

Collaborative Verification-Driven Engineering of Hybrid Systems

Stefan Mitsch · Grant Olney Passmore ·
André Platzer

Received: 1 November 2013 / Revised: 15 February 2014 / Accepted: 3 March 2014 / Published online: 24 April 2014
© Springer Basel 2014

Abstract Hybrid systems with both discrete and continuous dynamics are an important model for real-world cyber-physical systems. The key challenge is to ensure their correct functioning w.r.t. safety requirements. Promising techniques to ensure safety seem to be model-driven engineering to develop hybrid systems in a well-defined and traceable manner, and formal verification to prove their correctness. Their combination forms the vision of verification-driven engineering. Often, hybrid systems are rather complex in that they require expertise from many domains (e. g., robotics, control systems, computer science, software engineering, and mechanical engineering). Moreover, despite the remarkable progress in automating formal verification of hybrid systems, the construction of proofs of complex systems often requires nontrivial human guidance, since hybrid systems verification tools solve undecidable problems. It is, thus, not uncommon for development and verification teams to consist of many players with diverse expertise. This paper introduces a verification-driven engineering toolset that extends our previous work on hybrid and arithmetic verification with tools for (1) graphical (UML) and textual modeling of hybrid systems, (2) exchanging and comparing models and proofs, and (3) managing verification tasks. This toolset makes it easier to tackle large-scale verification tasks.

Keywords Formal verification · Hybrid system · Cyber-physical system · Model-driven engineering

Mathematics Subject Classification (2010) Mathematical modeling (engineering) 97M50 · Hybrid systems 34K34 · Theorem proving 68T15

S. Mitsch (✉) · A. Platzer
Computer Science Department, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
e-mail: smitsch@cs.cmu.edu

A. Platzer
e-mail: aplatzer@cs.cmu.edu

S. Mitsch
Department of Cooperative Information Systems, Johannes Kepler University, Linz, Austria

G. O. Passmore
LFCS, Edinburgh and Clare Hall, Cambridge, 10 Crichton Street, Edinburgh, UK
e-mail: grant.passmore@cl.cam.ac.uk

1 Introduction

Computers that control physical processes form so-called *cyber-physical systems* (CPS), which are pervasively embedded into our lives today. For example, cars equipped with adaptive cruise control form a typical CPS [42] that is responsible for controlling acceleration on the basis of distance sensors. Further prominent examples can be found in many safety-critical areas, such as in factory automation [51], medical equipment [39], power plants and grid [71], automotive [18,23], aviation [73], and railway industries [65]. From an engineering viewpoint, a CPS can be described as a *hybrid system* in terms of discrete control decisions (the cyber-part, e. g., setting the acceleration of a car) and differential equations modeling the entailed physical continuous dynamics (the physical part, e. g., motion) [57]. More advanced models of CPS include aspects of distributed hybrid systems or stochasticity [61], but are not addressed here.

The key challenge in engineering hybrid systems is the question of how to ensure their correct functioning in order to avoid incorrect control decisions w.r.t. safety requirements (e. g., a car with adaptive cruise control will never collide with a car driving ahead). Promising techniques to ensure safety seem to be *model-driven engineering* (MDE) [30,69] to incrementally develop systems in a well-defined and traceable manner and formal verification to mathematically prove their correctness. Together, these techniques form the vision of *verification-driven engineering* (VDE) [35].

Often, CPS are complex systems and require expertise in many different domains, for instance, in robotics, control systems, computer science, software engineering, and mechanical engineering. Despite the remarkable progress in automating formal verification of hybrid systems, many interesting and complex verification problems still remain that are hard to solve in practice with a single tool by a single person. It is, thus, not uncommon for serious hybrid systems development and verification teams to consist of many players, some with expertise in robotics, control theory and dynamical systems, some in software engineering, some in mathematical logic, some in real algebraic geometry, and so on. Hence, modeling languages that convey a model to a broad and possibly heterogeneous audience together with integrated tools in a toolchain and well-established project management techniques to coordinate team members are crucial to achieve effective collaborative large-scale verification of hybrid systems (cf., e. g., [13,74]). For example, the graphical nature of UML seems to be suitable to convey information even to UML novices and promote communication between team members (e. g., as observed in a large industry study for safety-critical process control [4]). Moreover, it can even increase comprehensibility and accessibility of formal notations [66].

Collaboration on CPS verification is important for yet another reason: Because hybrid systems are undecidable [3], hybrid systems verification tools work over an undecidable theory, and so verifying complicated systems within them often requires significant human guidance. This need for human guidance is true even for *decidable* theories utilized within hybrid systems verification [12], such as the first-order theory of non-linear real arithmetic (also called the theory of *real closed fields* or RCF), a crucial component of real-world verification efforts. Though decidable, RCF is fundamentally infeasible (it is worst-case doubly exponential [14]), which poses a problem for the automated verification of hybrid systems. Much expertise is needed to discharge arithmetical verification conditions in a reasonable amount of time and space, expertise requiring the use of deep results in real algebraic geometry. Successful examples of team-based large-scale verification of non-hybrid systems include the operating system kernel seL4 [32] in Isabelle/HOL and the Flyspeck project [25], and show that, indeed, collaboration is key for proving large systems. Similar effects are expected in CPS verification.

This paper introduces the VDE toolset Sphinx comprising modeling and verification tools for hybrid systems (including a backend deployment for project management and collaboration support). The toolset applies proof decomposition in-the-large across multiple verification tools, basing on the completeness of differential dynamic logic ($d\mathcal{L}$ [57,60]), which is a real-valued first-order dynamic logic for hybrid programs, a program notation for hybrid systems. Sphinx extends our previous work on the deductive verification tool KeYmaera [64] and on the nonlinear real arithmetic verification tools RAHD [54] and MetiTarski [55] with tools for (1) graphical (UML) modeling, model transformation, and textual modeling of hybrid systems, (2) exchanging and comparing models and proofs, and (3) exchanging knowledge and tasks through a project management backend.

Structure of the Paper. In the next section, we give an overview of related work. In Sect. 3 we introduce our architecture of a verification-driven engineering toolset, and describe implementation and features of its components. Section 4 introduces an autonomous robotic ground vehicle as application example. Finally, in Sect. 5 we conclude the paper with an outlook on real-world application of the toolset and possible directions for future work.

2 Related Work

Model-driven engineering in a collaborative manner has been successfully applied in the embedded systems community. Efforts, for instance, include transforming between different UML models and SysML [26], modeling in SysML and transforming these models to the simulation tool Orchestra [5], integration of modeling and simulation in Cosmic/Cadena [22], or modeling of reactive systems and integration of various verification tools in Syspect [19].

Recent surveys on verification methods for hybrid systems [2], modeling and analysis of hybrid systems [16], and modeling of cyber-physical systems [17], reveal that indeed many tools are available for modeling and analyzing hybrid systems, but in a rather isolated manner. Supporting collaboration on formal verification by distributing tasks among members of a verification team in a model-driven engineering approach has not yet been the focus. Although current verification tools for hybrid systems (e. g., PHAVer [20], SpaceEx [21]), as well as those for arithmetic (e. g., Z3 [15]) are accompanied by modeling editors of varying sophistication, they are not yet particularly well prepared for collaboration either. Developments in collaborative verification of source code by multiple complementary static code checkers [10], modular model-checking (e. g., [37]), and extreme verification [27], however, indicate that this is indeed an interesting field. Most notably, usage of online collaboration tools in the Polymath project has led to an elementary proof of a special case of the density Hales–Jewett theorem [24].

The Unified Modeling Language (UML [28]) is a standardized language for object-oriented modeling. But without extension it is not well suited for modeling hybrid systems [9]. Therefore, the profiling mechanism of UML was used to extend the standardized UML languages SysML [26] for modeling hardware and software components of complex systems and MARTE [44] for modeling real-time and embedded systems. These and other extensions [9,40,68] increase the support for hybrid modeling in UML. However, those profiles augment the UML Statechart formalism, since their languages base on hybrid automata as underlying principle. We, instead, use hybrid programs and therefore extend UML Activity Diagrams since they are a more natural way of modeling. Examples for integrating formal notations with informal ones can be found outside the hybrid systems community, for instance, UML-B [70], TRIO [38], or VeriAgent [50].

In summary, we address model-driven engineering and formal verification as follows.

- Unlike [5,22,26], who focus on exchanging models, we also facilitate collaboration on formal verification.
- Unlike [15,20,21], who focus on one aspect of verification, we provide modeling and collaboration tools that should make it easier for domain experts to work in verification teams and exchange models and verification results between different tools.
- Unlike [15,20], who focus on verification tools, we also work on modeling support and collaboration.
- Unlike [9,40,68], who define a hybrid automaton semantics for UML Statecharts, we define a hybrid program semantics for UML Activity Diagrams.
- Unlike [38,50,70], who combine formal models with semi-formal modeling in UML for discrete systems, we define a UML profile for hybrid systems modeling.

3 The VDE Toolset Sphinx

Verification teams often comprise experts with diverse heterogeneous background who are accustomed to different modeling and verification tools with heterogeneous notations. In order to integrate different modeling and verification tools, the verification-driven engineering toolset Sphinx¹ proposed in this paper follows a model-driven architecture:

¹ <http://www.cs.cmu.edu/~smitsch/sphinx.html>.

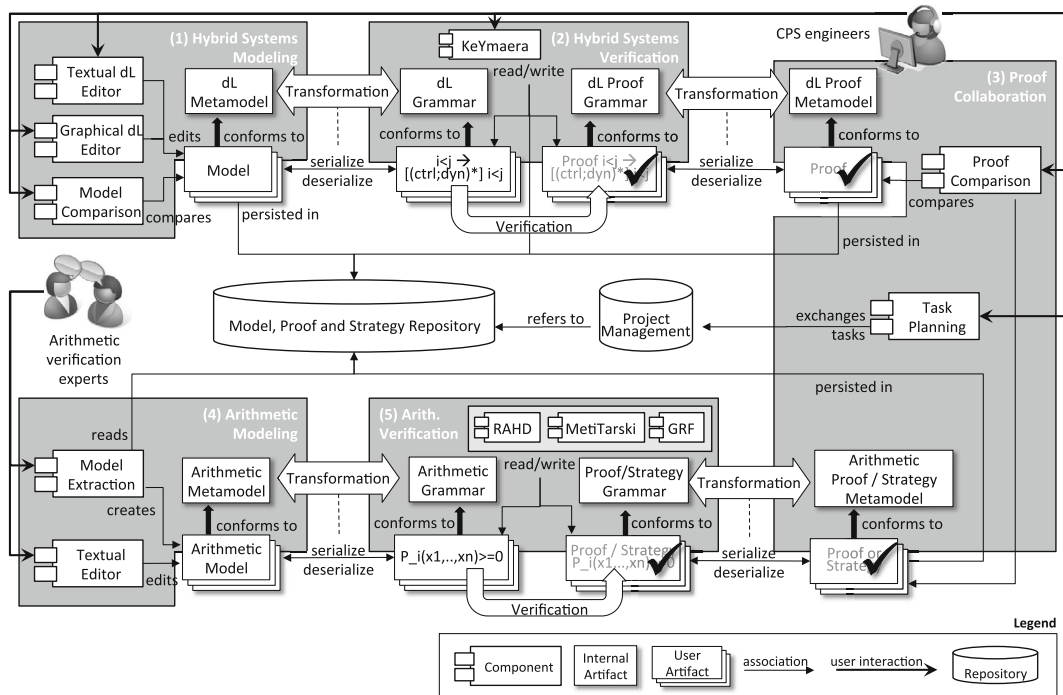


Fig. 1 Overview of components in the verification-driven engineering toolset. The toolset provides components for (1) hybrid systems modeling, (2) hybrid systems verification, (3) proof collaboration, (4) arithmetic modeling, and it provides or uses off-the-shelf components for (5) arithmetic verification. It uses further off-the-shelf components as repositories in the backend

Metamodels for different modeling and proof languages form the basis for manipulating, persisting, and transforming models. The idea is to provide general-purpose graphical and textual modeling languages for hybrid systems, while at the same time keeping the Sphinx platform open for additional languages. That way, we can still develop and integrate domain-specific languages (DSL), which are specifically tailored to the terminology used in a particular domain and with a semantics defined by transformation to hybrid programs [57–60]. Such DSLs would enable domain experts to express models that are suitable for verification purposes in their familiar terminology.

The notion of a model here denotes an instance of a metamodel, i. e., it comprises hybrid system models, proofs, and strategies. Following the definition of the Object Management Group (OMG²), a metamodel defines a language to formulate models: one example for a metamodel is the grammar of differential dynamic logic (dL [57–60]), which, among others, defines language elements for non-deterministic choice, sequential composition, assignment, repetition, and differential equations. An example for a model is given in Sect. 3.1: it describes a simple water tank as a set of formulas, differential equations, and other dL language elements. The model conforms to the grammar of dL, and thus is an instance of the dL metamodel. Figure 1 gives an overview of the toolset architecture: the dL metamodel, dL proof metamodel, arithmetic metamodel, and arithmetic proof metamodel each represent an interface between tools and to the backend.

dL metamodel The hybrid modeling components (textual and graphical editors for dL, as well as model comparison) manipulate models that conform to the dL metamodel. The dL models are serialized to and deserialized from their textual form that can be read by KeYmaera.

Hybrid Program UML The hybrid program UML profile extends UML with hybrid system concepts that can be translated to dL models.

² <http://www.omg.org>.

d \mathcal{L} proof metamodel The proof comparison component reads proofs that conform to the d \mathcal{L} proof metamodel.

