

# A Hybrid Programming Language for Formal Modeling and Verification of Hybrid Systems

Eduard Kamburjan<sup>1,3</sup>, Stefan Mitsch<sup>2</sup>, and Reiner Hähnle<sup>3</sup>

1 Department of Informatics, University of Oslo, Norway  
eduard@ifi.uio.no

2 Computer Science Department, Carnegie Mellon University, USA  
smitsch@cs.cmu.edu

3 Department of Computer Science, Technische Universität Darmstadt, Germany  
haehnle@cs.tu-darmstadt.de

## Abstract

Designing and modeling complex cyber-physical systems (CPS) faces the double challenge of combined discrete-continuous dynamics and concurrent behavior. Existing formal modeling and verification languages for CPS expose the underlying proof search technology. They lack high-level structuring elements and are not efficiently executable. The ensuing modeling gap renders formal CPS models hard to understand and to validate. We propose a high-level *programming*-based approach to formal

modeling and verification of hybrid systems as a hybrid extension of an Active Objects language. Well-structured hybrid active programs and requirements allow automatic, reachability-preserving translation into differential dynamic logic, a logic for hybrid (discrete-continuous) programs. Verification is achieved by discharging the resulting formulas with the theorem prover KeYmaera X. We demonstrate the usability of our approach with case studies.

**2012 ACM Subject Classification** Distributed programming languages, Model verification and validation, Logic and verification, Timed and hybrid models

**Keywords and phrases** Active Objects, Differential Dynamic Logic, Hybrid Systems

**Digital Object Identifier** 10.4230/LITES.xxx.yyy.p

**Received** Date of submission. **Accepted** Date of acceptance. **Published** Date of publishing.

**Editor** LITES section area editor

## 1 Introduction

Networked cyber-physical systems (CPS) are a main driving force of innovation in computing, from manufacturing to everyday appliances. But to design and model such systems poses a double challenge: first, their *hybrid* nature, with both continuous physical dynamics and complex computations in discrete time steps. Second, their *concurrent* nature: distributed, active components (sensors, actuators, controllers) execute simultaneously and communicate asynchronously. It is notoriously difficult to get CPS models right. *Formal* modeling languages, including hybrid automata [5], hybrid process algebra [27], and logics for hybrid programs [65], can be used to formally verify properties of CPS. Contrary to simulation frameworks, such as Ptolemy [71] or Simulink, however, these languages were *designed for verification* and are based on concepts of the underlying verification technology: automata, algebras, formulas. Their minimalist syntax lacks standard structuring elements of programming languages such as types, scopes, methods, complex commands, futures, etc. Thus it is hard to adequately represent concurrently executing, communicating, hybrid components with *symbolic* data structures and computations, for example, servers or cloud applications.

Moreover, “low-level” models are hard to *validate*, i.e. to ensure that a CPS model reflects the designer’s intention, because these formalisms are not (efficiently) executable. To bridge



© Eduard Kamburjan, Stefan Mitsch and Reiner Hähnle;  
licensed under Creative Commons License CC-BY

Leibniz Transactions on Embedded Systems, Vol. XXX, Issue YYY, pp. 1–34



Leibniz Transactions on Embedded Systems


Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

18 the modeling gap we propose a high-level *programming*-based approach to formal modeling and  
 19 verification of hybrid systems.

20 The basis of our approach is an *Active Objects* (AO) language [29] called **ABS** [48]. AO languages  
 21 combine OO programming with strong encapsulation as well as asynchronous, parallel execution.  
 22 Their concurrency model permits to decompose concurrent execution into sequential execution in  
 23 a compositional manner. We chose **ABS** for its formal semantics, its open source implementation  
 24 tool chain, and its demonstrated scaling on massively distributed systems [75], but our approach  
 25 is applicable to other AO languages. **ABS** is efficiently executable via compilation to **ERLANG** and  
 26 was used to model complex, real-world systems for cloud processing [3], virtualized services [49],  
 27 data processing [56], and railway operations [53]. However, it lacks the capability to model hybrid  
 28 systems. The *first main contribution* of this paper is the design of the *Hybrid ABS* (**HABS**) language,  
 29 a conservative (syntax and semantics preserving) extension of **ABS**, generalizing the Active Objects  
 30 paradigm to *Hybrid Active Objects* (**HAO**): AO with continuous dynamics. Obviously, it is  
 31 necessary to accordingly extend the formal semantics of **ABS** and its runtime environment. This is  
 32 our *second main contribution*. Our *third main contribution* is the implementation of **HABS** and a  
 33 formal verification tool for it.

34 Our approach to formal verification of **HABS** programs is based on reachability-preserving  
 35 translation into an existing verification formalism for hybrid programs. We choose differential  
 36 dynamic logic ( $d\mathcal{L}$ ) [66, 68, 69], as implemented in the KeYmaera X system [36], because it is  
 37 based on an imperative programming language that is a good match for the sequential fragment  
 38 of **HABS** and verification in  $d\mathcal{L}$  has been demonstrated to scale to realistic systems (e.g., [47]). The  
 39 translation from **HABS** to  $d\mathcal{L}$  involves to decompose a given **HABS** verification problem into a set of  
 40 independent *sequential*  $d\mathcal{L}$  problems. This is possible, because we impose an interaction pattern  
 41 for communication on **HABS** that is less restrictive than available component-based techniques [64],  
 42 yet is general enough to permit intuitive and concise modeling of relevant case studies. The  
 43 identification of this pattern, the generation of  $d\mathcal{L}$  verification conditions, and a reachability  
 44 preservation theorem constitute our *fourth main contribution*.

45 The overall approach is illustrated in Fig. 1: A CPS is modeled as an **HABS** program with the  
 46 aim to analyze its properties statically. One formulates desired properties as invariants that are  
 47 formally verified to hold under certain assumptions. Before verification is attempted, the model is  
 48 *validated* by executing it in the runtime environment to ensure that it behaves as intended. A  
 49 visualization component helps to analyze behavior over time. Subsequently, the verification claim  
 50 is automatically decomposed and translated into a set of  $d\mathcal{L}$  verification problems discharged in  
 51 KeYmaera X (optionally, formally verified runtime monitors [63] and formally verified machine  
 52 code is available from KeYmaera X through VeriPhy [18]). Both, unexpected runtime behavior  
 53 and failed verification attempts, serve to fix the model and/or the claimed properties.



■ **Figure 1** Structure of **HABS** workflow.

54 The paper is structured as follows. Sect. 2 gives an informal example of an **HABS** model with a  
 55 distributed water tank controller. Sect. 3 formally defines syntax and semantics of **HABS**. Sect. 4  
 56 describes modeling patterns. Sect. 5 gives theoretical background on  $d\mathcal{L}$ , the translation into  $d\mathcal{L}$ ,

57 the decomposition theorem, and tells how to prove correctness. It also contains a distributed  
58 controller case study. Finally, Sect. 6 discusses related and future work and concludes.

## 59 **2** Distributed Hybrid Systems by Example

60 Active Objects [29] are objects that realize actor-based concurrency [44] with futures [28] and  
61 cooperative scheduling: Active Objects communicate via asynchronous method calls. On the  
62 caller side, each method invocation generates a future as a handle to retrieve the call's result,  
63 once it is available. The caller may synchronize on that future, i.e. suspend and wait until it is  
64 resolved. At most one process is running on an Active Object at any time. That process suspends  
65 when it encounters the synchronization statement `await` on an unresolved future or a false Boolean  
66 condition. Once the guard becomes true, the process may be re-scheduled. All fields are strictly  
67 object-private.

68 Running a Hybrid Active Objects (HAO) model of a CPS can be pictured as follows: each  
69 object is capable of modeling a physical object, for example, a water tank. It may declare *physical*  
70 behavior via ordinary differential equations (ODEs) over “physical” fields, as well as *discrete*  
71 behavior via class and method declarations that can be used to control physical behavior. Once  
72 an HAO starts executing, the values of the physical fields evolve, governed by their ODEs, even  
73 when the controller is idle. This models the intuition that a physical system evolves independently  
74 of any observers and controllers.

75 Object orientation allows natural modeling of hybrid systems: continuous behavior is attached  
76 to an *object*, not a process. Processes realize discrete control behavior related to sensors and  
77 controllers. Specifically, the controller methods of an object may wait to execute until a certain  
78 physical state is reached (event-triggered control, for example, “tank is nearly full”). This “sensing”  
79 is modeled with getter methods of physical fields. Obviously, for validation the HABS runtime  
80 system must solve the differential equations in the physical model to determine the time point  
81 when such a waiting controller can start at the earliest; for verification, ODEs need not be solvable;  
82 they are analyzed with invariant-based techniques [67, 70]. Another communication pattern  
83 for controllers—time-triggered control—is provided by fixed sampling durations. More complex  
84 control patterns can be realized by waiting until the result of a subcomputation, i.e. a future, is  
85 ready.

86 Whenever a control process is activated, it can modify the physical state through actuators  
87 (for example, close a valve). In consequence, there are no timed race conditions, but the physical  
88 state might be changed by any process at the time it is scheduled. Actuation is modeled with  
89 setter methods of physical fields. Execution of control methods is assumed to take no physical  
90 time, unless explicitly modeled to do so.

91 Generally, a CPS can be modeled by several HAOs that communicate with each other via  
92 asynchronous method calls, for example, modeling a central controller. Often a controller object  
93 has no associated physical behavior; vice versa, an object that models physics, may not contain  
94 any control, but only sensor and actuator methods.


95 We demonstrate HAOs using three variants of water tank models. The first model, **TankMono**,  
96 is a single water tank that keeps its water level between two thresholds. It is modeled as a single  
97 object that integrates control and physics. The second model, **TankTick**, is also a single water tank,  
98 but it is modeled with two separate objects for tank and controller. The final model, **TankMulti**,  
99 is a distributed system of  $n$  **TankMono** tanks that, in addition to the local threshold, maintain a  
100 global threshold over the sum of all local water levels.

```

1 interface ISingleTank {
2   /*@ ensures 3 <= outLevel() <= 10 @*/
3   Real outLevel();
4   /*@ ensures -1/2 <= outDrain() <= 1/2 @
   */
5   Real outDrain();
6 }
7 /*@ requires 4 <= inVal <= 9 @*/
8 class CSingleTank(Real inVal)
9   implements ISingleTank {
10  /*@ invariant
11     3 <= level <= 10
12     & -1/2 <= drain <= 1/2
13     & (drain < 0 -> level > 3)
14     & (drain > 0 -> level < 10) @*/
15  physical {
16    Real level = inVal : level' = drain;
17    Real drain = -1/2 : drain' = 0;
18  }
19  Unit run() { this!ctrl(); }
20  Unit ctrl() {
21    await diff (level <= 3 & drain <= 0) | (level >= 10 & drain >= 0);
22    if (level <= 3) drain = 1/2;
23    else drain = -1/2;
24    this.ctrl();
25  }
26  Real outDrain() { return this.drain; }
27  Real outLevel() { return this.level; }
28 }

```

■ **Figure 2** TankMono: A water tank as a single HAO.



■ **Figure 3** Simulation Output of TankMono with  $\text{inVal} = 5$ .

## 101 2.1 Base System: TankMono

102 Fig. 2 shows an HAO model of a water tank whose **physical** section makes it either fill with  $\frac{1}{2}l/sec$   
 103 or drain at the same rate, according to the initial values and governing ODEs of the **level** and  
 104 **drain** fields. Method **ctrl()** realizes a control loop that switches the **drain** field between those  
 105 states so that the water level stays between  $3l$  and  $10l$ . The controller **ctrl** waits until the water  
 106 level reaches the upper or lower limit, i.e. until the condition in Fig. 2, Line 21 holds. Depending  
 107 on the case, it changes the state and calls itself recursively.

108 The JML style [20] comments in Fig. 2 contain an assumption on the initial state of **inVal**  
 109 and a conjectured safety invariant and conjectured output guarantees that, in this case, can be  
 110 proven: if the initial level is between  $4l$  and  $9l$ , then it always stays between  $3l$  and  $10l$ . Note  
 111 that Lines 13–14 express a safety invariant that must be *shown* to be true, rather than control  
 112 conditions. Intuitively, Line 13 expresses the property that the tank won't drain below a threshold  
 113 ( $\text{level} > 3$ ) even if water is leaking from it ( $\text{drain} < 0$ ). Similarly, Line 14 expresses that the tank  
 114 won't overflow ( $\text{level} < 10$ ) even if water is pumped into the tank ( $\text{drain} > 0$ ). Prior to formal  
 115 verification of this property one typically runs tests to see whether the model behaves as intended.  
 116 Our implementation allows to simulate and visualize an HAO model. The graph in Fig. 3 shows  
 117 the behavior of a **CSingleTank** object instantiated with  $\text{inVal} = 5$ . In Sect. 5 we show how the  
 118 class is translated into  $d\mathcal{L}$  and how to prove the safety invariant in KeYmaera X for *any* object  
 119 created with a parameter that satisfies the precondition. The only methods exposed to clients in  
 120 the interface are **outDrain()** and **outLevel()**.

