEE Computer Society, Los Alamitos (2003)
28. Rumbaugh, J., Jacobson, I., Booch, G.: *The Unified Modeling Language reference manual*. Addison-Wesley Longman, Essex (1999)
29. Snook, C., Butler, M.: UML-B: formal modeling and design aided by UML. *ACM Trans. Softw. Eng. Methodol.* **15**(1), 92–122 (2006)
30. Soeken, M., Wille, R., Kuhlmann, M., Gogolla, M., Drechsler, R.: Verifying UML/OCL models using Boolean satisfiability. In: *Design, Automation and Test in Europe*, pp. 1341–1344. IEEE Computer Society, New York (2010)
31. Soeken, M., Wille, R., Drechsler, R.: Verifying dynamic aspects of UML models. In: *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1–6 (2011)
32. Warmer, J., Kleppe, A.: *The Object Constraint Language: Precise Modeling with UML*. Addison-Wesley Longman, Boston, MA (1999)
33. Weilkens, T.: *Systems Engineering with SysML/UML: Modeling, Analysis, Design*. Morgan Kaufmann, San Francisco, CA (2008)
34. Woodcock, J., Davies, J.: *Using Z: Specification, Refinement, and Proof*. Prentice-Hall, Upper Saddle River, NJ (1996)