These proofs may either be closed ones (completed proofs, nothing else to be done) or partial proofs (to be continued). Again, proofs in Sphinx are serialized and deserialized from the textual form as generated by KeYmaera.

Arithmetic metamodel Arithmetic editors (not yet implemented) manipulate arithmetic models. Again transformations are performed between models expressed in terms of the arithmetic metamodel and the corresponding textual input (e. g., SMT-LIB syntax [6]) as needed by arithmetic tools, such as RAHD [54], MetiTarski [1], Z3 [15], GRF [56], or MathSat [11].

Arithmetic proof metamodel Finally, the proof comparison component reads arithmetic proofs expressed in terms of the arithmetic proof metamodel, which is serialized to and deserialized from the textual format of the arithmetic tool.

3.1 A Hybrid Water Tank Example

We illustrate the notion of hybrid systems and our hybrid programs and hybrid program UML profile by means of the classical water tank example: a water tank should not overflow when the flow in or out of the water tank is chosen once every time interval. The hybrid program UML model is shown in Fig. 2. We will use the hybrid program UML syntax informally here and later introduce it in detail in Sect. 3.5.3.

The system introduces a global clock c and a bound on the loop execution time ε , which must be strictly positive as indicated by the invariant $\varepsilon > 0$ attached to the class *World*. The system further consists of one agent, the *WaterTank*, which is characterized by the current water level x , the current flow f and the maximum level M . The maximum level is constant (*readOnly*) and non-negative, as defined in the attached invariant $M \geq 0$.

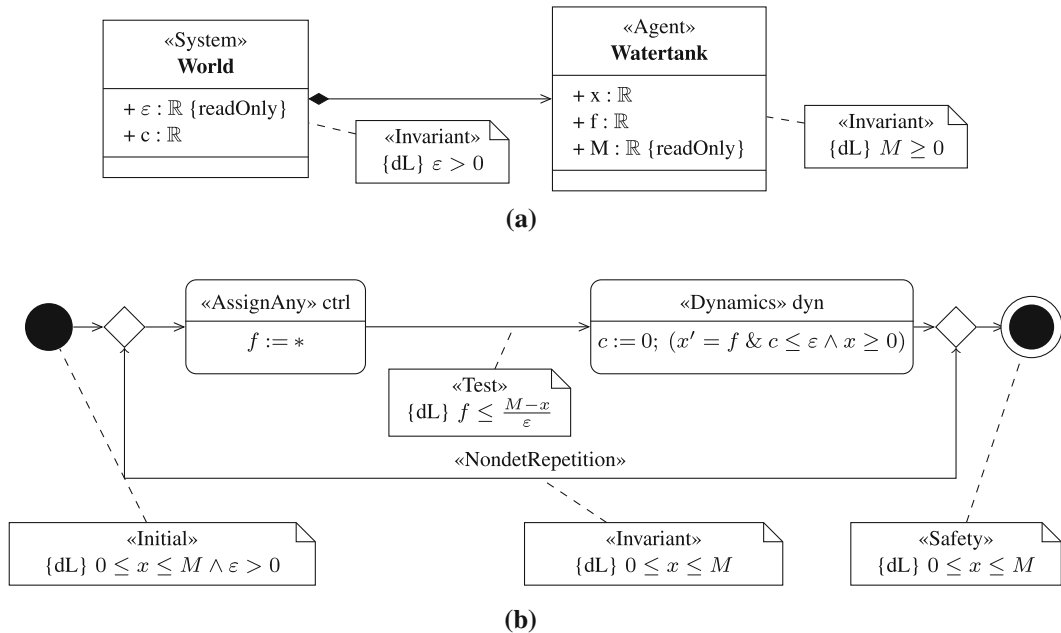


Fig. 2 Example of a hybrid system: a water tank. **a** The structure of the water tank model: the current water level x must not exceed the maximum level M when the flow f is chosen once every ε time units, which will be triggered by the clock c . **b** The behavior of the water tank model: the controller *ctrl* chooses a new nondeterministic value for the flow f , such that it satisfies the subsequent test before it is passed to the continuous dynamics *dyn*. Controller and continuous dynamics are repeated nondeterministically many times (which means it can be skipped entirely, cf. Table 1 and the hybrid program UML semantics in Table 2)

The behavior of the system is a single loop with two actions: the *ctrl* action chooses a new nondeterministic flow that will not exceed the water tank’s maximum capacity (cf. the test $\frac{M-x}{\varepsilon}$). The subsequent continuous evolution *dyn* resets the clock c and evolves the water level in the tank according to the chosen flow along the differential-algebraic equation $x' = v \ \& \ c \leq \varepsilon \wedge x \geq 0$ (the constraints ensure that the clock will not exceed a certain limit $c \leq \varepsilon$ and the water level will always be non-negative $x \geq 0$).

The specification about the water tank is annotated as constraints on the initial and the final node: when we start the water tank model in a state where the current level is within the limits of the water tank, then all runs of the model should keep the water level within the limits.

3.2 Development and Verification Process Overview

Let us exemplify the Sphinx toolset with a virtual walk through a collaborative verification scenario. We begin with modeling a hybrid system as in the water tank example above using the graphical and textual $d\mathcal{L}$ editors. The resulting model, which conforms to the $d\mathcal{L}$ metamodel, is transformed on-the-fly during editing to a textual input file, and loaded into KeYmaera. In KeYmaera, we apply various strategies for proving safety of our hybrid system model, but may get stuck at some difficult arithmetic problem. We mark the corresponding node in the partial proof and save it in KeYmaera’s textual output format. The proof collaboration tool transforms the partial proof text file into a model of the partial proof. We persist the hybrid model and the model of the partial proof in the model and proof repository. Then we create a request for arithmetic verification (ticket) in the project management repository using the task planning component. The assignee of the ticket accesses the linked partial proof, and extracts an arithmetic verification model from the marked proof node. Then a transformation runtime creates the textual input for one of the arithmetic verification tools. In this tool, a proof for the ticket can be created, along with a proof strategy that documents the proof. Such a proof strategy is vital for replaying the proof later, and for detecting whether or not the arithmetic proof still applies when the initial model has changed. Both, proof and proof strategy, are imported into the proof collaboration tool and persisted to the corresponding repository. The ticket is closed, together with the node on the original proof (if the arithmetic proof is complete; otherwise, the progress made is reported back). We fetch the new proof model version from the repository and inspect it using the proof comparison component. Then we transform the proof model into its textual form, load KeYmaera and continue proving our hybrid system from where we left off, but now with one goal closed. In case the corresponding arithmetic prover is connected to KeYmaera, we could even load the proof strategy from the strategy repository and repeat it locally to reduce proof effort on other subgoals.

3.3 KeYmaera: Hybrid System Verification

KeYmaera³ [64] is a verification tool for hybrid systems that combines deductive, real algebraic, and computer algebraic prover technologies. It is an automated and interactive theorem prover for a natural specification and verification logic for hybrid systems. KeYmaera supports differential dynamic logic ($d\mathcal{L}$) [57–60], which is a real-valued first-order dynamic logic for hybrid programs, a program notation for hybrid systems. KeYmaera supports hybrid systems with nonlinear discrete jumps, nonlinear differential equations, differential-algebraic equations, differential inequalities, and systems with nondeterministic discrete or continuous input.

For automation, KeYmaera implements a number of automatic proof strategies that decompose hybrid systems symbolically and prove the full system by proving properties of its parts [59]. This compositional verification principle helps scaling up verification, because KeYmaera verifies a big system by verifying properties of subsystems. Strong theoretical properties, including relative completeness results, have been shown about differential dynamic logic [57, 60].

³ <http://symbolaris.com/info/KeYmaera.html>.

KeYmaera implements fixedpoint procedures [62] that try to compute invariants of hybrid systems and differential invariants of their continuous dynamics, but may fail in practice. By completeness [57,60,61], this is the only part where KeYmaera’s automation can fail in theory. In practice, however, also the decidable parts of dealing with arithmetic may become infeasible at some point, so that interaction with other tools or collaborative verification via Sphinx is crucial.

At the same time, it is an interesting challenge to scale to solve larger systems, which is possible according to completeness but highly nontrivial. For systems that are still out of reach for current automation techniques, the fact that completeness proofs are compositional can be exploited by interactively splitting parts of the hybrid systems proof off and investigating them separately within Sphinx. If, for instance, a proof node in arithmetic turns out to be infeasible within KeYmaera, this node could be verified using a different tool connected to Sphinx.

KeYmaera has been used successfully for verifying case studies from train control [65], car control [41,42,46], air traffic management [43,63], robotic obstacle avoidance [45], and robotic surgery [36]. These verification results illustrate how some systems can be verified automatically while others need more substantial user guidance. The KeYmaera approach is described in detail in a book [59]. KeYmaera is linked to Sphinx by implementing extensions to the Eclipse launch configuration. These extensions hook into the context menu of Eclipse (models in $d\mathcal{L}$ and $d\mathcal{L}$ proof files in our case) and, on selection, launch KeYmaera as an external program. In the same fashion, further verification tools can be connected to Sphinx.

In order to guide domain experts in modeling discrete and continuous dynamics of hybrid systems, the case studies, further examples, and their proofs are included in the KeYmaera distribution. When applying proof strategies manually by selection from the context menu in the interactive theorem prover, KeYmaera shows only the applicable ones sorted by expected utility. Preliminary collaboration features include marking and renaming of proof nodes, as well as extraction of proof branches as new subproblems. These collaboration features are used for interaction with the arithmetic verification tools and the collaboration backend described below.

3.4 Real Arithmetic Verification

Proofs about hybrid systems often require significant reasoning about multivariate polynomial inequalities, i.e., reasoning within the *theory of real closed fields* (RCF). Though RCF is decidable, it is fundamentally infeasible (hyper-exponential in the number of variables). It is not uncommon for hybrid system models to have tens or even hundreds of real variables, and RCF reasoning is commonly the bottleneck for nontrivial verifications. Automatic RCF methods simply do not scale, and manual human expertise is often needed to discharge a proof’s arithmetical subproblems.

Real closed fields infeasibility is not just a problem for hybrid systems verification. Real polynomial constraints are pervasive throughout the sciences, and this has motivated a tremendous amount of work on the development of feasible proof techniques for various special classes of polynomial systems. In the context of hybrid systems verification, we wish to take advantage of these new techniques as soon as possible.

Given this fundamental infeasibility, how might one go about deciding large RCF conjectures? One approach is to develop a battery of efficient proof techniques for different practically useful fragments of the theory. For example, if an \exists RCF formula can be equisatisfiably transformed into an $\wedge\vee$ -combination of strict inequalities, then one can eliminate the need to consider any irrational real algebraic solutions when deciding the formula. Tools such as RAHD [54], Z3 [15] and MetiTarski [1] exemplify this heterogeneous approach to RCF, and moreover allow users to define *proof strategies* consisting of heuristic combinations of various specialized proof methods. When faced with a difficult new problem, one works to develop a proof strategy which can solve not only the problem at hand but also other problems sharing similar structure. Such strategies, though usually constructed by domain experts, can then be shared and utilized as automated techniques by the community at large.

As verification tools like KeYmaera progress, they accumulate a large database of RCF facts which pertain to the system being analyzed. As subsequent RCF subproblems are generated, they are tested for validity modulo this database of background facts. In practice, often only a small subset of the background RCF facts are needed

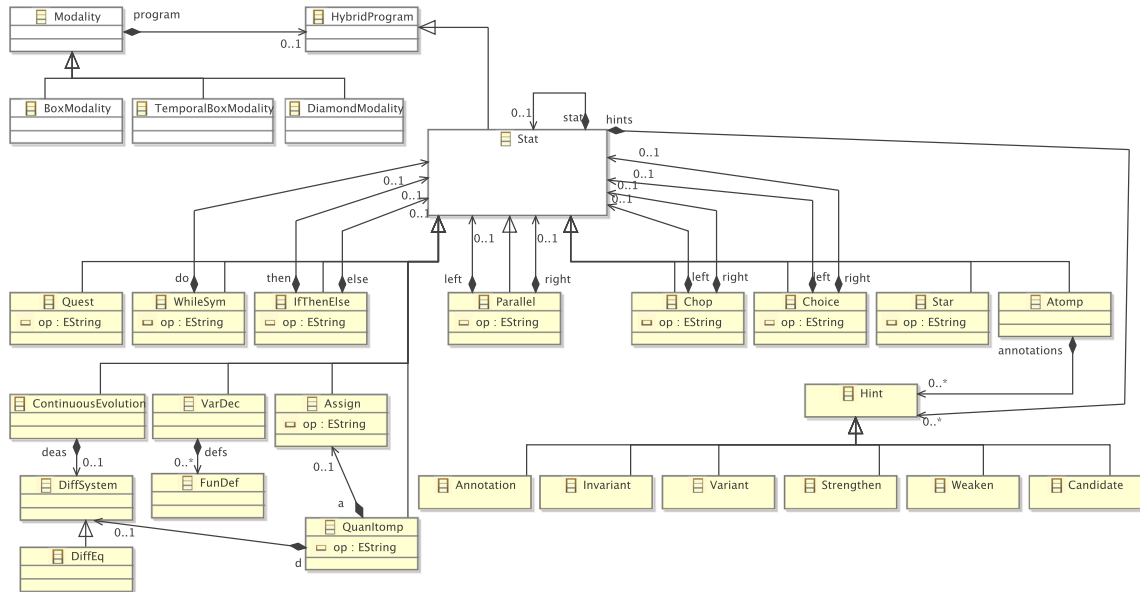


Fig. 3 The $d\mathcal{L}$ metamodel extracted from the input grammar of KeYmaera

to decide the generated subproblems. The difficulty lies in how the most relevant facts should be selected. The geometric relevance filtering (GRF) method [56] is an RCF decision method combining high-dimensional sampling techniques and incremental cell decomposition methods adapted from cylindrical algebraic decomposition (CAD) to use geometric information to select relevant background facts. GRF supports human experts in deciding which arithmetic subproblems to keep and which ones to discharge.