## 121 2.2 Discrete Controller: TankTick

122 The `ctrl()` method in **TankMono** corresponds to a perfect sensor/controller that physically  
 123 reacts to the water level and drain. **TankTick** splits controller and sensor into two objects and  
 124 uses a clock to read the water level at certain intervals. This corresponds to a closed-loop control  
 125 system with a discrete-time controller that samples the plant behavior.

```

1 interface Tank {
2   /* requires  $-1/2 \leq \text{newD} \leq 1/2$ ; */
3   Unit inDrain(Real newD);
4   /* ensures  $3 \leq \text{outLevel}() \leq 10$ ; */
5   Real outLevel();
6 }
7 class CTank(Real inVal) implements Tank {
8   physical {
9     Real level = inVal : level' = drain;
10    Real drain = -1/2 : drain' = 0;
11  }
12  Unit run() { }
13  /* requires  $\text{newD} > 0 \rightarrow \text{level} \leq 9.5$  */
14  /* requires  $\text{newD} < 0 \rightarrow \text{level} \geq 3.5$  */
15  /* timed_requires inDrain < 1 */
16  Unit inDrain(Real newD) { drain = newD; }
17  Real outLevel() { return level; }
18 }


19 /* requires  $0 < \text{tick} < 1 \ \& \ \text{inVal} > 3.5$  */
20 class FlowCtrl(Tank t, Real tick, Real inVal) {
21   /* invariant  $(\text{drain} > 0 \rightarrow \text{level} \leq 9.5)$ 
22     &  $(\text{drain} < 0 \rightarrow \text{level} \geq 3.5)$  */
23   Real drain = -1/2;
24   Real level = inVal;
25
26   Unit run() { this!ctrlFlow(); }
27
28   Unit ctrlFlow() {
29     await duration(tick,tick);
30     level = t.outLevel();
31     if (level <= 3.5) drain = 1/2;
32     if (level >= 9.5) drain = -1/2;
33     t!inDrain(drain);
34     this.ctrlFlow();
35   }
36 }

```

■ **Figure 4 TankTick:** A water tank modeled as two HAOs. Invariant and precondition of `CTank` are as in Fig. 2.

126 Fig. 4 shows a water tank realized by a controller `FlowCtrl` and a `Tank` implementation `CTank`.  
 127 The tank has an in-port (setter) method `inDrain()` and an out-port (getter) method `outLevel()`.  
 128 It has no active *discrete* behavior on its own (the `run` method is empty), but its state changes  
 129 nonetheless due to the *continuous* **physical** block. The `FlowCtrl` controller's fields `drain`, `level` are  
 130 its *local copies* of the state of the tank: `CTank.drain`, `CTank.level` are different fields from `FlowCtrl`  
 131 `.drain`, `FlowCtrl.level`, respectively, residing in different objects. The `ctrlFlow()` method first  
 132 updates `level`, decides on the state of `drain`, then pushes the (possibly changed) state of `drain` to  
 133 the tank. No time passes in the controller, which ensures that the copied fields are synchronized  
 134 at the end of the round. As the `Tank`'s fields are not directly accessible by the `FlowCtrl` instance,  
 135 it is not possible to wait on the `Tank`'s `level` with an **await diff** statement. Instead, the controller  
 136 uses **await duration** to run every `tick` seconds: `tick` is the sampling time of the controller.

137 The `Tank` interface specification declares an  
 138 input requirement and a guarantee on returned  
 139 values. The input requirement of the `inDrain`  
 140 `()` specification is a constraint on the input  
 141 parameter `newD`; specifically, it means that the  
 142 tank can only be instructed to fill if there is  
 143 sufficient capacity left (similar for draining).  
 144 The initial requirement is sufficient to establish  
 145 the controller's invariant, which in turn ensures  
 146 that the tank's requirements are met. The  
 147 `timed_requires` clause stipulates that `inDrain`  
 148 `()` is called at least once per second, which  
 149 suffices for the output guarantee. Fig. 5 shows  
 150 example output. We stress that all calls to `Tank` methods are *asynchronous*.

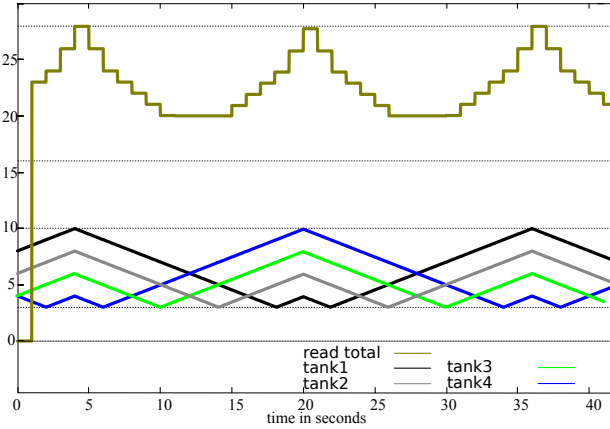


■ **Figure 5** Simulation Output of **TankTick** with `inVal = 5` for 30s.

```

1 class CControl(List<ISingleTank> tanks,
2               Real totalLower,
3               Real totalHigher,
4               Real tick)
5 implements IControl {
6   Unit run() {
7     await duration(tick, tick);
8     Real total = 0;
9     List<ISingleTank> lower = list[];
10    List<ISingleTank> higher = list[];
11    foreach ( next in tanks ) {
12      Real val = next.outLevel();
13      Real dir = next.outDrain();
14      if (dir < 0 && val > 3)
15        lower = Cons(next, lower);
16      if (dir > 0 && val < 10)
17        higher = Cons(next, higher);
18      total = total + val;
19    }
20    if (total <= totalLower+1)
21      foreach ( lnext in lower )
22        lnext!inDrain(1/2);
23    if (total >= totalHigher-1)
24      foreach ( hnext in higher )
25        hnext!inDrain(-1/2);
26    this.run();
27  }
28 }

```



■ **Figure 6 TankMulti:** A controller for  $n$  **TankMono** instances and an example simulation output. Interface omitted.

### 151 2.3 Distributed Tank Control: TankMulti

152 Consider a system where  $n$  water tanks are monitored by a central controller that aims to keep  
 153 the sum of all water levels between some thresholds. The code in Fig. 6 shows a controller that  
 154 monitors a list of **ISingleTank** (Fig. 2) instances. Each `tick` seconds the central controller iterates  
 155 over the list of tanks and if their combined level is almost at the upper threshold, the controller  
 156 drains all water tanks with rising levels (analogously for the lower threshold). Single water tanks  
 157 still ensure that their local thresholds are observed. To allow the **CControl** instance to manipulate  
 158 the **ISingleTank** instances, we add the following method to **CSingleTank** (and an analogous method  
 159 to the interface):

```

160 1 /* requires newD > 0 -> level < 10 */
161 2 /* requires newD < 0 -> level > 3 */
162 3 /* requires -1/2 <= newD <= 1/2 */
163 4 Unit inDrain(Real newD) { this.drain = newD; }

```

161 Contrary to the contract in **TankTick**, we do not need to specify how frequently the method  
 162 is called, because this information is available in the guard of the `ctrl` method of the instances.  
 163 The recursive call at the end of `ctrl` ensures that there is always one process executing `ctrl` for  
 164 each instance of **FlowCtrl**.

165 The graph in Fig.6 shows the simulation output for four water tanks with different initial  
 166 values. The upper thresholds are managed by the distributed controller and the water tanks  
 167 cooperatively: Only tanks 1 and 4 reach their local upper thresholds, the others are drained by the  
 168 distributed controller to maintain the global threshold. The lower local thresholds are managed  
 169 locally, the lower global threshold is never reached.

## 2.4 Futures

Future-based communication allows to decouple the call of a method from retrieving its result. For example, consider the code in Fig. 7. Class `Node` can perform some complex and time consuming computations on behalf of class `Client`. To enable load balancing the client has only a reference to an interface `Server`, which relays its request. The `Server` performs basic load balancing by a round-robin scheduling on a list of nodes. It then returns to the issuing client the future of the relayed request *without having to wait* for the computation to finish (Line 17). The client can then retrieve the future (Line 7) to synchronize on it without blocking the interface server (Line 8).

```

1 class Node {
2   Real compute_internal(Real r1, Real r2, Real r3){ ... }
3 }
4 class Client(Server s){
5   Unit run(){
6     Fut<Fut<Real>> ffr = s!compute(1,2);
7     Fut<Real> fr = ffr.get;
8     Real r = fr.get;
9     ...
10  }
11 }
12 class Server(Queue<Node> internal, Real param){
13   Fut<Real> compute(Real r1, Real r2){
14     Node n = internal.pop();
15     Fut<Real> fr = n!compute_internal(r1,r2,param);
16     internal.push(n);
17     return fr;
18   }
19 }

```

■ **Figure 7** An example for load balancing using futures. Interfaces omitted.

## 3 Hybrid Active Objects

An informal description of the intended semantics of Hybrid Abstract Objects in the Hybrid Abstract Behavioral Specification (HABS) language was provided in Section 2. The present section gives a formal account of its syntax and semantics. HABS is an extension of the Active Object language ABS [48]. ABS itself extends standard OO concepts as follows:

**Encapsulation.** All fields are strictly object-private.

**Cooperative Scheduling.** Active Objects cannot be preempted: a process running in an object may not be interrupted by other processes, unless the active process suspends itself or terminates.

**Asynchronous Calls, Futures.** All method calls to other objects are asynchronous. Every call not only generates a process on the callee side, but a future that points to that process. A process may pass around a future or synchronize with it to read the return value of the associated process once it has terminated.

As a *Timed* Active Object language, HABS also features:

**Simulation Time.** HABS allows to manipulate *simulation time* by explicitly advancing (and reading) an internal clock with specific statements. Simulation time is independent of the wall time.



### 193 3.1 Syntax

194 The syntax of HABS is given by the grammar in Fig. 8 and explained in the following section.  
 195 With  $e$  we denote standard expressions over fields  $f$ , variables  $v$  and operators  $|$ ,  $\&$ ,  $>=$ ,  $<=$ ,  $+$ ,  $-$ ,  $*$ ,  
 196  $/$ . Types  $T$  are all interface names, type-generic futures **Fut** $\langle T \rangle$ , lists **List** $\langle T \rangle$ , **Real**, **Int**, **Unit** and  
 197 **Bool**. We also assume the usual functions for lists, etc.

Prgm ::= $\overline{ID} \overline{CD} \text{Main}$	ID ::= <b>interface</b> $I$ [ <b>extends</b> $\overline{I}$ ]? $\{\overline{MS}\}$	Programs, Interfaces
Main ::= $\{s?\}$		Main
CD ::= <b>class</b> $C$ [ <b>implements</b> $\overline{I}$ ]? $[(\overline{T} \overline{f})]?\{\text{Phys? } \overline{FD} \overline{\text{Met}} \text{Run?}\}$		Classes
Run ::= <b>Unit</b> $\text{run}() \{s\}$	FD ::= $T \ f = e$	Run Method and Fields
Phys ::= <b>physical</b> $\{\overline{DED}\}$	DED ::= <b>Real</b> $f = e : f' = e$	Physical Block
MS ::= $T \ m(\overline{T} \ v)$	Met ::= $MS \ \{s; \text{return } e;\}$	Signatures, Methods
$s ::= \text{while } (e) \{s\} \mid \text{if } (e) \{s\} [\text{else } \{s\}]? \mid s; s$ $\mid \text{await } g \mid [T? \ e]? = \text{rhs}$		Statements
$g ::= \text{duration}(e, e) \mid \text{diff } e \mid e?$		Guards
$\text{rhs} ::= e \mid \text{new } C(\overline{e}) \mid e.\text{get} \mid e!m(\overline{e})$		RHS Expressions

■ **Figure 8** HABS grammar.  $T$  ranges over types,  $I$  over interfaces and  $C$  over classes. Differential expression  $de$  are normal expressions extended with a derivation operator  $e'$ .

198 A program contains a main method **Main**, interfaces  $\overline{ID}$  and classes  $\overline{CD}$ . Interfaces are standard,  
 199 the main method contains a list of object creations. Classes can have parameters  $\overline{Tf}$ , these are  
 200 fields being initialized during object creation. Classes have fields  $\overline{FD}$ , methods  $\overline{\text{Met}}$ , an optional  
 201 run method **Run** to start a process, and an optional physical block **Phys** that declares physical  
 202 fields. A declaration of a physical field is a field declaration followed by a differential equation.  
 203 A differential equation is an equation between two differential expressions, which are standard  
 204 expressions extended with a derivation operator  $e'$  for  $\frac{de}{dt}$ . HABS supports explicit autonomous  
 205 differential equations. The differential expressions and the field initialization form an initialized  
 206 ordinary differential equation, e.g., **Real**  $f = 0 : f' = 5-f$ . Note that  $f = 0$  specifies the initial  
 207 value of  $f$ , whereas the differential equation  $f' = 5-f$  is phrased in terms of the time-varying value  
 208 of  $f$ , so models logarithmic growth towards  $f = 5$ .

209 Methods and statements are mostly standard, we focus on HAO-specific constructs. Methods  
 210 are called asynchronously with  $e!m(\overline{e})$ , i.e., after the call, the caller continues execution without  
 211 waiting for the callee to finish. Instead, the caller generates a *future*. A future identifies the call  
 212 and can be passed around by the caller. A process interacts in two ways with a future: either by  
 213 awaiting its result with **await**  $e?$  on the guard  $e?$ , or by reading its value with  $e.\text{get}$ . Statements  
 214  $e.\text{get}$  block the reading *object*—no other process may run on it. In contrast, statements **await**  $g$   
 215 release the process control over the object while waiting for the guard  $g$  to hold. The guard is  
 216 either a future guard  $e?$ , a differential guard **diff**  $e$ , or a timed guard **duration**( $e1, e2$ ). The future  
 217 guard  $e?$  awaits the result of future  $e$ , the differential guard **diff**  $e$  suspends the process until the  
 218 expression  $e$  evaluates to true, and the timed guard **duration**( $e1, e2$ ) suspends the process for at  
 219 least  $e1$  time units<sup>1</sup>. The notation  $T \ v = o.m()$  is short for **Fut** $\langle T \rangle \ f = o!m()$ ;  $T \ v = f.\text{get}$ ; (a  
 220 call followed by a synchronization).

<sup>1</sup> The parameter  $e2$  is used by certain scheduling policies [16], and is not relevant for HABS.



## 3.2 Semantics of HABS

HABS extends the structural operational semantics (SOS) for Timed ABS [16] in three aspects: (i) it includes physical behavior in the object state; (ii) determines whether a differential guard holds and, if not, when it will at the earliest; (iii) updates the state whenever time passes. This affects only expression evaluation and auxiliary functions. *No new SOS rule is needed.* In the following we extend the core of the ABS SOS semantics [16] to hybrid systems.

### 3.2.1 States

The state of an object has three parts: (i) a store  $\rho$  that maps (physical and non-physical) fields to values, and the variables of the active process<sup>2</sup> to values; (ii)  $ODE$ , the differential equations from its physical block; (iii)  $F$ , the set of current solutions of  $ODE$ <sup>3</sup>. A solution  $f$  is a function from time to a store which only contains the physical fields. The set  $F$  may change, because the ODEs are solved as an initial-value problem with the current state of the physical fields as the initial values. For each  $f \in F$  and each physical field  $\mathbf{f}$  the following holds:  $f(0)(\mathbf{f}) = \rho(\mathbf{f})$ , i.e., the initial value  $f(0)(\mathbf{f})$  of physical field  $\mathbf{f}$  is the current value  $\rho(\mathbf{f})$  in the store  $\rho$ . We denote the solutions of  $ODE$  with initial values from  $\rho$  by  $\text{sol}(ODE, \rho)$ . We define runtime configurations formally:

$$\begin{aligned} tcn &::= \text{clock}(\mathbf{e}) \text{ } cn & cn &::= cn \text{ } cn \mid fut \mid msg \mid ob \\ ob &::= (o, \rho, \underline{ODE}, F, \underline{prc}, \underline{\overline{prc}}) & msg &::= \text{msg}(o, \bar{\mathbf{e}}, f) \\ prc &::= (\tau, f, \mathbf{rs}) \mid \perp & \mathbf{rs} &::= \mathbf{s} \mid \text{suspend}; \mathbf{s} & fut &::= \text{fut}(f, \mathbf{e}) \end{aligned}$$

■ **Figure 9** Runtime Syntax of HABS.

► **Definition 1** (Runtime Configuration [16]). The runtime syntax of HABS is summarized in Fig. 9:  $f$  ranges over future identities,  $o$  over object identities,  $\rho, \tau$  over stores, i.e., assignments from fields or variables to values. A timed configuration has a clock  $\text{clock}$  with the current time, as an expression of **Real** type and an object configuration  $cn$ . An object configuration  $cn$  consists of messages  $msg$ , futures  $fut$ , objects  $ob$ , and can be composed  $cn \text{ } cn$  (as usual, composition is commutative and associative). A message  $\text{msg}(o, \bar{\mathbf{e}}, f)$  records callee  $o$ , passed parameters  $\bar{\mathbf{e}}$  and the generated future  $f$ . A future configuration  $\text{fut}(f, \mathbf{e})$  connects the future  $f$  with its return value  $\mathbf{e}$ . An object  $(o, \rho, F, \underline{ODE}, \underline{prc}, \underline{\overline{prc}})$  has an identifier  $o$ , an object store  $\rho$ , the current solutions  $F$ , an active process  $prc$  and a queue of inactive processes.  $ODE$  is taken from the class declaration. A process is either terminated  $\perp$  or has the form  $(\tau, f, \mathbf{rs})$ : the process store  $\tau$  with current state of the local variables, its future  $f$ , and the statement  $\mathbf{rs}$  left to execute. The runtime syntax also allows the **suspend** statement, which is used to deschedule a process. Dotted underlined elements are an extension of HABS relative to ABS (also in Fig. 10 below).

Given a process store  $\tau$  and an object store  $\rho$  we use  $\sigma = \rho \circ \tau$  to denote the state of both fields and local variables. We first define the evaluation of expressions and guards.

<sup>2</sup> Recall that the active process executes the ABS methods, it does not relate to physical behavior.

<sup>3</sup> The solutions computed relative to the initial values (state) at the last suspension.

### 252 3.2.2 Evaluation of Expressions

253 Expressions  $e$  are evaluated with a function  $\llbracket e \rrbracket_\sigma^{F,t}$  over a store  $\sigma$  and a set of solutions  $F$  at  $t$  time  
254 units in the future. The semantics of expressions containing physical fields is as follows.

255 **► Definition 2 (Semantics of Expressions).** Let  $F$  be the set of solutions. Given a store  $\sigma$ , we can  
256 check whether  $F$  is a model of an expression  $e$  after  $t$  time units. Let  $\mathbf{f}_p$  be a physical field and  $\mathbf{f}_d$   
257 a non-physical field of  $o$ . The semantics of fields  $\mathbf{f}_p$ ,  $\mathbf{f}_d$ , unary operators  $\sim \in \{!, -\}$  and binary  
258 operators  $\oplus \in \{!, \&, >=, <=, +, -, *, /\}$  is defined as follows:

$$259 \quad \llbracket \mathbf{f}_d \rrbracket_\sigma^{F,t} = \sigma(\mathbf{f}_d) \quad \llbracket \mathbf{f}_p \rrbracket_\sigma^{F,t} = \begin{cases} v & \text{if } \forall f \in F. v = f(t)(\mathbf{f}_p) \\ \infty & \text{otherwise} \end{cases}$$

$$260 \quad \llbracket \sim e \rrbracket_\sigma^{F,t} = \sim \llbracket e \rrbracket_\sigma^{F,t} \quad \llbracket e_1 \oplus e_2 \rrbracket_\sigma^{F,t} = \llbracket e_1 \rrbracket_\sigma^{F,t} \oplus \llbracket e_2 \rrbracket_\sigma^{F,t}$$

262 Outside differential guards, only the evaluation in the current state  $\llbracket e \rrbracket_\sigma^{F,0}$  is needed, which is  
263  $\rho(\mathbf{f}_p)$  from  $f(0)(\mathbf{f}_p)$  and this expression is never  $\infty$ . We identify  $\llbracket e \rrbracket_\sigma^F$  and  $\llbracket e \rrbracket_\sigma$  with  $\llbracket e \rrbracket_\sigma^{F,0}$ .

### 264 3.2.3 Evaluation of Guards

265 The semantics of an **await**  $g$  statement is to suspend until the guard holds, i.e. until  $\llbracket g \rrbracket_\sigma^F$  evaluates  
266 to true. For example, a duration guard **duration**( $e_1, e_2$ ) evaluates to true if  $\llbracket e_1 \rrbracket_\sigma^F \leq 0$ . Defining  
267 the semantics of guards requires two operations: An extension of the *evaluation function* that  
268 returns true if the guard holds and the *maximal time elapse*  $mte_\sigma^F$  returning the time  $t$  that may  
269 elapse before the guard evaluates to true, or  $\infty$  if it never does.

270 First we define  $mte(e)$ : the *maximal* time that may elapse without missing an event is the  
271 *minimal* time needed by the system to evolve into a state where the guard is guaranteed to hold.  
272 This yields also the semantics of the guard itself.

273 **► Definition 3 (Semantics of Differential Guards).** Let  $F$  be the set of solutions of object  $o$  in  
274 state  $\sigma$ . Then we define:

$$275 \quad mte_\sigma^F(\mathbf{diff} \ e) = \underset{t \geq 0}{\mathbf{argmin}} \ (\llbracket e \rrbracket_\sigma^{F,t} = \mathbf{true})$$

277 **diff**  $e$  is evaluated to true if no time advance is needed:

$$278 \quad \llbracket \mathbf{diff} \ e \rrbracket_\sigma^{F,0} = \mathbf{true} \iff mte_\sigma^F(\mathbf{diff} \ e) = 0$$

280 If  $e$  contains no continuous variable then the differential guard semantics and the evaluation of  
281 expressions in Def. 2 coincides with condition synchronization and expression evaluation in the  
282 standard ABS semantics [48].

### 283 3.2.4 Transition System

284 Fig. 10 gives the most important rules for the semantics of a single object, the omitted rules  
285 are given in [16]. Rules (1)–(3) define the semantics of process suspension. An **await** statement  
286 suspends the current process and gives other processes in the queue  $q$  a chance to run, even if  
287 its guard is evaluated to true. Suspension is modeled in rule (1) simply by introducing a **suspend**  
288 statement in front of the **await**.<sup>4</sup> Rule (2) realizes a **suspend** statement by moving the current

<sup>4</sup> We follow the original ABS semantics, where suspension is handled with a separate **suspend** statement for reasons of uniformity—in principle, rules (1)+(2) could be combined.

$$\begin{aligned}
(1) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{await} \ g; \mathbf{s}), q) \rightarrow (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{suspend}; \mathbf{await} \ g; \mathbf{s}), q) \\
(2) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{suspend}; \mathbf{s}), q) \rightarrow (o, \rho, \underline{ODE}, \mathbf{sol}(ODE, \rho), \perp, q \circ (\tau, f, \mathbf{s})) \\
(3) \quad & (o, \rho, \underline{ODE}, F, \perp, q \circ (\tau, f, \mathbf{await} \ g; \mathbf{s})) \rightarrow (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{s}), q) \\
(4) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{v} = \mathbf{e}; \mathbf{s}), q) \rightarrow (o, \rho, \underline{ODE}, F, (\tau[\mathbf{v} \mapsto \llbracket \mathbf{e} \rrbracket_{\rho \circ \tau}], f, \mathbf{s}), q) \\
& \quad \text{if } \llbracket \mathbf{g} \rrbracket_{\rho \circ \tau} = \text{true} \\
& \quad \text{if } \mathbf{e} \text{ contains no call or } \mathbf{get} \\
(5) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{return} \ \mathbf{e};), q) \rightarrow (o, \rho, \underline{ODE}, \mathbf{sol}(ODE, \rho), \perp, q) \ \mathbf{fut}(f, \llbracket \mathbf{e} \rrbracket_{\rho \circ \tau}) \\
(6) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{v} = \mathbf{e}_1. \mathbf{get}; \mathbf{s}), q) \ \mathbf{fut}(f, \mathbf{e}_2) \rightarrow (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{v} = \mathbf{e}_2; \mathbf{s}), q) \\
& \quad \text{if } \llbracket \mathbf{e}_1 \rrbracket_{\rho \circ \tau} = f \\
(7) \quad & (o, \rho, \underline{ODE}, F, (\tau, f, \mathbf{v} = \mathbf{e}! \mathbf{m}(\mathbf{e}_1, \dots, \mathbf{e}_n); \mathbf{s}), q) \rightarrow \\
& \quad (o, \rho, \underline{ODE}, F, (\tau[\mathbf{v} \mapsto \tilde{f}], f, \mathbf{s}), q) \ \mathbf{msg}(\llbracket \mathbf{e} \rrbracket_{\rho \circ \tau}, (\llbracket \mathbf{e}_1 \rrbracket_{\rho \circ \tau}, \dots, \llbracket \mathbf{e}_n \rrbracket_{\rho \circ \tau}), \tilde{f}) \\
& \quad \text{where } \tilde{f} \text{ is fresh}
\end{aligned}$$

■ **Figure 10** Selected Rules for HABS objects.

289 process to the object's queue. As explained in Sect. 3.2.3, upon reactivation of a suspended  
290 process we must ensure its guard to be true, relative to the solution of *ODE* with *initial values at*  
291 *suspension time*. Therefore, rule (2) also recomputes the solutions *F*. Rule (3) can then re-activate  
292 a process beginning with an **await** statement, simply by checking whether its guard evaluates to  
293 true at current time (advancing time in timed configuration is explained below). An analogous  
294 rule (not shown in Fig. 10) activates a process with any other non-**await** statement. Rule (4)  
295 evaluates an assignment to a local variable. The rule for fields is analogous. Rule (5) realizes a  
296 termination (with solutions of the ODEs) and (6) a future read. Finally, (7) is a method call, the  
297 rule for transforming a message into a process is straightforward.