3.5 Modeling and Proof Collaboration

In order to interconnect the variety of specialized verification procedures introduced above, Sphinx follows a model-driven engineering approach: it introduces metamodels for the included modeling and proof languages. These metamodels provide a clean basis for model creation, model comparison, and model transformation between the formats of different tools. This approach is feasible, since in principle many of those procedures operate over the theory RCF, or at least share a large portion of symbols and their semantics. One could even imagine that very same approach for exchanging proofs between different proof procedures, since proofs in RCF, in theory, can all be expressed in the same formal system. Currently, proofs in Sphinx are exchanged merely for the sake of being repeated in the original tool (although KeYmaera already utilizes many such tools and hence is able to repeat a wide variety of proofs).

In the case of textual languages, Sphinx uses the Eclipse Xtext⁴ framework to obtain metamodels directly from the language grammars (cf. Fig. 3, obtained from the $d\mathcal{L}$ grammar [57]), together with other software artifacts, such as a parser, a model serializer, and a textual editor with syntax highlighting, code completion, and cross referencing. These metamodels are the basis for creating models in $d\mathcal{L}$, as well as for defining transformations between $d\mathcal{L}$ and other modeling languages. The models in $d\mathcal{L}$ make use of mathematical terms, and are embedded in KeY files since KeYmaera uses the KeY [49] format for loading models and saving proofs. In the following sections, we introduce $d\mathcal{L}$ in more detail and describe the support for creating $d\mathcal{L}$ models and working on proofs in Sphinx.

⁴ <http://www.eclipse.org/Xtext>.

3.5.1 Differential Dynamic Logic

For specifying and verifying correctness statements about hybrid systems, we use *differential dynamic logic* $\text{d}\mathcal{L}$ [57,59,60], which supports *hybrid programs* as a program notation for hybrid systems. The syntax of hybrid programs is summarized together with an informal semantics in Table 1; the metamodel introduced in Fig. 3 reflects this syntax. The sequential composition $\langle\alpha; \beta\rangle$ expresses that β starts after α finishes (e. g., first let a car choose its acceleration, then drive with that acceleration). The non-deterministic choice $\langle\alpha \cup \beta\rangle$ follows either α or β (e. g., let a car decide nondeterministically between accelerating and braking). The non-deterministic repetition operator $\langle\alpha^*\rangle$ repeats α zero or more times (e. g., let a car choose a new acceleration arbitrarily often). Discrete assignment $\langle x := \theta \rangle$ instantaneously assigns the value of the term θ to the variable x (e. g., let a car choose a particular acceleration), while $\langle x := * \rangle$ assigns an arbitrary value to x (e. g., let a car choose any acceleration). $\langle x' = \theta \ \& \ F \rangle$ describes a continuous evolution of x within the evolution domain F (e. g., let the velocity of a car change according to its acceleration, but always be greater than zero). The test $\langle ?F \rangle$ checks that a particular condition expressed by F holds, and aborts if it does not (e. g., test whether or not the distance to a car ahead is large enough). A typical pattern that involves assignment and tests, and which will be used subsequently, is to limit the assignment of arbitrary values to known bounds (e. g., limit an arbitrarily chosen acceleration to the physical limits of a car, as in $x := *; ?x \geq 0$). The deterministic choice $\langle \text{if}(F) \text{ then } \alpha \text{ else } \beta \rangle$ executes α if F holds, and β otherwise (e. g., let a car accelerate only when it is safe; brake otherwise). Finally, $\langle \text{while}(F) \text{ do } \alpha \text{ end} \rangle$ is a deterministic repetition that repeats α as long as F holds.

To specify the desired correctness properties of hybrid programs, differential dynamic logic ($\text{d}\mathcal{L}$) provides modal operators $[\alpha]$ and $\langle\alpha\rangle$, one for each hybrid program α . When ϕ is a $\text{d}\mathcal{L}$ formula (e. g., a simple arithmetic constraint) describing a state and α is a hybrid program, then the $\text{d}\mathcal{L}$ formula $[\alpha]\phi$ states that all states reachable by α satisfy ϕ . Dually, $\text{d}\mathcal{L}$ formula $\langle\alpha\rangle\phi$ expresses that there is a state reachable by the hybrid program α that satisfies $\text{d}\mathcal{L}$ formula ϕ . The set of $\text{d}\mathcal{L}$ formulas is generated by the following EBNF grammar (where $\sim \in \{<, \leq, =, \geq, >\}$ and θ_1, θ_2 are arithmetic expressions in $+, -, \cdot, /$ over the reals):

$$\phi ::= \theta_1 \sim \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \rightarrow \psi \mid \phi \leftrightarrow \psi \mid \forall x\phi \mid \exists x\phi \mid [\alpha]\phi \mid \langle\alpha\rangle\phi .$$

Table 1 Statements of hybrid programs

Statement	Metamodel element	Effect
$\alpha; \beta$	Chop	Sequential composition, performs HP α and then HP β afterwards
$\alpha \cup \beta$	Choice	Nondeterministic choice, follows either HP α or HP β
α^*	Star	Nondeterministic repetition, repeats HP α $n \geq 0$ times
$x := \theta$	Assign (term)	Discrete assignment of the value of term θ to variable x (jump)
$x := *$	Assign (wild card term)	Nondeterministic assignment of an arbitrary real number to x
$(x_1' = \theta_1, \dots,$ $x_n' = \theta_n \ \& \ F)$	ContinuousEvolution DiffSystem	Continuous evolution of x_i along differential equation system $x_i' = \theta_i$, restricted to maximum domain or invariant region F
$?F$	Quest	Check if formula F holds at current state, abort otherwise
$\text{if}(F) \text{ then } \alpha \text{ else } \beta$	IfThenElse	Perform HP α if F holds, perform HP β otherwise
$\text{while}(F) \text{ do } \alpha \text{ end}$	WhileSym	Perform HP α as long as F holds
$[\alpha]\phi$	BoxModality	$\text{d}\mathcal{L}$ formula ϕ must hold after all executions of HP α
$\langle\alpha\rangle\phi$	DiamondModality	$\text{d}\mathcal{L}$ formula ϕ must hold after at least one execution of HP α

3.5.2 Creating Models

Sphinx currently includes $d\mathcal{L}$ as generic modeling language to create models of hybrid and cyber-physical systems. The concrete textual $d\mathcal{L}$ editor is created from the $d\mathcal{L}$ metamodel and shown in Fig. 4, which also illustrates the graphical editor based on UML and the KeYmaera prover attached through the console.

In order to facilitate the creation of textual models in $d\mathcal{L}$, Sphinx includes templates of common model artifacts (e. g., ODEs of linear and circular motion). These templates, when instantiated, allow in-place editing and automated renaming of the template constituents. As usual in the Eclipse platform, such templates can be easily extended and shared between team members.

Since generic modeling languages, such as $d\mathcal{L}$ for hybrid systems, tend to incur a steep learning curve, the Sphinx platform can be extended with dedicated domain-specific languages (DSL). Such DSLs should be designed to meet the vocabulary of a particular group of domain experts. They can be included into Sphinx in a fashion similar to the generic modeling language $d\mathcal{L}$, i. e., in the form of Eclipse plugins that provide the DSL metamodel and the modeling editor.

In the next section we describe the Hybrid Program UML profile, an extension to UML for graphical modeling of hybrid programs that should make it easier to convey the main features of a hybrid system to a broader audience.

3.5.3 The Hybrid Program UML Profile

The Hybrid Program UML profile follows a fundamental principle of UML in that it separates modeling the structure of a hybrid system from modeling its behavior. Currently, Sphinx supports class diagrams for modeling the structure of a hybrid system, since hybrid programs do not yet support modules. Future work includes the addition of composite structure diagrams as used in [9,68], and the introduction of proof rules that exploit the

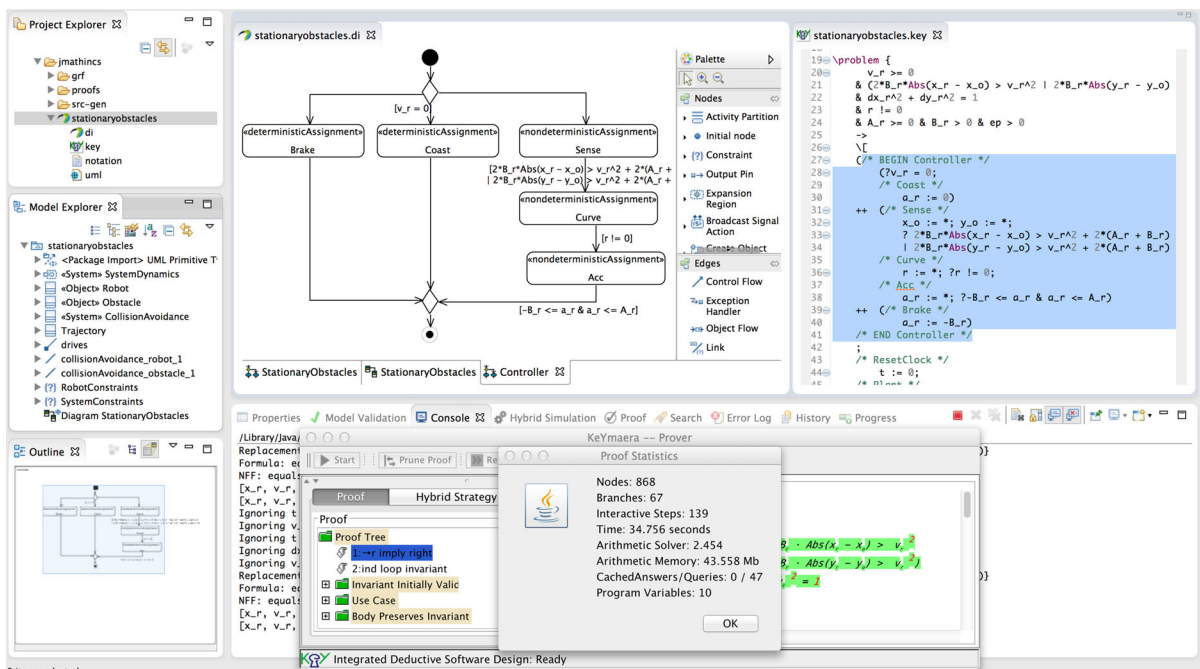


Fig. 4 Screenshot of the textual and graphical modeling editors and a proof in KeYmaera (details on the textual and graphical syntax are described in Sect. 3.5). The *left-most three panels*, from *top to bottom*, show the project explorer of Eclipse (access to models and proofs), a hierarchical tree view of the graphical model, and a miniature outline of the graphical model. The graphical editor is displayed in the *center* of the tool with the textual editor to the *right*. The text editor selection highlights the controller part that is displayed in the graphical editor. The *bottom-most panel* shows the KeYmaera theorem prover with its console output in the background

additional structural information during the verification process. For modeling behavior, we use activity diagrams instead of the UML statecharts used in existing hybrid system UML profiles [9,40,68], since activity diagrams model control flow more akin to hybrid programs. UML statecharts are a language to model system behavior as graphs where the vertices are the states of the system and the edges represent transitions between states. Thus, with some extension UML statecharts are suitable to represent hybrid automata (cf. [9,40,68]). UML activity diagrams, in contrast, are a language to model system behavior as graphs where the vertices are actions or decisions and the edges represent control flow. The notion of control flow between actions (statements) makes activity diagrams more suitable to model (computational) processes [28], such as our hybrid programs.

We use model transformation to define the semantics of the Hybrid Program UML profile relative to $d\mathcal{L}$. Besides defining the semantics of the Hybrid UML profile, model transformations can be implemented as model transformation specifications (e. g., using the Atlas transformation language ATL [29]) and executed to transform models back and forth between the Hybrid Program UML profile and their hybrid program counterparts. Since hybrid automata can be encoded in hybrid programs [59, Appendix C],⁵ we define both, a hybrid program semantics and a hybrid automaton semantics, for the Hybrid Program UML profile. Note, however, that hybrid automata, when encoded in $d\mathcal{L}$, are often less natural to express and also less efficient to verify than well-structured hybrid programs, because they lack program structure that could be exploited during the proof and require additional variables to identify the states of the automaton.

We use *UML profiles* as extension mechanism to provide hybrid system modeling concepts that are not yet present in standard UML. Profiles are the standard way to extend UML with domain-specific modeling concepts [28]. A UML profile is defined by specifying *stereotypes* and *constraints*. A stereotype is applicable to a particular element of the UML (e. g., a classifier) and adds additional modeling capabilities to the original UML element. For example, standard UML actions have a body that we can use to capture an atomic hybrid program, such as deterministic assignment or differential equations. When we want to describe additional information, such as differential invariant constraints or evolution domain constraints of a differential-algebraic equation, we can introduce a stereotype *Dynamics* for UML actions that adds the necessary modeling abilities to actions. Constraints can be added to profiles in order to restrict properties to only admissible values, derive property values from other properties, or otherwise check the consistency of a UML model. UML provides the Object Constraint Language (OCL, [53]) for defining such constraints. In the following paragraphs we describe our profiles for modeling the structure and the behavior of a hybrid system, which was already informally used in Sect. 3.1.