298 For configurations, there are two rules, shown in Fig. 11. Rule (i) realizes a step of some object  
299 without advancing time, Only if (i) is not applicable, i.e. all ABS processes are blocked, rule (ii)  
300 can be applied. It computes the global maximal time elapse *mte* and advances the time in the  
301 clock and all objects. In particular, it decreases syntactically the timed guards.

$$\begin{aligned}
(i) \quad & \mathbf{clock}(t) \ cn \ cn_1 \rightarrow \mathbf{clock}(t) \ cn_2 \ cn_1 \quad \text{with } cn \rightarrow cn_2 \\
(ii) \quad & \mathbf{clock}(t) \ cn \rightarrow \mathbf{clock}(t + \tilde{t}) \ \mathbf{adv}(cn, \tilde{t}) \quad \text{if (i) is not applicable and } \mathbf{mte}(cn) = \tilde{t} \neq \infty
\end{aligned}$$

■ **Figure 11** Timed Semantics of HABS configurations.

302 Fig. 12 shows the auxiliary functions and includes the full definition of *mte*. Note that *mte*  
303 is not applied to the currently active process, because, when (1) is not applicable, it is currently  
304 blocking and, thus, cannot advance time. *The characteristic feature of hybrid objects is that their*  
305 *physical state changes when time advances, even when no process is active*. This is expressed in  
306 the semantics by a function *adv*( $\sigma, t$ ) which takes a state  $\sigma$ , a duration  $t$ , and advances  $\sigma$  by  $t$   
307 time units. For non-hybrid Active Objects *adv*( $\sigma, t$ ) =  $\sigma$ . There, the function is needed only to  
308 modify the process pool of an object for scheduling, not its state, and is used exactly as in [16].

309 The *adv* auxiliary function handles uniqueness w.r.t. the solutions of the ODE at the points in

$$\begin{aligned}
mte(cn_1 \text{ } cn_2) &= \mathbf{min}(mte(cn_1), mte(cn_2)) & mte(msg) &= mte(fut) = \infty \\
mte(o, \rho, ODE, F, prc, q) &= \llbracket \mathbf{min}(mte(q), \infty) \rrbracket_\rho & mte(\tau, f, \mathbf{await} \text{ } g; s) &= \llbracket mte(g) \rrbracket_\tau \\
mte(\tau, f, s) &= \infty \text{ if } s \neq \mathbf{await} \text{ } g; \tilde{s} & mte(\mathbf{duration}(e_1, e_2)) &= e_1 \\
mte_\sigma^F(\mathbf{diff} \text{ } e) &= \mathbf{argmin}_{t \geq 0} (\llbracket e \rrbracket_\sigma^{F,t} = \text{true}) & mte(e?) &= \infty \\
adv(cn_1 \text{ } cn_2, t) &= adv(cn_1, t) \text{ } adv(cn_2, t) \\
adv(msg, F, t) &= msg & adv(fut, F, t) &= fut \\
adv((o, \rho, ODE, F, prc, q), F, t) &= (o, adv(\rho, t), ODE, F, adv(prc, F, t), adv(q, F, t)) \\
adv(\perp, F, t) &= \perp \\
adv((\tau, f, s), F, t) &= (\tau, f, s) \text{ if } s \neq \mathbf{await} \text{ } \mathbf{duration}(e_1, e_2); \tilde{s} \\
adv((\tau, f, \mathbf{await} \text{ } \mathbf{duration}(e_1, e_2); s), F, t) &= (\tau, f, \mathbf{await} \text{ } \mathbf{duration}(e_1+t, e_2+t); s) \\
adv(\sigma, t)(\mathbf{f}) &= \begin{cases} \sigma(\mathbf{f}) & \text{if } \mathbf{f} \text{ is not physical} \\ v & \text{if } \forall f \in F. v = f(t)(\mathbf{f}) \end{cases}
\end{aligned}$$

■ **Figure 12** Auxiliary functions. Lifting to lists is not shown.

310 time where the solutions are accessed: Note that the solutions are handled as a set  $F$ : at time  $t$   
311 function  $adv$  checks that all solutions coincide *at this point in time*. If this is not the case, or if no  
312 solution can be found by the implementation, a runtime error is thrown. Also, all solutions are  
313 computed without restrictions on the time domain (e.g., for how long they exists) because it is  
314 not known for how long the dynamics are followed at this point. Alternatively, one could either  
315 impose restrictions on the ODE to enforce uniqueness or non-deterministically choose one of the  
316 solutions.

317 We can now define *traces* of programs and objects.

318 ► **Definition 4** (Traces). Given a program  $\text{Prgm}$ , we denote with  $\text{clock}(0) \text{ } cn_0$  the initial state  
319 configuration [16]. A run of  $\text{Prgm}$  is a (possibly infinite) reduction sequence


$$320 \quad \text{clock}(0) \text{ } cn_0 \rightarrow \text{clock}(t_1) \text{ } cn_1 \rightarrow \dots$$

321 The trace  $\theta_o$  of an object  $o$  in a run is an assignment from the dense time domain  $\mathbb{R}^+$  to states.  
322 We say that  $\text{clock}(t_i) \text{ } cn_i$  is the final configuration at  $t_i$  in a run, if any other timed configuration  
323  $\text{clock}(t_i) \text{ } \tilde{cn}_i$  is before it. Fig. 13 gives a formal definition.

$$\theta_o(x) = \begin{cases} \text{undefined} & \text{if } o \text{ is not created yet} \\ \rho & \text{if } \text{clock}(x) \text{ } cn \text{ is the final configuration at } x \\ & \text{and } \rho \text{ is the state of } o \text{ in } cn \\ adv(\rho, F, x - y) & \text{if there is no configuration at } \text{clock}(x) \\ & \text{and the last configuration was at } \text{clock}(y) \\ & \text{with state } \rho \text{ and solutions } F \end{cases}$$

■ **Figure 13** Extraction of a trace  $\theta_o$  for an object  $o$  from a given run.

324 For any point in time  $x$ , the state of  $o$  is taken from the run, if a reduction step was made at  
325  $x$  and  $o$  was already created. The third case in the definition is illustrated in Fig. 14: At time



■ **Figure 14** Illustration of the state at time  $x$  and two discrete states with  $\text{clock}(y)$  and  $\text{clock}(z)$ .

326 points  $y$  and  $z$ , discrete steps are done, but none at  $x$ . The state  $\theta_o(x)$  is extrapolated from the  
 327 state  $\theta_o(y)$  by following solutions from the last step at point  $y$ , if  $o$  is created.

### 328 3.3 The Component Fragment

329 We define a sublanguage of HABS called *Component HABS* (CHABS) to model component-style  
 330 architectures with in- and out-ports, as well as dedicated controllers with a read-evaluate-write  
 331 cycle. Syntactically, a class is a *component* if it can be derived from the syntax in Fig. 8 with the  
 332 rule for Met replaced by the following:

```

333 Met ::= MS [OPort | IPort | Ctrl]
334 OPort ::= {return this.f;}    IPort ::= {this.f = v; return Unit;}
335 Ctrl ::= {sa; si; sc; so; this.m();}
336 sa ::= await duration(e,e) | await diff e
337 si ::= this.f = e.m() | si;si
338 sc ::= while (e) {sc} | if (e) {sc} [else {sc}]? | sc;sc | T? e = e | e!m(e)
339 so ::= e!m(this.f) | so;so
340

```

341 Additionally, we demand that the only numerical data types used are **Int**, **Real**. Out-ports return  
 342 the value of a field and in-ports copy a method parameter into a field. A controller method Ctrl  
 343 has a timed or differential guard sa, followed by reads si from the out-port methods of other  
 344 objects (recall that **this.f = e.m()** is a shortcut for an asynchronous call followed by a read, not  
 345 a synchronous call), computations sc, and writes so to the in-ports of other objects. In the  
 346 component fragment, we realize a component-based controller with a read-compute-write loop  
 347 by restricting the run method of Fig. 8 to start a controller with an asynchronous call to an  
 348 object's own controller method Ctrl and each controller ends with a recursive call to itself. The  
 349 **TankMono** and **TankTick** models are CHABS models, the central controller in **TankMulti** is  
 350 not. A controller method with a differential guard is an event-triggered controller, a controller  
 351 with a timed guard a time-triggered controller.

352 We model instantaneous controllers in CHABS: once controller is scheduled (i.e., after its guards  
 353 evaluates to true) no time can pass because all calls in Ctrl are to port methods that cannot block  
 354 the caller and neither suspensions nor future reads are allowed.

### 355 3.4 Simulation

356 The implementation of HABS extends the ABS compiler [81] to compute solutions for differential  
 357 guards, time elapse, and state advance. To compile differential guards correctly, it needs to  
 358 compute  $\text{mte}_\sigma^F(\text{diff } e)$  (Def. 3).

359 The ODEs of a class cannot be changed at runtime and are, therefore, represented as a string  
 360 in the class table. The simulator uses an external solver to solve initial value problems and  
 361 minimize/maximize duration between events.

362 **Solutions** To compute solutions  $F$ , the ODEs and the current state of the physical fields are passed  
 363 to Maxima [61] as an *initial value problem*. The solution is an equation system or an error. In  
 364 its default setting, the simulator neither supports non-unique solutions nor non-solvable ODEs.  
 365 The simulator, however, has the infrastructure to use solvers other than Maxima. This allows  
 366 us to handle non-linear ODEs: by prefixing the **physical** block with [1], the modeler can select  
 367 the solver `ic1` (instead of the default `desolve`), which can handle non-linear systems.

368 **Time elapse** After solving the initial value problem, Maxima is invoked with a *minimization*  
 369 *problem*: it minimizes the time  $t$  with the equation system representing  $F$  as the constraints  
 370 (this corresponds to eager mode switching in a hybrid automaton). The result is then handled  
 371 in the same way as a parameter to a timed guard by the runtime system. Once time has  
 372 passed and the suspended process is reactivated, the physical fields are updated according to  
 373  $F$ . This uses the Maxima function `fmin_coby1a`.

374 **State advance** To implement the advance function *adv*, if the state of the object changes any  
 375 physical field, the procedure used to compute time elapse is repeated for every currently  
 376 suspended differential guard to accumulate the result.

377 The output files used to visualize a program execution are of the form  $t_1, F_1, t_1, F_2, t_2, \dots, F_n, t_n$ .  
 378 Here  $t_i$  are the points in time where the object schedules a process and  $F_i$  the function describing  
 379 its physical behavior in the previous suspended state. Each time a differential guard is reactivated,  
 380 not only its state is updated, but the solution  $F_{i+1}$  and the reactivation time  $t_{i+1}$  are written to  
 381 the output. Each object has its own output file.

382 A Python script translates output files into a discrete dynamic graph in Maxima format which  
 383 in turn calls `gnuplot` that is responsible for creating the graph. The graphs in this work are slightly  
 384 beautified outputs.

## 385 4 Modeling with HABS

386 We give more examples of HABS models and discuss some design decisions in the language, as well  
 387 as modeling patterns in HABS for common phenomena in hybrid system control.

### 388 4.1 Non-Linear Dynamics

389 HABS can handle non-linear ODEs and non-linear dynamics to the extent the backends support  
 390 it. For an example, consider a resistor attached to an alternating current source that produces a  
 391 sine-formed current. This is described by the class in Fig 15.

392 We use the non-linear solver of Maxima (by annotating [1]). This solver requires the input to  
 393 satisfy certain syntax constraints, which entail the slightly awkward specification `r' = 0*t`. We  
 394 must give an explicit ODE for each non-constant variable for KeYmaera X and as HABS requires  
 395 an autonomous system, we add a clock variable `time` to express sine and cosine.

396 The example has a `run` method that illustrates validation. We check whether our simple model  
 397 is in fact a resistor and adheres to the law  $R = I/V$ : Even before visualization, we can use simple  
 398 command line output to check  $I/V$  by sampling every 1 second. The output for an instance  
 399 `Resistor(5)` is shown in Fig. 15, where `Time(n)` is the symbolic time at the point of time when  
 400 `now()` is evaluated. In the example this corresponds to seconds. As a next step, we can use the  
 401 visualization to observe longer trends in Fig. 16, again for a `Resistor(5)`.

```

class Resistor(Real init) {
  [1] physical {
    /* format expected by Maxima */
    Real t = 0: t' = 1;
    Real r = init: r' = 0*t;
    Real i = 0: i' = cos(t);
    Real v = 0: v' = r*cos(t);
  }
  Unit run() {
    await duration(1,1);
    println("step: " + toString(now()) +
      " with " + toString(v/i));
    if (timeValue(now()) < 60) this!run();
  }
}


```

```

step: Time(1) with 5
step: Time(2) with 4286450913523623 /
  ↪ 857290182704725
step: Time(3) with 1319812111494398 /
  ↪ 263962422298881
step: Time(4) with 1313376056981147 /
  ↪ 262675211396229
step: Time(5) with 295788950328081 /
  ↪ 59157790065616
step: Time(6) with 723097187038613 /
  ↪ 144619437407721
step: Time(7) with 758118670875062 /
  ↪ 151623734175013
step: Time(8) with 5
step: Time(9) with 5
...

```

■ **Figure 15** A resistor attached to an AC-circuit and its sine-formed current



■ **Figure 16** Example simulation output of a Resistor(5)

402 Finally, we can formally verify the behavior with our translation approach to KeYmaera X by  
 403 removing the `run` method and, thus, transforming it into a CHABS component.

## 404 4.2 Delays and Imprecision

405 Communication is imperfect in realistic models. We demonstrate how to model two such imper-  
 406 fections, delays and imprecision, in HABS. We use a simple platooning example, where a follower  
 407 car wants to follow a lead car at a certain distance. Follower cars are modeled in the CHABS class  
 408 `FollowerCar` in Fig. 17. For simplicity, the minimal (`minDist`) and maximal distance (`maxDist`) to  
 409 the lead car are independent of the speed and the controller sampling frequency, which means the  
 410 follower car will not provably stay in the desired distance interval. The time consuming statement  
 411 `await duration` can be used to model two kinds of delays:

- 412 1. Complex computations that take some time to finish.
- 413 2. Latency: By adding a time consuming statement as the last statement of a method before the  
 414 return, one can model delays in a network.



```

class FollowerCar
  (Real inita, Real start,
   Real tick, Real minDist,
   Real maxDist, ICar leadCar)
  implements ICar {
    Real next = start + minDist;
    physical {
      Real a = inita : a' = 0;
      Real v = 0      : v' = a;
      Real x = start : x' = v;
    }
    Unit run() {
      this!ctrlObserve();
    }
    Unit ctrlObserve() {
      await duration(tick, tick);
      next = leadCar.getPosition();
      if(next - x <= minDist) a = a/2;
      if(next - x >= maxDist) a = a*2;
      this.ctrlObserve();
    }
    Real getPosition() {
      return x;
    }
  }

```

■ **Figure 17** Simple platooning example for a follower car following safely behind a lead car

415 For example, we extend `getPosition()` in `FollowerCar` to model sensing latency as follows:

```

415 Real getPosition() {
416   Real oldVal = x;
417   await duration(1/10, 1/10);
418   return oldVal;
419 }

```

417 Like ABS, HABS has access to a (uniformly distributed) random number generator. There are  
 418 functions to generate other statistical distributions. This allows to model imprecision/uncertainty.  
 419 The following method adapts `getPosition()` to model sensor uncertainty:

```

420 Real getPosition() {
421   Real imp = (random(11) + 95)/100; // number between 0.95 and 1.05
422   return this.level * imp;
423 }

```

### 421 4.3 Variability Modeling

422 One of the main advantages of using a mature programming language as a host for hybrid behavior  
 423 is that we can use its structuring elements and concepts: HABS inherits the module system with  
 424 import/export clauses<sup>5</sup>, as well as the delta-oriented [73], feature-oriented [14] *product line* [8, 74]  
 425 (DFPL) mechanisms of ABS [25] to model variability.

426 DFPLs define not a single model, but a set of models which are variants of each other. From a  
 427 given *core* model, so-called code *deltas* define variants based on syntactic operations: removal,  
 428 modification and addition of classes, methods and fields. A variant is obtained from the core  
 429 model by applying modifications specified by the deltas to it.

430 To determine the relevant deltas, each delta has a set of features that activate its application.  
 431 A feature of a variant corresponds roughly to one implemented feature of the modified model. A  
 432 set of features is called a *product*. After selecting a product, the corresponding deltas are computed  
 433 and applied, resulting in an HABS model without variability.

<sup>5</sup> Omitted from the language syntax in Sec. 3 for brevity.

```

delta Delay;
modifies class Cars.FollowerCar {
  modifies Real getPosition() {
    Real old = original();
    await duration(1/10,1/10);
    return old;
  }
}
delta Imprecision;
modifies class Cars.FollowerCar {
  modifies Real getPosition() {
    return original()*(random(11)+95)/100;
  }
}

productline PL1;
features FDelay, FImprecision, FCruiseControl;
delta CruiseControl when FCruiseControl;
delta Delay when FDelay;
delta Imprecision after Delay when FImprecision;

delta CruiseControl;
modifies class Cars.FollowerCar {
  adds Real ccTick = this.tick*2;
  adds Unit cruise() {
    await duration(ccTick, ccTick);
    if ((v >= 5 || v <= 0) && a != 0) {
      a = 0;
    }
    this.cruise();
  }
}
modifies Unit run() {
  original();
  this!cruise();
}
}

```

■ **Figure 18** Product line based on Fig. 17 for variability in position readings and cruise control

434 We refrain from introducing the whole variability layer of ABS and refer to [25] for a detailed  
 435 and formal introduction. Instead, we use the platooning example in Fig. 17 to demonstrate  
 436 variability modeling in practice. The changes for imprecision and delay, as well as adding a cruise  
 437 control system can be modeled as a product line. This allows to select the appropriate car product  
 438 for a concrete system, as summarized in Fig. 18. The product line consists of three deltas (**Delay**,  
 439 **Imprecision** and **CruiseControl**), three features (**FDelay**, **FImprecision** and **FCruiseControl**) and  
 440 a knowledge base that defines which features select which delta (**delta D when F**) and in which  
 441 order deltas are applied if they modify the same method (**delta D after D2**).

442 The delta **Delay** modifies class `Cars.FollowerCar`<sup>6</sup> and its method `getPosition()`. The modified  
 443 method first calls the existing variant of the method via **original** and then waits before returning  
 444 the value. Delta **Imprecision** is similar. Both deltas modify the same method. There are numerous  
 445 desirable properties, and to make the product line outcome deterministic, we must fix the order in  
 446 which methods are applied that modify the same method. Here, we demand that **Imprecision** is  
 447 applied after **Delay**. Delta **CruiseControl** adds a field and method implementing a simple cruise  
 448 control system. Deltas may also remove methods and fields (not shown here). In our example we  
 449 represent each delta as a feature, and so any product that refers to a feature invokes its assigned  
 450 delta. The deltas are applied *syntactically* before type checking. As a result, a standard HABS  
 451 program is created. For example the product `{FDelay}` results in the code below.

```

class FollowerCar (...) implements ICar {
  ... // as above
  Real getPosition_core() { return x; }
  Real getPosition() { return this.getPosition_core()*(random(11) + 95)/100; }
}

```

<sup>6</sup> `Cars` is the module.

## 5 Formal Verification of HABS Models

As a prerequisite for formal verification of HABS, we briefly review *differential dynamic logic* ( $d\mathcal{L}$ ) [68, 69] as implemented in the hybrid systems theorem prover KeYmaera X [36]. We then discuss translation from HABS to  $d\mathcal{L}$ , and sketch formal verification in  $d\mathcal{L}$  with sequent proofs.

### 5.1 Background: Differential Dynamic Logic

Differential dynamic logic expresses the combined discrete and continuous dynamics of hybrid systems in a sequential imperative programming language called *hybrid programs*. Its syntax and informal semantics are in Table 1.

■ **Table 1** Hybrid programs in  $d\mathcal{L}$

Program	Informal semantics
$?\varphi$	Test whether formula $\varphi$ is true, abort if false
$x := \theta$	Assign value of term $\theta$ to variable $x$
$x := *$	Assign any (real) value to variable $x$
$\{x' = \theta \ \& \ H\}$	Evolve ODE system $x' = \theta$ for any duration $t \geq 0$ with evolution domain constraint $H$ true throughout
$\alpha; \beta$	Run $\alpha$ followed by $\beta$ on resulting state(s)
$\alpha \cup \beta$	Run either $\alpha$ or $\beta$ non-deterministically
$\alpha^*$	Repeat $\alpha$ $n$ times, for any $n \in \mathbb{N}$

Hybrid programs provide the usual discrete statements: assignment ( $x := \theta$ ), non-deterministic assignment ( $x := *$ ), test ( $?\varphi$ ), non-deterministic choice ( $\alpha \cup \beta$ ), sequential composition ( $\alpha; \beta$ ), and non-deterministic repetition ( $\alpha^*$ ). A typical modeling pattern combines non-deterministic assignment and test (e.g., “ $x := *; ?H$ ”) to choose any value subject to a  $d\mathcal{L}$  constraint  $H$ . Standard control structures are expressible, for example: (i) **if**  $H$  **then**  $\alpha$  **else**  $\beta \equiv (?H; \alpha) \cup (? \neg H; \beta)$ , (ii) **if**  $H$  **then**  $\alpha \equiv (?H; \alpha) \cup (? \neg H)$ , (iii) **while** ( $H$ )  $\alpha \equiv (?H; \alpha)^*; ? \neg H$ .

For continuous dynamics, the notation  $\{x' = \theta \ \& \ H\}$  represents an ODE system (derivative  $x'$  in time) of the form  $x'_1 = \theta_1, \dots, x'_n = \theta_n$ . Any behavior described by the ODE stays inside the evolution domain  $H$ , i.e. the ODE is followed for a non-deterministic, non-negative period of time, but stops before  $H$  becomes false. For example, a basic model of the water level  $x$  in a tank draining with flow  $-f$  is given by the ODE  $\{x' = -f \ \& \ x \geq 0\}$ , where the evolution domain constraint  $x \geq 0$  means the tank will not drain to negative water levels. With a careful modeling pattern, ODEs can be governed by  $H$  so that one can react to events, without restricting or influencing the continuous dynamics modeled in the ODE [72]: The pattern  $\{x' = \theta \ \& \ H\} \cup \{x' = \theta \ \& \ \tilde{H}\}$  permits control intervention to achieve different behavior triggered by an event  $H$ .  $\tilde{H}$  is the *weak* complement of  $H$ : they share exactly their *boundary* from which both behaviors are possible. For example,  $H \equiv x \leq 0$ ,  $\tilde{H} \equiv x \geq 0$ .

The  $d\mathcal{L}$ -formulas  $\varphi, \psi$  relevant for this paper are propositional logic operators  $\varphi \wedge \psi, \varphi \vee \psi, \varphi \rightarrow \psi, \neg \varphi$  and comparison expressions  $\theta \sim \eta$ , where  $\sim \in \{<, \leq, =, \neq, \geq, >\}$  and  $\theta, \eta$  are real-valued terms over  $\{+, -, \cdot, /\}$ . In addition, there is the  $d\mathcal{L}$  modal operator  $[\alpha]\varphi$ . The  $d\mathcal{L}$ -formula  $[\alpha]\varphi$  is true iff  $\varphi$  holds in all states reachable by program  $\alpha$ . The formal semantics of  $d\mathcal{L}$  [68, 69] is a Kripke semantics in which the states of the Kripke model are the states of the hybrid system. The semantics of a hybrid program  $\alpha$  is a relation  $\llbracket \alpha \rrbracket$  between its initial and final states. Specifically,  $\nu \models [\alpha]\varphi$  iff  $\omega \models \varphi$  for all states  $(\nu, \omega) \in \llbracket \alpha \rrbracket$ , so all runs of  $\alpha$  from  $\nu$  are safe relative to  $\varphi$ .

Proofs in  $d\mathcal{L}$  are sequent calculus proofs on the basis of  $d\mathcal{L}$  axioms. For example, validity of the  $d\mathcal{L}$  formula  $x \geq 0 \rightarrow [x := x + 1 \cup x := 2; \{x' = 3\}]x \geq 1$  over a simple program that either

487 increments the value of  $x$  or continuously evolves  $x$  with a constant slope  $x' = 3$  after setting the  
 488 initial value of the differential equation with  $x := 2$  is shown in the sequent proof below:

$$\begin{array}{c}
 \begin{array}{c}
 \text{QE} \frac{*}{x \geq 0 \vdash x + 1 \geq 1} \\
 \text{[:=]} \frac{}{x \geq 0 \vdash [x := x + 1]x \geq 1} \\
 \text{[}\cup\text{],}\wedge_R \frac{}{x \geq 0 \vdash [x := x + 1 \cup x := 2; \{x' = 3\}]x \geq 1} \\
 \rightarrow_R \frac{}{\vdash x \geq 0 \rightarrow [x := x + 1 \cup x := 2; \{x' = 3\}]x \geq 1}
 \end{array}
 \quad
 \begin{array}{c}
 \text{dI} \frac{\text{[:=],hideL} \frac{\text{[;]} \frac{\text{[;]} \frac{}{x \geq 0 \vdash [x := 2; \{x' = 3\}]x \geq 1}}{x \geq 0 \vdash [x := 2; \{x' = 3\}]x \geq 1}}{x \geq 0 \vdash [x := 2][\{x' = 3\}]x \geq 1}}{x = 2 \vdash [\{x' = 3\}]x \geq 1} \\
 \text{[;]} \frac{}{x \geq 0 \vdash [x := 2; \{x' = 3\}]x \geq 1} \\
 \text{[}\cup\text{],}\wedge_R \frac{}{x \geq 0 \vdash [x := x + 1 \cup x := 2; \{x' = 3\}]x \geq 1} \\
 \rightarrow_R \frac{}{\vdash x \geq 0 \rightarrow [x := x + 1 \cup x := 2; \{x' = 3\}]x \geq 1}
 \end{array}
 \end{array}$$

490 Sequent proofs proceed bottom-up but validity transfers top-down, i.e., from the subgoals  
 491 above the horizontal bar, the axiom or proof rule annotated to the left of the bar implies the  
 492 sequent below the horizontal bar. In each step, assumptions are listed to the left of the  $\vdash$ , and the  
 493 alternatives to prove to the right of it. The proof starts with step  $\rightarrow_R$  to make the left-hand side  
 494  $x \geq 0$  of the implication available as an assumption. Next, the non-deterministic choice step  $[\cup]$   
 495 means that both choices must ensure the postcondition  $x \geq 1$ , so with conjunction splitting  $\wedge_R$  we  
 496 get two subgoals: a left subgoal for the increment program  $x := x + 1$  and a right subgoal for the  
 497 differential equation program. On the increment program branch, we execute the assignment in  
 498 step  $[\text{:=}]$  and the result follows by real arithmetic in step QE. On the differential equation branch,  
 499 step  $[\text{;}]$  splits the sequential composition into nested box modalities, and then step  $[\text{:=}], \text{hideL}$   
 500 executes the assignment and weakens the now obsolete assumption  $x \geq 0$ . The branch closes by  
 501 differential induction dI (intuitively, the dI step expresses that  $x \geq 1$  stays true along the flow of  
 502 the differential equation, see [67]). This concludes the example proof.