System Structure Hybrid programs in $d\mathcal{L}$ use variables and functions over the reals as modeling primitives. Further structuring mechanisms, such as classes, are not yet supported in $d\mathcal{L}$. In order to still capture and communicate the intended structure of a hybrid program, we provide stereotypes for UML class diagrams.⁶ Figure 5 shows the stereotypes currently available in the Hybrid UML profile for modeling the structure of a hybrid system.

A hybrid system usually consists of multiple *entities* [52], which are either *objects* that may evolve through manipulation but not by themselves, or *agents* whose evolution is driven by the decisions of a controller (e. g., the robot agent that has to avoid obstacles).

Entities are usually characterized by some *properties* (e. g., the robot's position) [7,33]. These properties can be discerned into *constant* properties (cf. *Constant*: their value can be read but not written, e. g., the position of a stationary obstacle), whereas others can change and are therefore called *fluent* [67] or *variable* (cf. *Variable*: their value can both be read and written, e. g., the position of the robot). These stereotypes can be equivalently modeled using standard UML notation *readOnly* for properties. Additional constraints may apply to properties (e. g., a minimum positive braking force $B > 0$, or bounds on the acceleration $-B \leq a \leq A$). We can model these constraints in the structure of the system using the stereotype *Invariant*, if the constraints have to be satisfied throughout model

⁵ The transformation from hybrid automata into hybrid programs follows the same principle as implementing a finite automaton in a programming language. The converse transformation from hybrid programs into hybrid automata is based on the transition structure induced by the semantics of hybrid programs [59,61].

⁶ In principle, a single class would already be a valid structure for a hybrid program. It is, however, useful to split the system into multiple separate classes corresponding to different entities in the system.

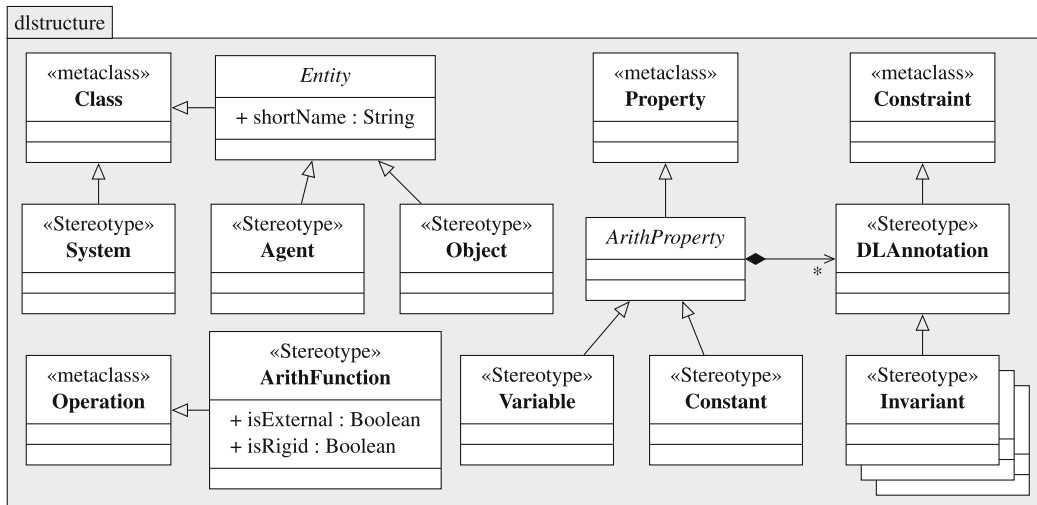


Fig. 5 The Hybrid Program UML profile: structure of hybrid programs

execution; otherwise, they are part of the system behavior. We use $d\mathcal{L}$ to formalize those constraints, since OCL does not support arbitrary arithmetic expressions. Some properties in a hybrid program are shared among all entities (e. g., time). A class marked with the stereotype *System* can capture such shared knowledge.

We allow further decomposition of agents into multiple classes. These classes are linked to their respective agent via the *association* concept of UML. If we want to emphasize that an instance of some class is owned by at most one agent at a time, we use *composition* (e. g., a robot’s internal control variables could be factored into a dedicated control state, which no other robot has access to). If we want to share instances of a class between multiple agents, we use a standard *association* instead. No further annotation with stereotypes is necessary.

System Behavior Hybrid programs know essentially two kinds of actions that can change the state of a system: instantaneous jumps (i. e., assignment) are part of the discrete control structure of a hybrid system, and differential equations are part of the continuous dynamics of a hybrid system. Figure 6 shows the stereotypes for modeling the discrete and continuous dynamics of a hybrid system.

Activity diagrams, as introduced briefly above, provide modeling concepts to represent actions (opaque actions are essentially atomic blackbox actions), decisions, guarded control flow, and constraints. In hybrid program UML we distinguish the following actions (actions are represented as rounded rectangles with a name compartment and a body compartment in UML): Nondeterministic assignment (*AssignAny*) chooses any value for the variable. Deterministic assignment (*AssignTerm*) chooses the value defined by an arithmetic term for the variable. Continuous evolution (*Dynamics*) evolves the values of variables along a differential equation system that stays within a maximum evolution domain. The differential invariant of the differential equation system, if known, can be annotated as a constraint to the dynamics action. The common features of deterministic and nondeterministic assignment are factored into the abstract base class *Assignment*.

The control flow of a hybrid program can be defined with three composition operations for hybrid program statements: nondeterministic choice, sequential composition, and nondeterministic repetition. Nondeterministic choice can be modeled with standard UML notation for decisions (splitting and merging nodes), while sequential composition is control flow. We introduce a stereotype *NondetRepetition* for control flow (i. e., nondeterministic repetitions will be backwards edges), which can be annotated with a constraint to specify an inductive loop invariant.

Tests in hybrid programs ensure that a particular condition is satisfied in subsequent program statements. They are modeled as algebraic constraints on control flows. Further useful constraints are either part of the specification language (*Initial*, *Safety*, *Liveness*) to express correctness criteria, or guide the verification process but do not influence system behavior (*InductiveInvariant*, *DiffInvariant*, *Convergence*, *DiffVariant*).

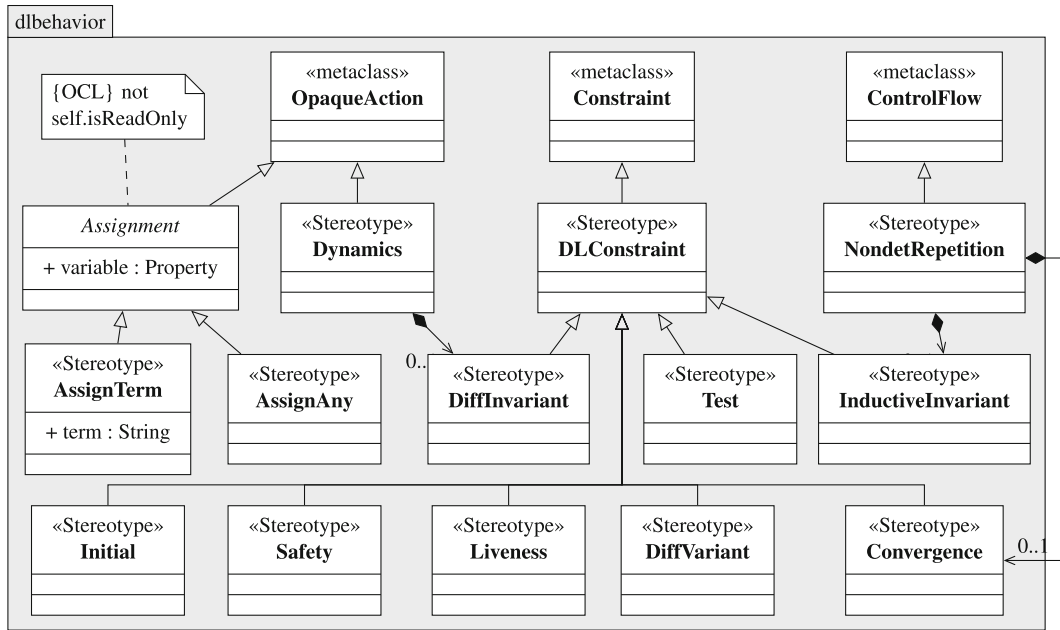


Fig. 6 The Hybrid Program UML profile: behavior of hybrid programs

We define the semantics of Hybrid Program UML relative to $d\mathcal{L}$ by transformation specifications from the Hybrid Program UML profile to the hybrid program metamodel. Currently, we support two kinds of transformations: *well-structured* activity diagrams⁷ can be transformed into well-structured hybrid programs with loops, whereas arbitrary activity diagrams can be transformed into $d\mathcal{L}$ automata, which are hybrid automata embedded into hybrid programs with additional variables and tests to represent the states.

The Hybrid Program Semantics of Hybrid Program UML Currently, all constants and variables are handled globally as a flat structure (i. e., the structuring mechanisms present in a Hybrid Program UML model are not yet directly reflected in a hybrid program). The hybrid programs α and β in Table 2 represent either one of the atomic actions in Hybrid Program UML (assignment, nondeterministic assignment, or continuous dynamics) or a well-structured part of an activity diagram. An edge between two actions α and β corresponds to a sequential composition of the corresponding transformed hybrid programs α and β , cf. (1). A guard on an edge is transformed into a sequential composition with an intermediate test, cf. (2). A decision node and a matching merge node with a forward edge and a backward edge is translated into a nondeterministic repetition, cf. (3). If the forward edge is missing this means at least one repetition, cf. (4). Decision nodes with tests are either translated into if-statements (5) or if-else-statements (6). Finally, a decision node with a matching merge node (but without a back edge) is transformed into a nondeterministic choice, cf. (7). This transformation in Table 2 gives perfectly structured hybrid programs for well-nested activity diagrams.

The Hybrid Automaton Embedding Semantics of Hybrid Program UML The hybrid automaton embedding in hybrid programs defined in Table 3 matches smaller patterns in an activity diagram compared to the hybrid program transformation in Table 2. It is thus applicable to a wider range of activity diagrams, which do not even have to be well-structured (i. e., arbitrary state jumps are allowed). As a downside, the transformation preserves no explicit program structure (e. g., sequence of statements) that could be exploited during verification. This means that sufficient information about the program structure has to be conveyed in the system invariant. This practice significantly increases the verification effort.

⁷ Well-structured activity diagrams consist of properly nested loops and branches and define a unique initial and final node.

Table 2 The hybrid program semantics of Hybrid Program UML












HP UML	Hybrid program	Description
(1) 	$\alpha; \beta$	Direct control flow is a sequential composition
(2) 	$\alpha; ?F; \beta$	Guarded control flow is a sequential composition with intermediate test
(3) 	α^*	Decision node and merge node linked with backedge and forward edge is a nondeterministic repetition
(4) 	$\alpha; \alpha^*$	Decision node and merge node linked with backedge is at least one repetition
(5) 	$\text{if}(F) \alpha \text{ fi}$	Decision node and merge node linked with forward edge is a conditional branch
(6) 	$\text{if}(F) \alpha \text{ else } \beta \text{ fi}$	Decision node and merge node with actions and mutually exclusive guards on each branch are if-else conditional
(7) 	$\alpha \cup \beta$	Decision node and matching merge node are a nondeterministic choice (analogous for more than two branches)

Table 3 The hybrid automaton embedding semantics of Hybrid Program UML

HP UML	Hybrid automaton embedding in hybrid program	Description
(1) 	$?s = id(\alpha); \alpha; s := id(\beta)$	Action with direct control flow to another action is a sequential composition of a location test and the actual atomic action, with the transition modeled as assignment of a new location ID
(2) 	$?s = id(\alpha); \alpha; s := id(\diamond)$	Control flow between an action and a decision/merge node is a sequential composition of a state test and assignment of a new state ID
(3) 	$?s = id(\alpha); \alpha; ?F; s := id(\beta)$	Guarded control flow is a sequential composition with an intermediate test (analogous for decision/merge node)
(4) 	$?s = id(\diamond); (s := id(\alpha) \cup s := id(\beta))$	Control flow between a decision node and actions is a nondeterministic choice (analogous for more than two branches)

The hybrid automaton embedding constructs an automaton-structure from hybrid program notation instead of explicit program structure [59]. It uses an additional variable s to keep track of the current location of the automaton. A unique identifier per vertex and edge of the activity diagram identifies the automaton location. The hybrid automaton embedding is then a nondeterministic choice over all the locations embedded in a single

nondeterministic repetition,⁸ and the control flow is translated into updates of the current location with the respective follow-up location, as summarized in Table 3. The construction is analogous to the embedding of hybrid automata into hybrid programs [59, App. C], which follows exactly the same principles of implementing finite automata as programs.

3.5.4 Hybrid System Simulation

An interesting opportunity for inspecting the behavior of a hybrid system during the modeling phase (prior to verification) is provided by Mathematica 9, which is able to simulate and plot hybrid system behavior using a combination of *NDSolve* and *WhenEvent* conditions.⁹ We transform corresponding excerpts of *dL* to Mathematica for visualizing plots of the dynamic behavior of a hybrid program over time in Sphinx. Simulations can be useful for debugging system models and quickly conveying intuitions about their behavior to the respective members of the collaborative verification-driven engineering team.