## 5.2 Formal Verification of Components

503 To establish system-wide properties, hybrid active objects  
 504 must be shown to satisfy their class *invariants*, provided  
 505 that the constraints expressed in the preconditions are  
 506 met. We make this precise now. A *class specification* is  
 507 a tuple  $(\text{inv}, \text{pre}, \text{TReq}, \text{Req}, \text{Ens})$ , where  $\text{inv}$  is the class  
 508 invariant (annotated */\* invariant ... \*/*, see Fig. 19, lines  
 509 20–21), a  $d\mathcal{L}$  formula over the fields and parameters of  
 510 the class;  $\text{pre}$  is the precondition (annotated to class  
 511 declarations with */\* requires ... \*/*, see Fig. 19, line 17),  
 512 a  $d\mathcal{L}$  formula over the initial values of fields and class  
 513 parameters.  $\text{TReq}$  is the set of timed input requirements  
 514 for in-port methods (annotated with */\* timed\_requires*  
 515 *... \*/*, see Fig. 19, line 12):  $d\mathcal{L}$  formulas over a dedicated  
 516 program variable with the method's name.  $\text{Req}$  is the  
 517 set of input requirements for in-port methods (annotated  
 518 with */\* requires ... \*/*, see Fig. 19, lines 2, 10–11):  $d\mathcal{L}$   
 519 formulas over fields and method parameters.  $\text{Ens}$  is the  
 520 set of output guarantees for out-port methods (annotated  
 521 with */\* ensures ... \*/*, see Fig. 19, line 4):  $d\mathcal{L}$  formulas over a dedicated program variable with the  
 522 method's name.

524 To verify a class  $C$  against a class specification, both are translated into  $d\mathcal{L}$ -formula (1) that  
 525 expresses safety.

$$526 \quad \text{assumptions}_C \rightarrow [(\text{code}_C; \text{plant}_C)^*] \text{safety}_C \quad (1)$$

```

1 interface Tank {
2   /* requires -1/2 <= newD <= 1/2; */
3   Unit inDrain(Real newD);
4   /* ensures 3 <= outLevel <= 10; */
5   Real outLevel();
6 }
7
8 class CTank(Real inVal)
9   implements Tank {
10  /* requires newD > 0 -> level <= 9.5 */
11  /* requires newD < 0 -> level >= 3.5 */
12  /* timed_requires inDrain < 1 */
13  Unit inDrain(Real newD) { ... }
14  ...
15 }
16
17 /* requires 0 < tick < 1 & inVal > 3.5 */
18 class FlowCtrl(Tank t, Real tick,
19               Real inVal) {
20  /* invariant (drain > 0 -> level <= 9.5)
21   & (drain < 0 -> level >= 3.5) */
22  ...
23 }

```

Figure 19 Annotations in the TankTick model, repeated from Fig. 4

527 The placeholders `assumptionsC`, `codeC`, `plantC`, and `safetyC` (defined formally in Sect. 5.3  
 528 below) encode class `C` and its specification (`inv`, `pre`, `TReq`, `Req`, `Ens`) as follows: The formula  
 529 `assumptionsC` is the conjunction of `pre` and conditions on variables that keep track of time. As  
 530 usual in controller verification, the program repeats a control part `codeC` followed by the continuous  
 531 behavior `plantC`. The condition `safetyC` must hold after an arbitrary number of iterations. It  
 532 combines `inv` with input requirements of in-port methods of referred objects and guarantees of  
 533 own out-port methods.

534 Even though formula (1) `safetyC` is a postcondition that must hold only in the final states of  
 535 the system, we stress that this means at *every real time point* during the continuous dynamics,  
 536 because ODEs advance for a non-deterministic duration while discrete statements take no time.  
 537 The modality, therefore, expresses that whenever `codeC` executes completely, the invariant holds.  
 538 In particular, the invariant holds at the beginning of and throughout the evolution of the continuous  
 539 dynamics in `plantC`. Thus, validity of formula (1) expresses safety of every correctly created  
 540 object (with respect to its specification).

541 The following translation of an HABS class and its specification defines formally how the  
 542 placeholders are composed. The translation is fully automatic and verification is compositional:  
 543 only classes whose code changed explicitly need re-verification, not the whole system.

### 544 5.3 Translation from CHABS to $d\mathcal{L}$

545 We use two operations on sets of programs  $P$ . Operation  $\sum P$  constructs a program that non-  
 546 deterministically executes one of the elements. Operation  $\prod P$  constructs all permutations of  
 547 sequential element-wise execution. Let  $|P| = n$ :

$$548 \quad \sum P = \sum \{p_1, \dots, p_n\} = p_1 \cup p_2 \cup \dots \cup p_n$$

$$549 \quad \prod P = \{p_1; \dots; p_n \mid \forall i, j \leq n. p_i, p_j \in P \wedge (i \neq j \rightarrow p_i \neq p_j)\}$$
 550

551 We translate classes `C` with the following design restrictions: (1) All controllers update their  
 552 local caches of other objects before providing information to those objects (for example, read the  
 553 current water level before instructing the tank to drain or fill); local caches, once updated, are not  
 554 modified later. (2) In-port methods with a timed input requirement are only called from timed  
 555 controllers (for example, a tank that expects to be filled every 5 s is governed by a controller  
 556 running at a corresponding frequency). (3) Duration statements are exact (have two identical  
 557 parameters). (4) Local variable names are unique. The first two constraints fix the interaction  
 558 pattern between components, the last two simplify the presentation. For classes following these  
 559 restrictions, the translation has four phases, each discussed in detail in subsequent paragraphs:  
 560 (i) provision of program variables, (ii) generation of assumptions and safety condition, (iii) control  
 561 code generation, (iv) provision of ODEs and constraints.

#### 562 5.3.1 Program Variables

563 For each field, parameter, and local variable in `C` we create a program variable with the same  
 564 name. For each method `m` we create a time variable  $t_m$ , for each in-port method `m` a tick variable  
 565  $tick_m$ , both type `Real`;  $tick_m$  models the unknown time when an in-port method is *called* next.  
 566 Time variables are local time for each method and determine when a time-triggered controller or  
 567 an in-port is *executed* the next time. We denote the set of all tick variables with `Tick` and the set  
 568 of all time variables with `Time`.

$$\begin{array}{l}
\left. \begin{array}{l}
\text{trans}(f) \equiv f, \text{ where } f \text{ is a } d\mathcal{L} \text{ variable representing field } f \\
\text{trans}(v) \equiv v, \text{ where } v \text{ is a } d\mathcal{L} \text{ variable representing variable } v \\
\text{trans}(e_1 \text{ op } e_2) \equiv \text{trans}(e_1) \text{ op } \text{trans}(e_2)
\end{array} \right\} \text{expressions } e \\
\left. \begin{array}{l}
\text{trans}(\text{if}(e)\{s\}[\text{else}\{s'}]) \equiv \text{if}(\text{trans}(e)) \text{ then } \text{trans}(s)[\text{else } \text{trans}(s')] \\
\text{trans}(\text{while}(e)\{s\}) \equiv \text{while}(\text{trans}(e))\text{trans}(s) \quad \text{trans}(s_1; s_2) = \text{trans}(s_1); \text{trans}(s_2) \\
\text{trans}([T] v = e) \equiv \text{trans}(v) := \text{trans}(e) \quad \text{trans}(f = e) \equiv \text{trans}(f) := \text{trans}(e) \\
\text{trans}(e!m()) \equiv ?\text{true} \quad \text{trans}(f = e.m()) \equiv \text{trans}(f) := *; ?\varphi_m \\
\text{where } \varphi_m \text{ is the postcondition of } m, \text{ with the method name replaced by } \text{trans}(f)
\end{array} \right\} \text{statements } s
\end{array}$$

■ **Figure 20** Translation of expressions  $e$  and statements  $s$

### 5.3.2 Assumptions and Safety Condition

The formula  $\text{assumptions}_{\mathcal{C}}$  (2) is  $\mathcal{C}$ 's precondition  $\text{pre}$  plus all initializations  $\text{init}$  plus conditions on the time and tick variables: in the beginning, each time variable starts at zero and the tick variables have an unknown positive value. Each tick variable  $\text{tick}$  has a method  $m_{\text{tick}}$  that is responsible for its generation. We refer to the timed input requirement of this method with  $\psi(\text{tick})$ , where the method name  $m_{\text{tick}}$  has been replaced with  $\text{tick}$ . The initial value of the tick variable is also described by the timed input requirement and describes when the method is issued for the first time at the latest.

$$\text{assumptions}_{\mathcal{C}} \equiv \text{pre} \wedge \bigwedge_{\varphi \in \text{init}} \varphi \wedge \bigwedge_{t \in \text{Time}} t \doteq 0 \wedge \bigwedge_{\text{tick} \in \text{Tick}} (0 < \text{tick} \wedge \psi(\text{tick})) \quad (2)$$

The formula  $\text{safety}_{\mathcal{C}}$  (3) captures the guarantees of class  $\mathcal{C}$ : we need to show that  $\mathcal{C}$  (i) preserves its own invariant  $\text{inv}$ ; (ii) provides guarantees  $\text{Ens}$  about own out-port methods (shows what others can rely on); (iii) respects timed preconditions  $\text{TReq}^s$ ; and, (iv) when writing to in-port methods of callees, respects their input requirements  $\text{Req}^s$ . If class  $\mathcal{C}$  comes with a time-triggered controller with guard  $\text{duration}(e, e)$ , technical constraint (1) above ensures that at the moment the controller calls an in-port of another object, it has a correct copy of the callee state.  $\text{Req}^s$  are input requirements of used in-port methods of other classes than  $\mathcal{C}$ , where the method parameter is replaced by the field passed to it.  $\text{Ens}$  are guarantees of all out-port methods of  $\mathcal{C}$ . Some special care needs to be taken for timed input requirements. With  $\text{TReq}^s$ , we denote the set of timed input requirements (constructed over  $\text{tick}$ , as above) of all called in-ports where such a clause is given.

$$\text{safety}_{\mathcal{C}} \equiv \text{inv} \wedge \bigwedge_{\varphi \in \text{Req}^s} \varphi \wedge \bigwedge_{\tau \in \text{TReq}^s} \tau \wedge \bigwedge_{\psi \in \text{Ens}} \psi \quad (3)$$

The safety condition expresses that the controllers of class  $\mathcal{C}$  respect the input requirements when writing to the in-port methods of other components and call in-port methods with a timed input requirement sufficiently open. The structure of controllers in CHABS per Sect. 3.3 enforces that these calls occur last in the controller bodies.

### 5.3.3 Control Code

The translation of ABS statements to hybrid programs is defined in Fig. 20. We discuss the non-obvious rules: Calls  $e!m()$  to in-port methods of other objects are mapped to  $?true$  (i.e. skip), because there is no effect on the caller object. A read  $f=e.m()$  from an out-port method is mapped

598 to  $\text{trans}(f) := *; ?\varphi_m$ : a non-deterministic assignment, restricted with a subsequent test for the  
 599 guarantee of the called out-port method.

600 The translation of ports and control methods has the *general form* **if** (check) **then** {exec; cleanup}.  
 601 This pattern is instantiated per method type as follows:

- 602 ■ Time-triggered controller  $m$  with method body **await duration**( $e, e$ );  $s$ ; **this.m()**: check makes  
 603 sure the correct duration elapsed and cleanup resets time, so  $\text{check} \equiv t_m \doteq \text{trans}(e)$ ,  $\text{exec} \equiv$   
 604  $\text{trans}(s)$ ,  $\text{cleanup} \equiv t_m := 0$ .
- 605 ■ Event-triggered controller  $m$  with body **await diff**  $e$ ;  $s$ ; **this.m()**: check tests the guard, so  
 606  $\text{check} \equiv \text{trans}(e)$ ,  $\text{exec} \equiv \text{trans}(s)$ ,  $\text{cleanup} \equiv ?\text{true}$ .
- 607 ■ In-port method  $m$  with body **this.f** =  $v$ , input requirement  $\varphi$  and timed input requirement  $\psi$ :  
 608 check ensures the correct duration elapsed, so  $\text{check} \equiv t_m \doteq \text{tick}_m$ ; exec chooses a value consistent  
 609 with  $\varphi$ , so  $\text{exec} \equiv f := *; ?\varphi$ ; finally, cleanup does the same for a new duration consistent with  
 610  $\psi$  (method name replaced by  $\text{tick}_m$ ), so  $\text{cleanup} \equiv \text{tick}_m := *; ?\text{tick}_m > 0; ?\psi; t_m := 0$ .
- 611 ■ Out-port methods and the **run** method are not translated. Out-port methods have no effect  
 612 on object state and their guarantees (included in (1) in  $\text{safety}_C$ ) must be shown to hold  
 613 throughout plant execution. The **run** method initializes the system and ensures that every  
 614 controller can run once before the first plant execution, which is guaranteed in (1) through  
 615 sequential composition of  $\text{code}_C$ ;  $\text{plant}_C$ .