3.5.5 During the Proof

Collaboration support in Sphinx includes model and proof comparison tools, both locally and with the model and proof repositories maintained in a central source code repository. For this, not only textual comparison is implemented, but also structural comparison of models expressed in terms of the *dL* metamodel and proofs expressed in terms of the *dL* proof metamodel is supported (cf. Fig. 7). Exchanging proofs and inspecting updates on partial proofs is vital especially when multiple team members collaborate on a proof. We inherit textual comparison integrated into our textual modeling editors from Eclipse and the Xtext framework. Textual comparison of proofs, however, may not be the most efficient way of pointing to the relevant changes in a proof, because proofs, in addition to the relevant proof steps, often contain information for bookkeeping to mechanically check the proof. Thus, we additionally inherit a structural comparison tool from the EMF Compare Diff/Merge framework¹⁰ on the basis of a graph diff algorithm [34], so that differences are categorized by the relationships in the metamodel. The relevant changes in the actual proof section of the proof file are then compared by their metamodel elements, which represent proof rules; if necessary, the additional bookkeeping code can be overlaid on demand. Figure 7 shows a structural comparison overview in the top pane. Four changes are in the proof section, which is expanded; three of them are located in the part of the inductive proof of a loop, where we have to show that the loop body preserves the invariant. The first applied proof rule is selected and, therefore, the two lower panes show where this new proof step was inserted into the proof.

Specific unsolved subproblems of a proof (e. g., complex arithmetic problems) can be flagged in KeYmaera and extracted to other tools to further facilitate knowledge and expertise exchange. That makes it easier to partition the verification effort and collaborate in jointly coming up with a solution. An open question, however, concerns the merging of the partial verification results into a single coherent proof without recourse to external verification steps. In a first step, in Sphinx we only allow exchanging proof strategies that can be executed by KeYmaera. Sphinx injects these proof strategies directly into a *dL* proof instead of an open goal, and KeYmaera checks the injected proof steps for correctness. That way, external (arithmetic) solvers can replace manual verification effort without compromising proof trustworthiness.

Later, actual proof certificates and further proof strategies will be exchanged to further increase trust, and more sophisticated comparisons of proof goals are envisioned to more robustly support replaying proofs.

⁸ The final location does not change the location variable s . Thus, the system remains in the final location despite the fact that the nondeterministic repetition is allowed to execute arbitrarily many times.

⁹ <http://www.wolfram.com/mathematica>.

¹⁰ <http://www.eclipse.org/emf/compare/>.

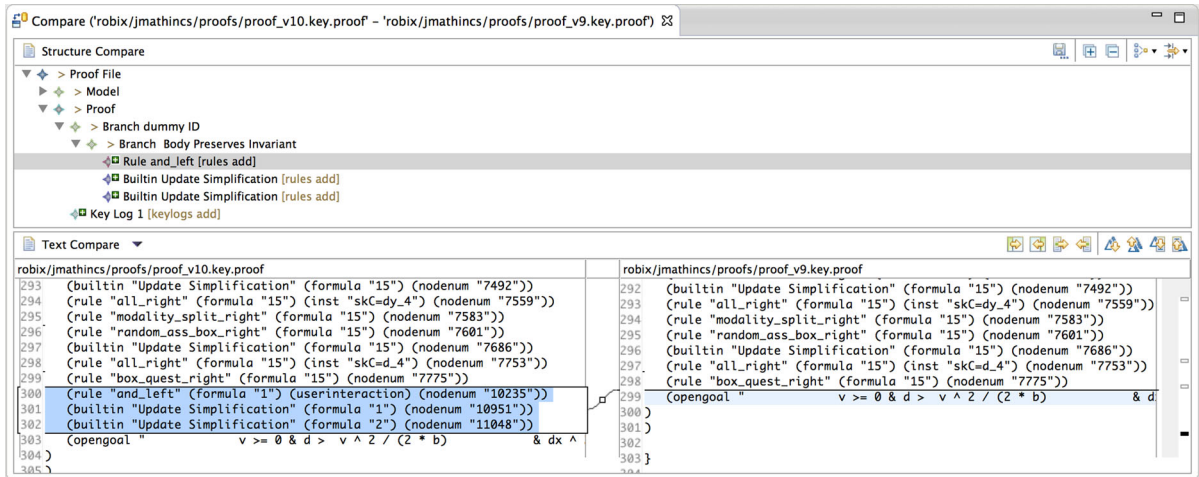


Fig. 7 Comparison of the structure of two proof versions

3.6 The Collaboration Backend

Technical Details The Sphinx modeling tool uses existing Eclipse plugins to connect to a variety of backend source code repositories and online project management tools. As source code repository we currently use Subversion¹¹ and the Eclipse plugin Subclipse,¹² but any other source code repository that is connected to Eclipse would work just as well. Currently, Mylyn¹³ and its connectors are used for accessing online project management tools (e. g., Bugzilla,¹⁴ Redmine,¹⁵ or any web-based tool via Mylyn’s Generic Web templates connector) and exchanging tickets (i. e., requests for verification). These tickets are the organizational means for collaborating on verification problems and tasks within a working group. Exchange of models and proofs may then be conducted either by attaching files to tickets, or by linking tickets directly to models and proofs in the source code repository. In the latter case, one benefits from the model and proof comparison capabilities of Sphinx. The collaboration through source code repositories and online project management systems seems useful because those are generally well-accepted by computer scientists and in other engineering domains; what still needs to be determined, however, is whether these collaboration tools are similarly well received by other domain experts.

Horizontal Splitting An important question for collaboration in cyber-physical systems verification is how domain experts with abstract, high-level understanding of a model can work with experts in control and dynamics to fill in details, and finally with verification experts to detect constraints that are necessary for verification. Such a *horizontal splitting* is in accordance with the distinction between high-level platform-independent models (PIM) and more detailed platform-specific models (PSM) as promoted by the OMG. In our tool suite, domain experts can introduce placeholders into graphical models, which are essentially activity nodes with a descriptive name but without the formal definition of their underlying meaning (e. g., the dynamics of the water tank as a differential equation). Sphinx translates placeholders into comments in the formal language, so that other experts (e. g., in modeling motion with differential equations) can fill in those gaps. Moreover, if Sphinx encounters a placeholder in a graphical model, it augments the safety condition with *false* so that all proof attempts on incomplete models will fail. Domain experts

¹¹ subversion.apache.org.

¹² subclipse.tigris.org.

¹³ <http://www.eclipse.org/mylyn>.

¹⁴ <http://www.bugzilla.org>.

¹⁵ <http://www.redmine.org>.

can even omit tests (constraints) on transitions. The absence of constraints that are essential for safety will be detected by the theorem prover KeYmaera anyway, as described below.

Constraint Detection Often, a development team has a good overall understanding of the desired safety conditions, but not of all of the corner cases and constraints that make the system safe w.r.t. these safety conditions. Therefore, one of the most challenging problems in cyber-physical systems verification is to find such constraints (e. g., loop invariants or switching constraints). KeYmaera can help finding constraints in various ways, as done in a prior case study on train control [65]. KeYmaera syntactically decomposes hybrid programs, so that only arithmetic proof obligations remain towards the leaves of a proof. These arithmetic proof obligations, if being unprovable, reveal specific valuations as counterexamples and the assumptions that were collected up to that point.

- The path leading to an unprovable goal encodes the specific location in the model; it may thus guide a domain expert to an important corner case in the model.
- A counterexample may reveal necessary initial and invariant conditions when choosing an uncontrolled dynamics model [65] as input to KeYmaera.

The discovered constraints can then be communicated to the domain expert for review before the verification continues.

4 Application Example

In this section we illustrate a verification example of an autonomous robot [45] that we collaboratively developed and solved using KeYmaera and geometric relevance filtering [56]. We compare the effort of using KeYmaera interactively, KeYmaera fully automated, and KeYmaera together with geometric relevance filtering connected via Sphinx.

With the increased introduction of autonomous robotic ground vehicles as consumer products—such as autonomous hovers and lawn mowers, or even accepting driverless cars on regular roads in California—we face an increased need for ensuring product safety not only for the good of our consumers, but also for the sake of managing manufacturer liability. One important aspect in building such systems is to make them scrutable, in order to mitigate unrealistic expectations and increase trust [72]. In the design stage of such systems, formal verification techniques ensure correct functioning w.r.t. some safety condition, and thus, increase trust. In the course of this, formal verification techniques can help to make assumptions explicit and thus clearly define what can be expected from the system under which circumstances (before the system is built and executed).

We are going to illustrate a design and verification process that encourages collaboration from high-level graphical models which convey intuition about the system to a broad and possibly heterogeneous audience to detailed formal models, which are suitable for formal verification. For this we will discuss our formal model of an autonomous robotic ground vehicle and its proof. More details on the model and case studies, as well as extensions for moving obstacles, sensor uncertainty, sensor failure, and actuator disturbance can be found in [45].

We will begin with a hierarchically structured graphical model that defines the high-level system behavior, the expected operating environment and the initial conditions under which the robot can be activated safely together with the invariants and safety conditions that the robot will then guarantee. We will complement the high-level model with a more detailed robot controller model. Together, these models are translated into $d\mathcal{L}$ and formally verified. Finally, we will discuss a simple one-dimensional model of the robot to exemplify how modeling decisions in $d\mathcal{L}$ can make verification easier.

4.1 Hierarchical Graphical Modeling

First, we construct a high-level model of the structure and the behavior of the robotic ground vehicle and of the assumptions about the environment it is operating in. Figure 8 uses the Hybrid Program UML profile to model the structure of the robotic obstacle avoidance algorithm.

The system class *World* provides a global clock c and ensures a cycle time of at most ε time units (i. e., any controller in the system will run at least once every ε time units). The state of a robot is characterized by its position (x, y) and orientation (d_x, d_y) in two dimensions and its linear velocity (v) . The robot can control its linear acceleration within certain bounds $(a \in [-B, A])$ and choose a new trajectory. It measures the position of the nearest obstacle to make decisions about its trajectory.

The high-level behavior of the robotic obstacle avoidance algorithm is modeled in a hierarchical activity diagram using our Hybrid Program UML profile. Figure 9 shows the high-level behavior with the controller and the dynamics. In this example, the dynamics is a non-linear differential-algebraic equation that describes the robot's motion on a circular segment: $x' = vd_x, y' = vd_y, d_x' = -\frac{vd_y}{r}, d_y' = \frac{vd_x}{r}, v' = a \ \& \ v \geq 0 \wedge c \leq \varepsilon$. The high-level behavior further details the initial condition under which the obstacle avoidance algorithm is safe to start and the safety condition that we want to be true for all executions (in these conditions we use $p_r = (x, y)$ to denote the position of the robot and p_o to denote the position of the obstacle in two dimensions).

A model of such high-level behavior is useful to communicate major design decisions, such as the expected operating environment and the most important constraints that the system must obey. It also consolidates more detailed models that may have been produced by different members of a verification team. As future work we will integrate composite structure diagrams, as in [9], to make the interfaces between those detailed models explicit. A more detailed model of the controller complements the high-level *ctrl* block with detailed implementation-specific decisions, as shown in Fig. 10.

The robot has three control options (top to bottom in Fig. 10): If the robot's current state is safe with respect to the sensed position of the nearest obstacle, then the robot may choose a new curve and accelerate with any rate within its physical bounds. For this, we utilize the modeling pattern introduced above: we assign an arbitrary value to the robot's acceleration state ($a := *$), which is then restricted to any value from the interval $[-B, A]$ using a test ($? - B \leq a \leq A$). The robot can brake ($a := -B$), which we want to be an emergency action that should be executed with minimal time delay (i. e., we want braking to be safe even when the robot relies on previously

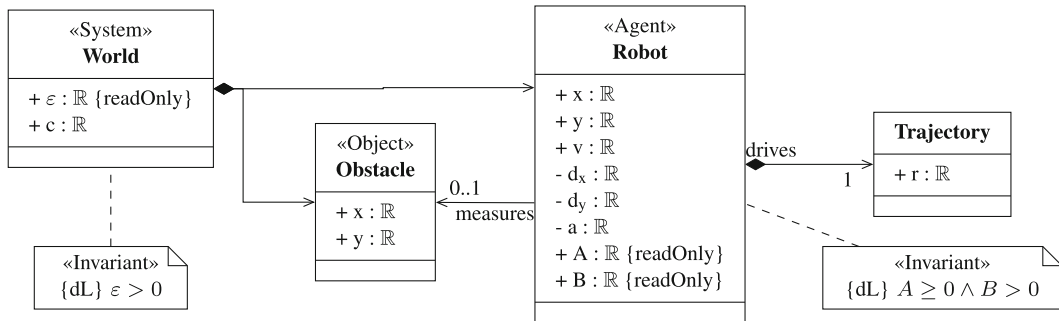


Fig. 8 The structure of the robotic obstacle avoidance model

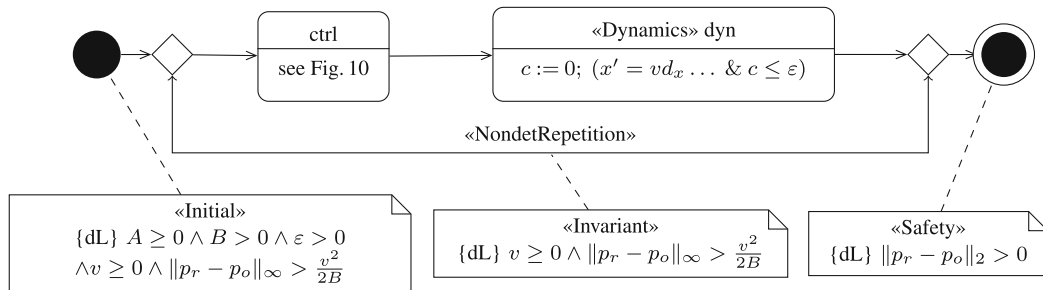


Fig. 9 Overview of the behavior of the robotic obstacle avoidance model. The model is structured hierarchically, with details on *ctrl* specified in Fig. 10