616 Let  $M$  be the set of all translations of in-port methods and controllers, then:

$$617 \quad \text{code}_C \equiv \left( \sum \prod M \right); \left( \sum M \right)^* \quad (4)$$

618 The controller  $\text{code}_C$  first executes all controllers in a non-deterministically chosen order  
 619  $(\sum \prod M)$ , then allows each controller/in-port to repeat  $(\sum M)^*$ . The latter replicates eager ABS  
 620 behavior on satisfied guards: when an event-triggered controller is triggered and its guard still  
 621 holds after its execution, then in ABS the controller is run again.

622 Note that  $(\sum M)^*$  safely overapproximates all possible orders, including the behavior of the  
 623 first part  $\sum \prod M$ . However, including  $\sum \prod M$  in  $\text{code}_C$  simplifies practical proofs, because in  
 624 typical models that disable the check guards at the end of control and in-port method bodies (e.g.,  
 625 a time-triggered controller that resets time in cleanup so that it becomes re-enabled only after  
 626 some time passes), every method is executed at most once before time advances. The structure of  
 627 the controller  $\text{code}_C$  mirrors this with the first part  $\sum \prod M$  to simplify practical proofs as follows:  
 628 (i) the proof obligations of enabled control and in-port methods (i.e., whose check is true) are  
 629 easier because the outer loop is dropped, and additionally the proof obligations of all the disabled  
 630 control and in-port methods can be easily disposed of by contradiction with their check guards;  
 631 (ii) finding a loop invariant for the second part  $(\sum M)^*$  is easy when no method is executed twice  
 632 before time advances: in that case, the loop invariant for  $(\sum M)^*$  must simply imply that none of  
 633 the check guards holds. Further note that  $\sum \prod M$  does not exclude runs, because the general  
 634 form **if** (check) **then** {exec; cleanup} of control methods and ports in  $M$  ensures that there is  
 635 progress through the implicit **else**  $?true$  even if all controllers and in-ports are disabled.

### 636 5.3.4 Plant

637 The plant of a class  $C$  has the form

$$638 \quad \text{plant}_C \equiv \sum \{(\text{ode}, \text{ode}_t \ \& \ c) \mid c \in \mathcal{C}\} \ , \quad (5)$$

639 where  $\text{ode}$  is the ODE from its physical block,  $\text{ode}_t$  describes the time variables, and the constraints  
 640  $c \in \mathcal{C}$  partition the domain of the physical fields. The boundaries of the subdomains overlap



641 exactly where the differential guards hold.<sup>7</sup> This models guards as events in  $d\mathcal{L}$ , following the  
 642 modeling pattern described in Sect. 5.1. To ensure that no differential guard is omitted, it is  
 643 necessary that no two differential guards share a program variable. This is not a restriction, as  
 644 two controllers can be merged with a disjunction: see the guard in Fig. 2.

645 To define  $\mathcal{C}$  let  $e_1, \dots, e_m$  be the translations of differential guards in the class and  $\tilde{e}_i$  the weak  
 646 complement of  $e_i$ . Let  $t_1, \dots, t_l$  be all time variables introduced for time-triggered controllers with  
 647  $e_{t_i}$  the expression in the **duration** statement. Let  $pt_1, \dots, pt_k$  be all time variables introduced for  
 648 in-port methods and  $tick_{pt_i}$  the associated tick variable. We set  $\text{ode}_t \equiv \{t'_1 = 1, \dots, t'_l = 1, pt'_1 =$   
 649  $1, \dots, pt'_k = 1\}$  and define:

$$650 \quad \mathcal{C} \equiv (\{e_1, \tilde{e}_1\} \times \{e_2, \tilde{e}_2\} \times \dots \times \{e_m, \tilde{e}_m\}) \\ 651 \quad \cup \{t_1 \leq e_{t_i}\}_{i \leq l} \cup \{t_1 \geq e_{t_i}\}_{i \leq l} \cup \{pt_i \leq tick_{pt_i}\}_{i \leq n} \cup \{pt_i \geq tick_{pt_i}\}_{i \leq n}$$

### 653 5.3.5 On the Random Number Generator

654 We do not translate the **random(i)** expression from HABS to  $d\mathcal{L}$ , because its semantics is that it  
 655 returns an *integer* below  $i$ . However, integer arithmetic is undecidable, which is the reason why  $d\mathcal{L}$   
 656 opts to embed its modality into a decidable first-order logic over the reals [66]. A straightforward  
 657 overapproximation with a translation to a variation of **random** that returns a real value is:

$$658 \quad \text{trans}(\mathbf{f} = \text{random}(\mathbf{x})) \equiv \text{trans}(\mathbf{f}) := *; ?(0 \leq \text{trans}(\mathbf{f}) < \text{trans}(\mathbf{x}))$$

## 659 5.4 Compositional Verification

660 We can now state our main theorem: If we can prove safety of all classes, i.e., close all proof  
 661 obligations, then the whole system is safe, i.e., every class indeed preserves its invariant. Verification  
 662 is compositional: if we change the code or invariant of one class, only the proof obligation of this  
 663 class has to be reproven. If we change a method precondition, additionally the proof obligations  
 664 of all calling classes have to be reproven.

665 ► **Theorem 5.** *Let  $\mathbf{P}$  be a set of classes, with each  $\mathbf{c} \in \mathbf{P}$  associated with  $\varphi_{\mathbf{c}}$  per formula (1). If*  
 666 *all the  $\varphi_{\mathbf{c}}$  are valid, then for every main block that creates objects satisfying  $\text{pre}_{\mathbf{c}}$  all reachable*  
 667 *states of all objects satisfy  $\text{inv}_{\mathbf{c}}$ .*

668 **Proof Sketch.** Recall that the trace of an HAO is an assignment of time to stores (Def. 4). For  
 669 the proof, each store is indexed by its time and the trace starts with 0 (i.e., the possible offset  
 670 caused by the delayed object creation is removed):

$$671 \quad \theta_o(t) = (\rho_t)_{t \in \mathbb{R}^+} = \rho_0 \dots$$

672 We are going to use that there are only countably many discrete steps in a run and partition the  
 673 trace into countably many substraces. Then we show by induction on these discrete steps that the  
 674 invariant is always preserved.

675 Let  $D$  be the set of all time points with discrete steps of  $o$  in the run that generates  $\theta_o$ . Note  
 676 that  $0 \in D$  and that  $\theta_o(d)$  is the last store defined by the SOS semantics, if several such stores  
 677 share the same time; further note that this is reflecting the reachability relation of  $d\mathcal{L}$ .

678 We define  $\theta_o^d$  as the subtrace of  $\theta_o$  starting with  $d$  and ending at the next time point of a  
 679 discrete step. Let  $\text{next}(d)$  be the next time point of a discrete step after  $d$ , if such a time point

<sup>7</sup> Expressions contain only  $>=$ ,  $<=$ , so weak complement ensures a boundary overlap.

680 exists, and  $\infty$  otherwise:

$$681 \quad \mathbf{dom}(\theta_o^d) = [d..\mathbf{next}(d)] \quad \text{with} \quad \theta_o^d(t) = \theta_o(t)$$

682 We observe that each state in the HABS semantics is also a state in the Kripke structure of  
 683 the semantics if all class parameters are removed. We show that  $\mathbf{trans}$  preserves reachability: if  
 684 from a state  $\rho$  state  $\rho'$  is reachable by an HABS statement  $\mathbf{s}$  in the HABS semantics, then state  $\rho'$  is  
 685 reachable from state  $\rho$  by  $\mathbf{trans}(\mathbf{s})$ . This is justified as follows:

- 686 1. The  $d\mathcal{L}$  program omits no events, because each event is at a boundary of two evolution domain  
 687 constraints on a variable and no two events share a variable (each controller has its own time  
 688 variable).
- 689 2. The evolution domain constraints cover all possible states, so no run is rejected for a domain  
 690 being too small.
- 691 3. Each test in  $d\mathcal{L}$  formula  $\varphi_{\mathcal{C}}$  that discards runs does so using a condition that is provably  
 692 guaranteed by other objects. For example, the test that discards all runs of an in-port method  
 693 for inputs not satisfying its input requirements is safe, because on the caller side this condition  
 694 is part of the safety condition (3).
- 695 4. The observation also relies on technical constraint (1) above and the recursive call being at  
 696 the end of a controller. Together, this guarantees that *at that moment* the caller copy of the  
 697 callee's state is consistent with the callee's actual state.

698 Let  $D = (d_i)_{i \in \mathbb{N}}$  be an enumeration of the discrete time points and  $\hat{\theta}_o^{d_i}$  the union of all subtraces  
 699 of  $\theta_o$  up to  $d_i$ :

$$700 \quad \mathbf{dom}(\hat{\theta}_o^{d_i}) = \bigcup_{j \leq i} \mathbf{dom}(\theta_o^{d_j}) \quad \text{with} \quad \hat{\theta}_o^{d_i}(t) = \theta_o(t)$$

701 We show by induction on  $i$  that every state in  $\hat{\theta}_o^{d_i}$  is safe, i.e., a model for the invariant  $\mathbf{inv}_{\mathcal{C}}$ .

702 **Induction Base:**  $i = 0$ . It is explicitly checked that  $\theta_o^{d_0}$  is safe. By assumption, the object is  
 703 created in a state  $\theta_o^{d_0}$  such that the precondition  $\mathbf{pre}_{\mathcal{C}}$  holds. From axiom 1 of  $d\mathcal{L}$  [68] we know  
 704 that the safety condition must be true in the beginning of the loop, thus validity of  $\varphi_{\mathcal{C}}$  implies  
 705 validity of  $\mathbf{pre}_{\mathcal{C}} \rightarrow \mathbf{inv}_{\mathcal{C}}$ . Since all the formulas  $\varphi_{\mathcal{C}}$  are proved in isolated component proofs,  
 706 we conclude  $\mathbf{inv}_{\mathcal{C}}$  holds for all reachable states of all objects as by the correctness argument  
 707 reachability is preserved.

708 **Induction Step.**  $i > 0$ . This is analogous to the base case, but instead of an explicit check that  $d_i$   
 709 is safe, we use the induction hypothesis that every state in  $\hat{\theta}_o^{d_{i-1}}$  is safe and that the statement  
 710 for  $d_i$  is executed in a state at time  $t \in \mathbf{dom}(\hat{\theta}_o^{d_{i-1}})$ .  $\blacktriangleleft$

711  $\blacktriangleright$  **Remark.** The theorem states soundness of safety properties in  $d\mathcal{L}$  proof obligations and does not  
 712 prove semantic equivalence between the contained  $d\mathcal{L}$ -program and the HABS class. This approach  
 713 stands in the tradition of modular deductive verification of object-oriented software, in particular,  
 714 it follows the structure of systems for distributed object-oriented programs [52]. The main reason  
 715 to pursue this approach is that the form of proof obligations and the translation of statements  
 716 cannot be disentangled: the translation of method calls includes the postcondition of the called  
 717 methods: soundness of the translation relies on the fact that all other proof obligations can be  
 718 established. This is already the case for discrete, sequential languages [41]. Note that this is *not*  
 719 circular. As the proof of Theorem 5 shows, we can order all method executions in a run such that  
 720 we have a well-founded induction on them. The first method execution in every object relies only

721 on the state precondition which is guaranteed at creation. These in turn are guaranteed in the main  
 722 block, which has no assumptions. Another reason is that each  $d\mathcal{L}$  proof obligation corresponds to  
 723 the (symbolic) execution of one *object* in a class. To model all permissible evolutions of several  
 724 method executions in a proof, therefore, it is necessary to encode the scheduler. This requires a  
 725 form of proof obligation that assumes the object invariant (which contains scheduling constraints).  
 726 This effect is well-known in deductive verification of distributed programs [31, 32, 52].

## 727 5.5 Case Study

728 We illustrate the HABS-to-KeYmaera X translation defined above with the **TankTick** system in  
 729 Fig. 4. The example, the implementation of the translation and the simulation, as well as the  
 730 mechanical proofs of the translation are available in the supplementary material.<sup>8</sup> We start with  
 731 the two-object water tank, whose behavior for an initial level of 5 *l* is plotted in Fig. 5.