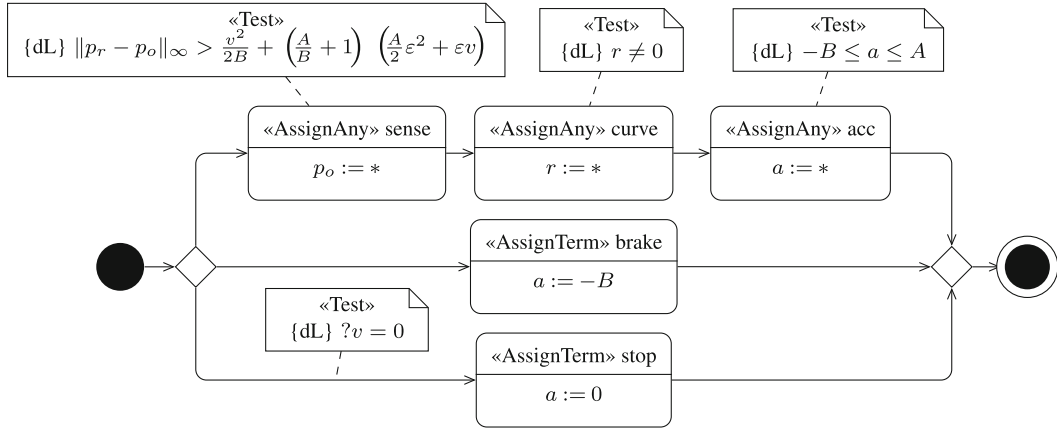


Fig. 10 The controller of the robotic obstacle avoidance model

sensed obstacle positions). Finally, if the robot is stopped ($?v = 0$), it may choose to remain in its current spot ($a := 0$).

To stay always safe the robot must account for (1) its own braking distance ($\frac{v^2}{2B}$), (2) the distance it may travel with its current velocity (ϵv) until it is able to initiate braking, and (3) the distance needed to compensate the acceleration A that may have been chosen in the worst case. For a complete model of the robotic obstacle avoidance algorithm and further variants as a hybrid program we refer to [45].

4.2 The Effect of Collaboration on Arithmetic Verification Effort

In the following paragraphs we discuss how the structure of the robotic obstacle avoidance algorithm and the resulting proof structure can be exploited to facilitate collaboration during the proof. Furthermore, we describe variants of the proof with varying degree of manual guidance and with/without collaboration using geometric relevance filtering [56]. This way, we are able to give a comparison on the proof effort that is necessary to discharge arithmetic proof obligations with and without collaboration.

The $d\mathcal{L}$ proof calculus provides proof rules to syntactically decompose a hybrid program into smaller, easier provable pieces. Such a proof unfolds into many subgoals that often can be handled separately. Proof 1 sketches the proof structure of the robot obstacle avoidance safety proof together with the proof rules used in the proof sketch.¹⁶ The names in the proof are the abbreviations that we introduced in the graphical model as placeholders for more complicated formulas, which get expanded when necessary. The three control choices of *ctrl* are transformed by the proof rule $[U]r$ into a conjunction, which is further split by the proof rule $\wedge r$ into separate branches in the proof.

These branches can be handled separately by different verification team members, who apply further proof rules of the $d\mathcal{L}$ proof calculus to continue the proof (cf. branches **expert A**, **B** and **C**). Towards the leaves of a branch the proof rules of $d\mathcal{L}$ increasingly eliminate hybrid program elements by turning them into first-order real arithmetic formulas. These formulas are often hard to prove, because the $d\mathcal{L}$ proof rules are not designed to automatically identify and eliminate unnecessary context information (e. g., ϕ still contains information about acceleration, even though it is irrelevant to prove braking safety). Quantifier elimination, which is the final step to proof correctness of a first-order real arithmetic formula, is doubly exponential in the formula size [14]. This means that we want to reduce the number of variables at the leaves of the proof as much as possible. At this stage collaboration across different verification tools becomes possible: we can ship off the formulas to an arithmetic tool or expert to discover

¹⁶ The $d\mathcal{L}$ proof calculus is explained in detail in [59,60].

Proof 1 Proof sketch of the robot obstacle avoidance algorithm using indicated proof rules

$$\begin{array}{c}
 \text{([}\cup\text{]}_r) \quad \frac{\Gamma \vdash [\alpha]\phi \wedge [\beta]\phi, \Delta}{\Gamma \vdash [\alpha \cup \beta]\phi, \Delta} \quad (\wedge_r) \quad \frac{\Gamma \vdash \phi, \Delta \quad \Gamma \vdash \psi, \Delta}{\Gamma \vdash \phi \wedge \psi, \Delta} \quad ([:]) \quad \frac{[\alpha][\beta]\phi}{[\alpha; \beta]\phi} \quad (\text{Wr}) \quad \frac{\vdash}{\vdash \phi} \quad (\text{Wl}) \quad \frac{\vdash}{\phi \vdash} \\
 \\
 \text{QE} \frac{\dots}{\tilde{\phi} \vdash \tilde{\psi}} \quad \text{B} \frac{\dots}{\phi \vdash [\text{brake}][\text{dyn}]\psi} \quad \text{C} \frac{\dots}{\phi \vdash [?v = 0; \text{stop}][\text{dyn}]\psi} \\
 \text{Wl, Wr} \frac{\dots}{\phi \dots \vdash \psi \dots} \\
 \text{expert A} \frac{\dots}{\phi \vdash [\text{sense}; \text{curve}; \text{acc}][\text{dyn}]\psi} \\
 \text{[}\cup\text{]}_r, \wedge_r \frac{\dots}{\phi \vdash [(\text{sense}; \text{curve}; \text{acc}) \cup (\text{brake}) \cup (?v = 0; \text{stop})][\text{dyn}]\psi} \\
 \text{[}] \frac{\phi \vdash [\text{ctrl}][\text{dyn}]\psi}{\phi \vdash [\text{ctrl}; \text{dyn}]\psi}
 \end{array}$$

Table 4 Proof effort in KeYmaera with and without collaboration

Variant	Proof size		Manual steps		Exported (solved)	Time (s)	Mem. (MB)
	Branches	Steps	All	Wl, Wr			
(1) Baseline	67	868	139	84	0	34.8 (2.4)	51.7
(2) Guided (interactive finish)	67	935	202	148	0	45.8 (11.6)	52.9
(3) Guided (auto finish)	Aborted after >2h						
(4) Guided (GRF finish)	67	980	108	54	18 (13)	38.1 (4)	52.2
(5) Mechanic (interactive finish)	87	1,193	440	356	0	46.9 (3.7)	52.2
(6) Mechanic (auto finish)	Aborted after >2h						
(7) Mechanic (GRF finish)	87	1,230	139	55	32 (28)	46.4 (4.2)	52.4

what information is unnecessary and can then weaken the formulas in the sequent (**Wl**, **Wr**) before we invoke the quantifier elimination procedure (**QE**).

We compare different proof variants of the robot obstacle avoidance algorithm to highlight the potential reduction in proof effort when developers with different expertise collaborate on a proof. Table 4 compares the number of proof branches, the total number of proof steps, the number of manually executed steps, the number of manually executed weaken operations, the number of exported goals and goals solved by the external tool, the proof execution duration, and the memory used during the arithmetic in the proof. The baseline (line 1 in Table 4) is a proof with manual optimization to reduce branching. We created two further variants of the proof: the guided variant manually weakened obviously unnecessary contextual information to reduce the number of branches in the proof; the mechanic variant branches fully automated by KeYmaera. Both variants were finished fully interactively (cf. lines 2 and 5 in Table 4), fully automated in KeYmaera (cf. lines 3 and 6 in Table 4), and automated with real-arithmetic formulas exported to geometric relevance filtering (cf. lines 4 and 7 in Table 4).

The interesting result is that geometric relevance filtering can solve many of the cases introduced by the fully mechanic branching, while it fails on the same highly complex problems as in the partly mechanic case. This means that the external tool directs the manual effort that is still needed in both variants to the interesting cases, while it takes care of much of the tedious work. Thus, although less manual effort was put into guiding the automated branching of KeYmaera, the effort for reducing arithmetic goals to a manageable size for quantifier elimination procedures was reduced by prior collaboration involving geometric relevance filtering [56].

4.3 The Effect of Model Variants on Proof Structure

Since it is hard to come up with a fully verifiable model that includes all the details right from the beginning, the models discussed in the previous section and in our previous case studies [45–47] are the result of different modeling

and verification variants. In the process of creating these models, different assumptions and simplifications were applied until we reached the final versions. We developed proof-aware refactoring methods to carry over verified properties about an original model to a refactored model [48], in order to reduce proof effort.

In this section, we discuss how various design decisions influence the structure of a proof and, in turn, the verification effort.

4.3.1 Modeling

We use a simplified model of a robot on a one-dimensional track [47]. In this example, navigation of a robot is considered safe, if the robot is able to stay within its assigned area (e. g., on a track) and does not actively crash with obstacles. Since we cannot guarantee reasonable behavior of obstacles, however, the robot is allowed to passively crash (i. e., while obstacles might run into the robot, the robot will never move into a position where the obstacle could not avoid a collision).

Model 1 shows a textual dL model of a hybrid system comprising the control choices of an autonomous robotic ground vehicle, the control choices of a moving obstacle, and the continuous dynamics of the system. The system represents the common controller-plant model: it repeatedly executes control choices followed by dynamics, cf. (1). The control of the robot is executed in parallel to that of the obstacle, cf. (2).

Once again, the robot has three options: it can brake unconditionally, cf. (3). If its current state is safe, as defined by (6), then the robot may accelerate with any rate within its physical bounds, cf. (4). Finally, if the robot is stopped, it may choose to remain in its current spot and may or may not change its orientation while doing so, cf. (5). This is expressed again by arbitrary assignment with subsequent test: this time, the test $?o_r^2 = 1$, however, restricts the orientation value to either forwards or backwards ($o_r \in \{1, -1\}$).

To stay safe the robot must account for the worst case braking, travel, and acceleration distance, cf. (6). This safety margin applies to either the upper or the lower bound of the robot's area, depending on the robot's orientation: when driving forward (i. e., towards the upper bound), we do not need a safety margin towards the lower bound, and vice versa. This is expressed by the factors $\frac{1-o_r}{2}$ and $\frac{1+o_r}{2}$, which mutually evaluate to zero (e. g., $\frac{1-o_r}{2} = 0$ when driving forward with $o_r = 1$). The distance between the robot and the obstacle must be large enough to (1) allow the robot to brake to a stand-still, (2) compensate its current velocity and worst-case acceleration, and (3) account

Model 1 Single wheel drive without steering (one-dimensional robot navigation)

$$swd \equiv (ctrl; dyn)^* \tag{1}$$

$$ctrl \equiv (ctrl_r \parallel ctrl_o) \tag{2}$$

$$ctrl_r \equiv (a_r := -B) \tag{3}$$

$$\cup (?safe; a_r := *; ? - B \leq a_r \leq A) \tag{4}$$

$$\cup (?v_r = 0; a_r := 0; o_r := *; ?o_r^2 = 1) \tag{5}$$

$$safe \equiv x_b + \frac{1-o_r}{2} \left(\frac{v_r^2}{2B} + \left(\frac{A}{B} + 1 \right) \left(\frac{A}{2} \varepsilon^2 + \varepsilon v_r \right) \right) < x_r < x_b - \frac{1+o_r}{2} \left(\frac{v_r^2}{2B} + \left(\frac{A}{B} + 1 \right) \left(\frac{A}{2} \varepsilon^2 + \varepsilon v_r \right) \right) \tag{6}$$

$$\wedge \|x_r - x_o\| \geq \frac{v_r^2}{2B} + \left(\frac{A}{B} + 1 \right) \left(\frac{A}{2} \varepsilon^2 + \varepsilon v_r \right) + V \left(\varepsilon + \frac{v_r + A\varepsilon}{B} \right) \tag{7}$$

$$ctrl_o \equiv \left(?v_o = 0; o_o := *; ?o_o^2 = 1 \right) \tag{8}$$

$$\cup (v_o := *; ?0 \leq v_o \leq V) \tag{9}$$

$$dyn \equiv (t := 0; x'_r = o_r v_r, v'_r = a_r, x'_o = o_o v_o, t' = 1 \ \& \ v_r \geq 0 \wedge v_o \geq 0 \wedge t \leq \varepsilon) \tag{10}$$

for the obstacle moving towards the robot with worst-case velocity V while the robot is still not stopped, cf. (7). Note, that we have to be more conservative towards the obstacle than towards the bounds, because we want to be able to stop even when the obstacle approaches the robot from behind.

The obstacle, essentially, has similar control options as the robot (with the crucial difference of not having to care about safety): it may either remain in a spot and possibly change its orientation (8), or choose any velocity up to V , cf. (9).

4.3.2 Verification

We verify the acceleration and orientation choices as modeled in Model 1 above are safe, using a formal proof calculus for $\text{d}\mathcal{L}$ [57,59]. The robot is safely within its assigned area and at a safe distance to the obstacle, if it is able to brake to a complete stop at all times.¹⁷ The following condition captures this requirement as an invariant $\langle r \text{ stoppable } (o, b) \rangle$ that we want to hold at all times during the execution of the model:

$$\begin{aligned} r \text{ stoppable } (o, b) \equiv & \|x_r - x_o\| \geq \frac{v_r^2}{2B} + \frac{v_o V}{b} \wedge x_{\underline{b}} + \frac{1 - o_r}{2} \frac{v_r^2}{2B} < x_r < x_{\bar{b}} - \frac{1 + o_r}{2} \frac{v_r^2}{2B} \\ & \wedge v_r \geq 0 \wedge o_r^2 = 1 \wedge o_o^2 = 1 \wedge 0 \leq v_o \leq V \end{aligned}$$

The formula $\langle r \text{ stoppable } (o, b) \rangle$ states that the distance between the robot to both the obstacle and the bounds is safe, if there is still enough distance for the robot to brake to a complete stop before it reaches either. Also, the robot must drive with non-negative velocity, the chosen directions of robot and obstacle must be either forwards ($o_r = 1$) or backwards ($o_r = -1$), and the obstacle must use only non-negative velocities up to V .

Theorem 1 (Safety of single wheel drive) *If a robot is inside its assigned area and at a safe distance from the obstacle's position x_o initially, then it will not actively collide with the obstacle and stay within its area while it follows the swd control model (Model 1), as expressed by the provable $\text{d}\mathcal{L}$ formula:*

$$r \text{ stoppable } (o, b) \rightarrow [\text{swd}]((v_r > 0 \rightarrow \|p_r - p_o\| > 0) \wedge x_{\underline{b}} < x_r < x_{\bar{b}}) \quad (11)$$

We proved Theorem 1 using KeYmaera. With respect to making autonomous systems more scrutable, such a proof may help in a twofold manner: on the one hand, it may increase trust in the implemented robot (given the assumption that the actual implementation and execution can be traced back to the abstract model). On the other hand, it makes the behavior of the robot more understandable. In this respect, the most interesting properties of the proven model are the definition of *safe* and the invariant, which allow us to analyze design trade-offs and tell us what is always true about the system regardless of its state. As an example, let us consider the distance between the robot and the obstacle that is considered safe: $\|x_r - x_o\| \geq \frac{v_r^2}{2B} + (\frac{A}{B} + 1)(\frac{A}{2}\varepsilon^2 + \varepsilon v_r) + V(\varepsilon + \frac{v_r + A\varepsilon}{B})$. This distance can be interpreted as the minimum distance that the robot's obstacle detection sensors are required to cover (e. g., as done in [46]); it is a function of other robot design parameters (maximum velocity, braking power, worst-case acceleration, sensor/processor/actuator delay) and the parameters expected in the environment (obstacle velocity). The distance $\|x_r - x_o\|$ can be optimized w.r.t. different aspects: for example, to find the most cost-efficient combination of components that still guarantees safety, to specify a safe operation environment given a particular robot configuration, or to determine time bounds for algorithm optimization.

With respect to the manual guidance and collaboration needed in such a proof, we had to apply knowledge in hybrid systems and in-depth understanding of the robot model to find a system invariant, which is the most important manual step in the proof above. We further used arithmetic interactions, such as the hiding of superfluous terms

¹⁷ The requirement that the robot has to ensure an option for the obstacle to avoid a collision is ensured trivially, since the obstacle in this model can choose its velocity directly. In a more realistic model the obstacle would choose acceleration instead; then the robot had to account for the braking distance of the obstacle, too.

to reduce arithmetic complexity, transforming and replacing terms (e. g., substitute the absolute function with two cases, one for negative and one for positive values).

4.3.3 The Proof Structure of Model Variants

We now want to discuss the proof structure of different model variants: For example, one can make explicit restrictions on particular variables, such as first letting the robot start in a known direction (instead of an arbitrary direction). Such assumptions and simplifications, of course, are not without implications on the proof. While in some aspect a proof may become easier, it may become more laborious or more complex in another. In this section, we discuss five variants of the single wheel drive model (without obstacle) to demonstrate implications on the proof structure and on the entailed manual guidance needed to complete a proof in KeYmaera.

The following model variants are identical in terms of the behavior of the robot. However, assumptions on the starting direction were made in the antecedent of a provable $d\mathcal{L}$ formula, and the starting direction as well as the orientation of the robot were explicitly distinguished by disjunction or non-deterministic choice, or implicitly encoded in the arithmetic, as described below (we denote a changed formula by using primed versions of the original formula reference).

Assumed Starting Direction, Orientation by Disjunction In the first variant, the robot is assumed to start in a known direction, specified in the antecedent of (11').

$$o_r = 1 \dots \rightarrow [swd](x_b < x_r < x_b) . \quad (11')$$

Also, the robot had an explicit choice on turning during stand-still in (5').

$$ctrl_r \equiv \dots \cup (?v_r = 0; a_r := 0; o_r := -o_r) \cup (?v_r = 0; a_r := 0) \quad (5')$$

The orientation of the robot is explicitly distinguished by disjunction in (6').

$$safe \equiv (o_r = -1 \wedge x_b + \dots < x_r) \vee (o_r = 1 \wedge x_r < x_b - \dots) . \quad (6')$$

Orientation by Arithmetic In the second variant, we kept the assumed starting direction of the first variant. However, the orientation by disjunction in the definition of *safe* was replaced by using o_r as discriminator value encoded in the arithmetic, resulting in (6) of Model 1.

Arbitrary Starting Direction by Disjunction The third variant relaxes the assumption on the starting direction by introducing a disjunction of possible starting directions in the antecedent of formula (11') to get (11'').

$$(o_r = 1 \vee o_r = -1) \dots \rightarrow [swd](x_b < x_r < x_b) \quad (11'')$$

Arbitrary Starting Direction by Arithmetic The fourth variant replaces the disjunction in the antecedent of (11'') by stating the two orientation options as $o_r^2 = 1$ to get the antecedent in (11).

Replace Non-Deterministic Choice with Arithmetic Finally, we replace the non-deterministic turning choice of (5') with $(?v_r = 0; a_r := 0; o_r := *; ?o_r^2 = 1)$ to get (5) of Model 1. This last variant proves correctness of (11) with *swd* as in Model 1.

Table 5 summarizes the proof structures of the five variants.

Unsurprisingly—when considering the rules of the $d\mathcal{L}$ proof calculus [58] as listed in Table 5—disjunctions in the antecedent ($\vee l$) or in tests of hybrid programs, as well as non-deterministic choices ($[U]$) increase the number of proof branches and with it the number of manual proof steps. Verification experts, who are familiar with the proof

Table 5 Nodes, branches, and manual proof steps of variants [47]

	Variant	Nodes	Branches	Manual steps	Avoids
(i)	Assumed starting direction, orientation by disjunction	387	34	24	
(ii)	Orientation by arithmetic	331	28	25	($\vee I$)
(iii)	Arbitrary starting direction by disjunction	650	56	44	
(iv)	Arbitrary starting direction by arithmetic	185	17	22	($\vee I$)
(v)	Replace non-deterministic choice	160	14	29	($[\cup](\wedge r)$)

	$\frac{\Gamma, \phi \vdash \Delta \quad \Gamma, \psi \vdash \Delta}{\Gamma, \phi \vee \psi \vdash \Delta}$	($\vee I$)
--	--	--------------

calculus and the branching behavior of KeYmaera, can in some cases express the same model with an alternative encoding using arithmetic expressions. The number of proof branches can be reduced, if we can replace disjunctions in the antecedent (but also conjunctions in the consequent) or non-deterministic choices in the hybrid program by an equivalent arithmetic encoding. Conversely, this means that some arithmetic problems can be traded for easier ones with additional proof branches.

5 Conclusion

In this paper, we gave a vision of a verification-driven engineering toolset including hybrid and arithmetic verification tools, and introduced modeling and collaboration tools with the goal of making formal verification of hybrid systems accessible to a broader audience. The current implementation of Sphinx¹⁸ includes a textual and a graphical modeling editor for $d\mathcal{L}$; automated transformation from the graphical models to textual models; simulation of the valuations of variables in hybrid programs over time; integration of KeYmaera as a hybrid systems verification tool; model and proof comparison; and connection to various collaboration backend systems.

We applied the tool suite to examples based on prior case studies on robotic obstacle avoidance [45]. We made the following two main observations on the effects of collaboration.

Arithmetic Verification Although less manual effort was put into guiding the automated branching of KeYmaera, the effort for reducing arithmetic goals to a manageable size for quantifier elimination procedures was reduced by prior collaboration involving geometric relevance filtering [56].

Domain Model Decisions The effect of some modeling decisions—which may come natural to domain experts, such as using disjunction or non-deterministic choice for modeling logical or program alternatives—on the verification effort can be reduced by verification experts that are familiar with the proof calculus by introducing an alternative arithmetic encoding.

As a vision of extending collaboration support, it is planned to integrate Wikis and other online collaboration tools (currently, we use Redmine both as project management repository and for knowledge exchange) for exchanging knowledge on proof tactics. Additionally, collaboration with experts outside the own organization can be fostered by linking to Web resources, such as MathOverflow.¹⁹ In such a platform, requests could be forwarded to those experts whose knowledge matches the verification problem best.

Another interesting research direction are refactoring operations, which systematically construct incremental model variants without the need for re-verification of the entire model [48]. In a naive way, after a refactoring was applied we would have to reprove all properties about a model. But often a refactoring operation changes only

¹⁸ Available for download at <http://www.cs.cmu.edu/~smitsch/sphinx.html>.

¹⁹ mathoverflow.net.

fragments of a model while it leaves the remainder of the model untouched, or the refactored model and properties are systematically reducible to previous proofs by side deduction. A refactoring operation should systematically reduce verification effort by creating new artifacts that are less effort to prove than the complete model. In a verification-driven engineering process, a refactoring operation creates (1) a refactored model and properties with links to the original model and a description of the applied refactoring; (2) a correctness conjecture, together with verification tickets in the project management backend. This way, all changes applied to models and their properties via refactoring operations can be provably traced back.

The VDE toolset is currently being tested in a collaborative verification setting between Carnegie Mellon University, the University of Cambridge, and the University of Edinburgh.

Acknowledgments This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246, NSF EXPEDITION CNS-0926181, and under Grants CNS-1035800 and CNS-0931985, by DARPA under agreement number FA8750-12-2-0291, and by the US Department of Transportation's University Transportation Center's TSET grant, award# DTRT12GUTC11. Passmore was also supported by the UK's EPSRC [Grants Numbers EP/I011005/1 and EP/I010335/1]. This work was also supported by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under grant FFG FIT-IT 829598, FFG BRIDGE 838526, and FFG Basisprogramm 838181. The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n° P10F-GA-2012-328378.

References

1. Akbarpour, B., Paulson, L.C.: MetiTarski: an automatic theorem prover for real-valued special functions. *J. Autom. Reason.* **44**(3), 175–205 (2010)
2. Alur, R.: Formal verification of hybrid systems. In: Chakraborty, S., Jerraya, A., Baruah, S.K., Fischmeister, S., (eds.) *Proceedings of the 11th International Conference on Embedded Software (EMSOFT)*, pp. 273–278. ACM (2011)
3. Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T.A., Ho, P.-H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. *Theor. Comput. Sci.* **138**(1), 3–34 (1995)
4. Anda, B., Hansen, K., Gullesten, I., Thorsen, H.K.: Experiences from introducing UML-based development in a large safety-critical project. *Empir. Softw. Eng.* **11**(4), 555–581 (2006)
5. Bajaj, M., Scott, A., Deming, D., Wickstrom, G., Spain, M.D., Zwemer, D., Peak, R.: Maestro—a model-based systems engineering environment for complex electronic systems. In: *Proceedings of the 22nd Annual INCOSE International Symposium*. INCOSE, Rome (2012)
6. Barrett, C., Stump, A., Tinelli, C.: The SMT-LIB standard: Version 2.0 (2012). <http://smtlib.cs.uiowa.edu/papers/smt-lib-reference-v2.0-r12.09.09.pdf> (last accessed 2013-01-09)
7. Baumgartner, N., Mitsch, S., Müller, A., Salfinger, A., Retschitzegger, W., Schwinger, W.: A tour of BeAware: a situation awareness framework for control centers. *Inf. Fusion* (2014). doi:10.1016/j.inffus.2014.01.008
8. Belta, C., Ivancic, F., (eds.): *Hybrid Systems: Computation and Control* (part of CPS Week 2013), HSCC'13, ACM, Philadelphia (2013)
9. Berkenkötter, K., Bisanz, S., Hannemann, U., Peleska, J.: The HybridUML profile for UML 2.0. *STTT* **8**(2), 167–176 (2006)
10. Christakis, M., Müller, P., Wüstholtz, V.: Collaborative verification and testing with explicit assumptions. In: Giannakopoulou, D., Méry, D. (eds.) *FM*, Volume 7436 of LNCS, pp. 132–146. Springer, Berlin (2012)
11. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The MathSAT5 SMT solver. In: Piterman, N., Smolka, S.A. (eds.) *TACAS, Lecture Notes in Computer Science*, vol. 7795, pp. 93–107. Springer, Berlin (2013)
12. Collins, P., Lygeros, J.: Computability of finite-time reachable sets for hybrid systems. In: 44th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC), pp. 4688–4693. IEEE (2005)
13. Craigen, D., Gerhart, S.L., Ralston, T.: Formal methods reality check: industrial usage. *IEEE Trans. Softw. Eng.* **21**(2), 90–98 (1995)
14. Davenport, J.H., Heintz, J.: Real quantifier elimination is doubly exponential. *J. Symb. Comput.* **5**(1–2), 29–35 (1988)
15. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS, Lecture Notes in Computer Science*, vol. 4963, pp. 337–340. Springer, Berlin (2008)
16. De Schutter, B., Heemels, W., Lunze, J., Priour, C.: Survey of modeling, analysis, and control of hybrid systems. In: Lunze, J., Lamnabhi-Lagarriague, F. (eds.) *Handbook of Hybrid Systems Control—Theory, Tools, Applications*, chap. 2, pp. 31–55. Cambridge University Press, Cambridge (2009)
17. Derler, P., Lee, E.A., Sangiovanni-Vincentelli, A.: Modeling cyber-physical systems. *Proc. IEEE* **100**(1), 13–28 (2012)
18. Deshpande, A., Göllü, A., Varaiya, P.: Shift: A formalism and a programming language for dynamic networks of hybrid automata. In: Antsaklis, P.J., Kohn, W., Nerode, A., Sastry, S., (eds.) *Hybrid Systems, Lecture Notes in Computer Science*, vol. 1273, pp. 113–133. Springer, Berlin (1996)

19. Faber, J., Linker, S., Olderog, E.-R., Quesel, J.-D.: Syspect—modelling, specifying, and verifying real-time systems with rich data. *Int. J. Softw. Inf.* **5**(1–2), 117–137 (2011)
20. Frehse, G.: PHAVer: algorithmic verification of hybrid systems past HyTech. In: Morari, M., Thiele, L., (eds.) *Hybrid Systems: Computation and Control*, 8th International Workshop, HSCC 2005, Zurich, Proceedings, LNCS, vol. 3414, pp. 258–273. Springer, Berlin (2005)
21. Frehse, G., Le Guernic, C., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: SpaceEx: scalable verification of hybrid systems. In: Ganesh Gopalakrishnan, S.Q. (ed.) *CAV*, LNCS. Springer, Berlin (2011)
22. Gokhale, A.S., Balasubramanian, K., Krishna, A.S., Balasubramanian, J., Edwards, G., Deng, G., Turkey, E., Parsons, J., Schmidt, D.C.: Model driven middleware: a new paradigm for developing distributed real-time and embedded systems. *Sci. Comput. Program.* **73**(1), 39–58 (2008)
23. Goswami, D., Schneider, R., Masrur, A., Lukasiewicz, M., Chakraborty, S., Voit, H., Annaswamy, A.: Challenges in automotive cyber-physical systems design. In: *ICSAMOS*, pp. 346–354. IEEE (2012)
24. Gowers, T., Nielsen, M.: Massively collaborative mathematics. *Nature* **461**, 879–881 (2009)
25. Hales, T.C., Harrison, J., McLaughlin, S., Nipkow, T., Obua, S., Zumkeller, R.: A revision of the proof of the Kepler conjecture. *Discrete Comput. Geom.* **44**(1), 1–34 (2010)
26. Hause, M.C., Thom, F.: An integrated MDA approach with SysML and UML. In: *Proceedings of the 13th International Conference on Engineering of Complex Computer Systems, ICECCS '08*, pp. 249–254. IEEE Computer Society, Washington (2008)
27. Henzinger, T.A., Jhala, R., Majumdar, R., Sanvido, M.A.A.: Extreme model checking. In: Dershowitz, N., (ed.) *Verification: Theory and Practice, Essays Dedicated to Zohar Manna on the Occasion of His 64th Birthday*, LNCS, vol. 2772, pp. 332–358. Springer, Berlin (2003)
28. Hitz, M., Kappel, G., Kapsammer, E., Retschitzegger, W.: *UML @ Work*. dpunkt (2005)
29. Jouault, F., Allilaire, F., Bézivin, J., Kurtev, I.: ATL: a model transformation tool. *Sci. Comput. Program.* **72**(1–2), 31–39 (2008)
30. Kent, S.: Model driven engineering. In: Butler M.J., Petre L., Sere K., (eds.) *IFM*, LNCS, vol. 2335. pp. 286–298. Springer, Berlin (2002)
31. Kerber, M., Lange, C., Rowat, C., (eds.): *Enabling Domain Experts to use Formalised Reasoning—Symposium AISB, Do-Form 2013*, Exeter. Proceedings. AISB (2013)
32. Klein, G., Elphinstone, K., Heiser, G., Andronick, J., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., Winwood, S.: seL4: formal verification of an OS kernel. In: *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*, pp. 207–220. ACM, New York (2009)
33. Kokar, M.M., Matheus, C.J., Baclawski, K.: Ontology-based situation awareness. *Int. J. Inf. Fusion* **10**(1), 83–98 (2009)
34. Kolovos, D.S., Di Ruscio, D., Pierantonio, A., Paige, R.F.: Different models for model matching: An analysis of approaches to support model differencing. In: *Proceedings of the 2009 ICSE Workshop on Comparison and Versioning of Software Models, CVSM '09*, pp. 1–6. IEEE Computer Society, Washington (2009)
35. Kordon, F., Hugues, J., Renault, X.: From model driven engineering to verification driven engineering. In: *Proc. of the 6th IFIP Int. Workshop on Software Technologies for Embedded and Ubiquitous Systems*, pp. 381–393. Springer, Berlin (2008)
36. Kouskoulas, Y., Renshaw, D., Platzer, A., Kazantzides, P.: Certifying the safe design of a virtual fixture control algorithm for a surgical robot. In: *Belta and Ivancic* [8]
37. Kupferman, O., Vardi, M.Y.: Modular model checking. In: *Revised Lectures from the International Symposium on Compositionality: The Significant Difference, COMPOS'97*, pp. 381–401. Springer, London (1998)
38. Lavazza, L., Quaroni, G., Venturelli, M.: Combining UML and formal notations for modelling real-time systems. In: *ESEC/SIGSOFT FSE*, pp. 196–206. ACM (2001)
39. Lee, I., Sokolsky, O., Chen, S., Hatcliff, J., Jee, E., Kim, B., King, A.L., Mullen-Fortino, M., Park, S., Roederer, A., Venkatasubramanian, K.K.: Challenges and research directions in medical cyber-physical systems. *Proc. IEEE* **100**(1), 75–90 (2012)
40. Liu, J., Liu, Z., He, J., Mallet, F., Ding, Z.: Hybrid MARTE statecharts. *Front. Comput. Sci.* **7**(1), 95–108 (2013)
41. Loos, S.M., Platzer, A.: Safe intersections: at the crossing of hybrid systems and verification. In: Yi, K. (ed.) *ITSC*, pp. 1181–1186 (2011)
42. Loos, S.M., Platzer, A., Nistor, L.: Adaptive cruise control: Hybrid, distributed, and now formally verified. In: *FM*, LNCS, vol. 6664, pp. 42–56. Springer, Berlin (2011)
43. Loos, S.M., Renshaw, D., Platzer, A.: Formal verification of distributed aircraft controllers. In: *Belta and Ivancic* [8]
44. Mallet, F., de Simone, R.: MARTE: a profile for RT/E systems modeling, analysis—and simulation? In: Molnár, S., Heath, J.R., Dalle, O., Wainer, G.A. (eds.) *SimuTools*, p. 43. ICST (2008)
45. Mitsch, S., Ghorbal, K., Platzer, A.: On provably safe obstacle avoidance for autonomous robotic ground vehicles. In: *Robotics: Science and Systems* (2013)
46. Mitsch, S., Loos, S.M., Platzer, A.: Towards formal verification of freeway traffic control. In: Lu, C. (ed.) *Proc. of the 2nd Int. Conference on Cyber-Physical Systems (ICCPs)*, pp. 171–180. IEEE (2012)
47. Mitsch S., Passmore, G.O., Platzer, A.: A vision of collaborative verification-driven engineering of hybrid systems. In: Kerber et al. [31], pp. 8–17
48. Mitsch, S., Quesel, J.-D., Platzer, A.: Refactoring, refinement, and reasoning—a logical characterization for hybrid systems. In: *FM*, LNCS. Springer, Berlin (2014, to appear)
49. Mostowski, W.: The KeY syntax. In: Beckert, B., Hähnle, R., Schmitt, P.H., (eds.) *Verification of Object-Oriented Software. The KeY Approach*, Lecture Notes in Computer Science, vol. 4334, pp. 599–626. Springer, Berlin (2007)

50. Mota, E., Clarke, E.M., Groce, A., Oliveira, W., Falcão, M., Kanda, J.: VeriAgent: an approach to integrating UML and formal verification tools. *Electr. Notes Theor. Comput. Sci.* **95**, 111–129 (2004)
51. Niemueller, T., Ewert, D., Reuter, S., Karras, U., Ferrein, A., Jeschke, S., Lakemeyer, G.: Towards benchmarking cyber-physical systems in factory automation scenarios. In: Timm, I.J., Thimm, M., (eds.) *KI, Lecture Notes in Computer Science*, vol. 8077. Springer, Berlin, pp. 296–299 (2013)
52. Niles, I., Pease, A.: Towards a standard upper ontology. In: *Proc. of the 2nd Int. Conf. on Formal Ontology in Information Systems (FOIS '01)*, pp. 2–9. ACM, Ogunquit (2001)
53. Object Management Group. *OMG object constraint language (OCL)*. Technical Report formal/2012-01-01, OMG (2012)
54. Passmore, G.O.: *Combined Decision Procedures for Nonlinear Arithmetics, Real and Complex*. PhD thesis, University of Edinburgh (2011)
55. Passmore, G.O., Paulson, L.C., de Moura, L. M.: Real algebraic strategies for MetiTarski proofs. In: Jeuring, J., Campbell, J.A., Carette, J., Reis, G.D., Sojka, P., Wenzel, M., Sorge, V., (eds.) *AISC/MKM/Calculemus, LNCS*, vol. 7362, pp. 358–370. Springer, Berlin (2012)
56. Passmore, G.O., Platzer, A., Zawadzki, E., Avigad, J.: *Geometric relevance filtering for real closed field arithmetic* (2013, in preparation)
57. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reas.* **41**(2), 143–189 (2008)
58. Platzer, A.: Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.* **20**(1), 309–352 (2010)
59. Platzer, A.: *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg (2010)
60. Platzer, A.: The complete proof theory of hybrid systems. In: *LICS*, pp. 541–550. IEEE (2012)
61. Platzer, A.: Logics of dynamical systems. In: *LICS*, pp. 13–24. IEEE (2012)
62. Platzer, A., Clarke, E.M.: Computing differential invariants of hybrid systems as fixedpoints. *Formal Methods Syst. Design* **35**(1), 98–120 (2009)
63. Platzer, A., Clarke, E.M.: Formal verification of curved flight collision avoidance maneuvers: a case study. In: Cavalcanti, A., Dams, D. (eds.) *FM, LNCS*, vol. 5850, pp. 547–562. Springer, Berlin (2009)
64. Platzer, A., Quesel, J.-D.: KeYmaera: a hybrid theorem prover for hybrid systems. In: Armando, A., Baumgartner, P., Dowek, G. (eds.) *IJCAR, LNCS*, vol. 5195, pp. 171–178. Springer, Berlin (2008)
65. Platzer, A., Quesel, J.-D.: European train control system: a case study in formal verification. In: Breitman, K., Cavalcanti, A. (eds.) *ICFEM, LNCS*, vol. 5885, pp. 246–265. Springer, Berlin (2009)
66. Razali, R., Snook, C.F., Poppleton, M.R.: Comprehensibility of UML-based formal model: a series of controlled experiments. In: *Proceedings of the 1st ACM International Workshop on Empirical Assessment of Software Engineering Languages and Technologies: Held in Conjunction with the 22Nd IEEE/ACM International Conference on Automated Software Engineering (ASE), WEASEL Tech'07*, pp. 25–30. ACM, New York (2007)
67. Reiter, R.: *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Cambridge (2001)
68. Schäfer, W., Wehrheim, H.: Model-driven development with Mechatronic UML. In: Engels, G., Lewerentz, C., Schäfer, W., Schürr, A., Westfechtel, B. (eds.) *Graph Transformations and Model-Driven Engineering, Lecture Notes in Computer Science*, vol. 5765, pp. 533–554. Springer, Berlin (2010)
69. Schmidt, D.C.: Guest editor's introduction: model-driven engineering. *IEEE Comput.* **39**(2), 25–31 (2006)
70. Snook, C.F., Butler, M.J.: UML-B: formal modeling and design aided by UML. *ACM Trans. Softw. Eng. Methodol.* **15**(1), 92–122 (2006)
71. Sridhar, S., Hahn, A., Govindarasu, M.: Cyber-physical system security for the electric power grid. *Proc. IEEE* **100**(1), 210–224 (2012)
72. Tintarev, N., Oren, N., Deemter, K.V., Kutlak, R., Green, M., Masthoff, J., Vasconcelos, W.: SAsSy—scrutable autonomous systems. In: Kerber et al. [31], pp. 1–3
73. Tomlin, C., Pappas, G., Sastry, S.: Conflict resolution for air traffic management: a study in multiagent hybrid systems. *IEEE Trans. Autom. Control* **43**(4), 509–521 (1998)
74. Woodcock, J., Larsen, P. G., Bicarregui, J., Fitzgerald, J.: Formal methods: practice and experience. *ACM Comput. Surv.* **41**(4), 19:1–19:36 (2009)