### 732 5.5.1 Class `CTank`

733 The in-port method `inDrain()` of the `CTank` class gives rise to a time variable  $t_{\text{inDrain}}$  and a tick  
 734 variable  $tick_{\text{inDrain}}$ . Following (2),  $\text{assumptions}_{\text{Tank}}$  is:

$$\begin{aligned} \text{assumptions}_{\text{Tank}} \equiv & 4 \leq \text{inVal} \leq 9 \\ & \wedge t_{\text{inDrain}} \dot{=} 0 \wedge 0 < tick_{\text{inDrain}} \\ & \wedge \text{level} \dot{=} \text{inVal} \wedge \text{drain} \dot{=} -1/2 \end{aligned} \quad (6)$$

736 The safety condition says the tank level stays within its limits and that `level` adheres to its  
 737 contract which happen to be identical. No in-port methods of other classes are used, hence:

$$\text{safety}_{\text{Tank}} \equiv 3 \leq \text{level} \leq 10 . \quad (7)$$

739 The `CTank` class has no controller method, so the `inDrain` method, which has a timed input  
 740 requirement, per (4) results in  $\text{code}_{\text{Tank}}$  below

$$\text{code}_{\text{Tank}} \equiv \text{p}; (\text{p})^* \quad (8)$$

742 where  $\text{p} \equiv \text{trans}(\text{inDrain})$  below is translated from Fig. 4 using the translation of Fig. 20:

$$\begin{aligned} \text{p} \equiv & \text{if } (t_{\text{inDrain}} \dot{=} tick_{\text{inDrain}}) \text{ then} \\ & \text{drain} := *; \\ & ?-1/2 \leq \text{drain} \leq 1/2 \wedge (\text{drain} < 0 \rightarrow \text{level} \geq 3.5) \\ & \quad \wedge (\text{drain} > 0 \rightarrow \text{level} \leq 9.5); \\ & tick_{\text{inDrain}} := *; ?0 < tick_{\text{inDrain}} < 1; t_{\text{inDrain}} := 0 \end{aligned}$$

749 The plant  $\text{plant}_{\text{Tank}}$ , following shape (5), is based on the physical block and the new clock  
 750 variable (there are no differential guards), with the evolution domain constraint split along the  
 751 new time variable  $t_{\text{inDrain}}$ . ODEs of the form  $v' = 0$  are default and omitted.

$$\begin{aligned} \text{plant}_{\text{Tank}} & \equiv \text{plant}_{\text{Tank}}^{\leq} \cup \text{plant}_{\text{Tank}}^{\geq} \\ \text{plant}_{\text{Tank}}^{\leq} & \equiv \{\text{level}' = \text{drain}, t'_{\text{inDrain}} = 1 \ \& \ t_{\text{inDrain}} \leq tick_{\text{inDrain}}\} \\ \text{plant}_{\text{Tank}}^{\geq} & \equiv \{\text{level}' = \text{drain}, t'_{\text{inDrain}} = 1 \ \& \ t_{\text{inDrain}} \geq tick_{\text{inDrain}}\} \end{aligned} \quad (9)$$

<sup>8</sup> <https://doi.org/10.5281/zenodo.5973904>

753 ► **Lemma 6.** *Class Tank is safe, i.e., formula  $\varphi_{\text{Tank}}$ —obtained per (1) referring to tank assumptions*  
 754 *assumptions<sub>Tank</sub> (6), postcondition safety<sub>Tank</sub> (7), code code<sub>Tank</sub> (8), and plant plant<sub>Tank</sub> (9)—is*  
 755 *valid.*

$$756 \quad \varphi_{\text{Tank}} \equiv \text{assumptions}_{\text{Tank}} \rightarrow [(\text{code}_{\text{Tank}}; \text{plant}_{\text{Tank}})^*] \text{safety}_{\text{Tank}}$$

758 **Proof.** See KeYmaera X proofs in the supplementary material. The proof sketch here serves as  
 759 an illustration of how sequent proofs in KeYmaera X systematically use the invariant annotations  
 760 in HABS. In the proof, we show the inductive loop invariant  $\text{inv}_{\text{Tank}}^{\leq}$ , which expresses that the level  
 761 always stays within limits and that the next input will be supplied before exceeding the timed input  
 762 requirement as follows:  $3 \leq \text{level} \leq 10 \wedge -1/2 \leq \text{drain} \leq 1/2 \wedge 3 \leq \text{level} + \text{drain}(\text{tick}_{\text{inDrain}} -$   
 763  $t_{\text{inDrain}}) \leq 10 \wedge \text{tick}_{\text{inDrain}} \leq t_{\text{inDrain}}$ .

764 The proof starts in step  $\rightarrow_R$  to make the left-hand side assumptions<sub>Tank</sub> of the implication  
 765 available as assumptions. Next, [\*] uses the loop invariant  $\text{inv}_{\text{Tank}}^{\leq}$  for induction: the base case  
 766 in the left-most subgoal and the use case in the right-most subgoal follow by real arithmetic  
 767 automation; the induction step in the middle subgoal continues with [;] to split the sequential  
 768 composition into nested box modalities.

$$\begin{array}{c}
 \begin{array}{c}
 \text{auto;} \frac{*}{\text{inv}_{\text{Tank}}^{\leq} \vdash [\text{plant}_{\text{Tank}}^{\leq}] \text{inv}_{\text{Tank}}^{\leq}} \quad \text{contradiction;} \frac{*}{\text{inv}_{\text{Tank}}^{\leq} \vdash [\text{plant}_{\text{Tank}}^{\geq}] \text{inv}_{\text{Tank}}^{\leq}} \\
 \hline
 [\cup], \wedge_R \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{plant}_{\text{Tank}}^{\leq} \cup \text{plant}_{\text{Tank}}^{\geq}] \text{inv}_{\text{Tank}}^{\leq} \\
 \hline
 \text{expand} \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{plant}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq}
 \end{array} \\
 \\
 \begin{array}{c}
 \text{auto;} \frac{*}{\text{inv}_{\text{Tank}}^{\leq} \vdash [\text{p}] \text{inv}_{\text{Tank}}^{\leq}} \quad \text{auto;} \frac{*}{\text{inv}_{\text{Tank}}^{\leq} \vdash [\text{p}^*] \text{inv}_{\text{Tank}}^{\leq}} \\
 \hline
 [;], \text{MR} \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{p}; (\text{p}^*)] \text{inv}_{\text{Tank}}^{\leq} \\
 \hline
 \text{expand} \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{code}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq} \quad \dots \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{plant}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq} \\
 \hline
 \text{MR} \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{code}_{\text{Tank}}] [\text{plant}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq} \\
 \hline
 [;] \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{code}_{\text{Tank}}; \text{plant}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq}
 \end{array} \\
 \\
 \begin{array}{c}
 \text{auto;} \frac{*}{\text{assumptions}_{\text{Tank}} \vdash \text{inv}_{\text{Tank}}^{\leq}} \quad \dots \quad \text{inv}_{\text{Tank}}^{\leq} \vdash [\text{code}_{\text{Tank}}; \text{plant}_{\text{Tank}}] \text{inv}_{\text{Tank}}^{\leq} \quad \text{auto;} \frac{*}{\text{inv}_{\text{Tank}}^{\leq} \vdash \text{safety}_{\text{Tank}}} \\
 \hline
 [*] \quad \text{assumptions}_{\text{Tank}} \vdash [(\text{code}_{\text{Tank}}; \text{plant}_{\text{Tank}})^*] \text{safety}_{\text{Tank}} \\
 \hline
 \rightarrow_R \quad \text{assumptions}_{\text{Tank}} \rightarrow [(\text{code}_{\text{Tank}}; \text{plant}_{\text{Tank}})^*] \text{safety}_{\text{Tank}}
 \end{array}
 \end{array}$$

769 The main insight now is that code<sub>Tank</sub> reacts at the latest when  $\text{tick}_{\text{inDrain}} = t_{\text{inDrain}}$  and  
 770 will reset the timer using  $\text{tick}_{\text{inDrain}} := 0$ , so that the timing requirement  $\text{tick}_{\text{inDrain}} \leq t_{\text{inDrain}}$   
 771 can be strengthened to a strict inequality  $\text{tick}_{\text{inDrain}} < t_{\text{inDrain}}$  in the inductive loop invariant.  
 772 The resulting intermediate condition  $\text{inv}_{\text{Tank}}^{\leq}$  is used in step MR to split into two subgoals: in the  
 773 left subgoal of MR, we show that code<sub>Tank</sub> guarantees the intermediate condition  $\text{inv}_{\text{Tank}}^{\leq}$ . In the  
 774 right subgoal of MR we show that plant<sub>Tank</sub> preserves the loop invariant from that intermediate  
 775 condition: the plant listens for the event  $\text{tick}_{\text{inDrain}} = t_{\text{inDrain}}$  with a choice between two  
 776 differential equations, whose evolution domain constraints exactly overlap at the event. On  
 777 evolution domain  $\text{tick}_{\text{inDrain}} \leq t_{\text{inDrain}}$  in  $\text{plant}_{\text{Tank}}^{\leq}$ , the differential equation preserves the loop  
 778 invariant, whereas on evolution domain  $\text{tick}_{\text{inDrain}} \geq t_{\text{inDrain}}$  in  $\text{plant}_{\text{Tank}}^{\geq}$  the contradiction  
 779 shows that the controller reacts such that the plant can never enter this unsafe behavior. ◀  
 780

## 781 5.5.2 Time-Triggered Controller FlowCtrl

782 Assumptions assumptions<sub>FlowCtrl</sub> of FlowCtrl constructed per (2) and plant plant<sub>FlowCtrl</sub> con-  
 783 structed per (5) are straightforward. The latter is created for the sake of observing time events,

784 even though no physical block is present:

$$785 \quad \text{assumptions}_{\text{FlowCtrl}} \equiv 0 < \text{tick} < 1 \quad (10)$$

$$786 \quad \text{plant}_{\text{FlowCtrl}} \equiv \{t'_{\text{ctrlFlow}} = 1 \ \& \ t_{\text{ctrlFlow}} \geq \text{tick}\} \quad (11)$$

$$787 \quad \cup \{t'_{\text{ctrlFlow}} = 1 \ \& \ t_{\text{ctrlFlow}} \leq \text{tick}\}$$

789 The safety condition  $\text{safety}_{\text{FlowCtrl}}$  constructed per (3) is the timed input requirement of the  
790 called `inDrain` method and the class invariant (subsumed by the input requirement of `inDrain`):

$$791 \quad \text{safety}_{\text{FlowCtrl}} \equiv -1/2 \leq \text{drain} \leq 1/2 \wedge \text{tick} < 1$$

$$\wedge (\text{drain} < 0 \rightarrow \text{level} \geq 3.5) \quad (12)$$

$$\wedge (\text{drain} > 0 \rightarrow \text{level} \leq 9.5)$$

792 Finally, the code  $\text{code}_{\text{FlowCtrl}}$  is translated as

$$793 \quad \text{code}_{\text{FlowCtrl}} \equiv \text{q}; (\text{q})^* \quad (13)$$

794 with

```
795 q ≡ if (tctrlFlow = tick) then
796   level := *; ?3 ≤ level ≤ 10;
797   if (level ≤ 3.5) then {drain := 1/2};
798   if (level ≥ 9.5) then {drain := -1/2};
799   tctrlFlow := 0
800
```

801 ► **Lemma 7.** *Class `FlowCtrl` is safe, i.e., formula  $\varphi_{\text{FlowCtrl}}$ —obtained per (1) referring to*  
802 *assumptions  $\text{assumptions}_{\text{FlowCtrl}}$  (10), postcondition  $\text{safety}_{\text{FlowCtrl}}$  (12), code  $\text{code}_{\text{FlowCtrl}}$  (13),*  
803 *and plant  $\text{plant}_{\text{FlowCtrl}}$  (11)—is valid.*


$$804 \quad \varphi_{\text{FlowCtrl}} \equiv \text{assumptions}_{\text{FlowCtrl}} \rightarrow [(\text{code}_{\text{FlowCtrl}}; \text{plant}_{\text{FlowCtrl}})^*] \text{safety}_{\text{FlowCtrl}}$$

806 **Proof.** See KeYmaera X-proofs in the supplementary material. ◀

### 807 5.5.3 Event-Triggered Controller `CSingleTank`

808 Translation of class `CSingleTank` from Fig. 2 illustrates the handling of event-triggered controllers.

809 The plant and code interact. The plant separates the evolution  
810 domain into two parts, with the guard of the event-triggered  
811 controller (the white areas in Fig. 21) defining their boundary. The  
812 gray areas are *larger* than the safe region defined by  $3 \leq \text{level} \leq$   
813  $10$ . This is necessary to avoid Zeno behavior in the eager execution  
814 semantics of HABS: If we used simply the weak complement of the  
815 safe region  $\text{level} \leq 3 \mid \text{level} \geq 10$  as a guard and happen to  
816 be in a program state at the boundary (the lower of the states  
817 indicated with a star in Fig. 21), then the controller changes the state as shown by the arrow.  
818 But if the next state is again *on the boundary*, which is the case when the safe region is too small,  
819 then the guard is triggered, the controller loops back to the first state, etc., without physical time  
820 being able to advance. The guard in Fig. 2 ensures that after the controller has run, the state is



821 **Figure 21** Avoiding Zeno-behavior in `TankMono`.

821 *not* on the boundary anymore. This behavior is exhibited by our implementation, see Fig. 3. The  
822 code `codeCSingleTank` has the form  $r; (r)^*$  with  $r$  being:

```
823   r ≡ if (level ≤ 3 ∧ drain ≤ 0) ∨ (level ≥ 10 ∧ drain ≥ 0) then
       if (level ≤ 3) then drain := 1/2 else drain := -1/2
```

824 The plant of `CSingleTank` with sufficiently large regions is as follows:

```
825   plantCSingleTank ≡
826     {level' = drain ∧ (level ≤ 3 ∧ drain ≤ 0) ∨ (level ≥ 10 ∧ drain ≥ 0)}
827     ∪ {level' = drain ∧ (level ≥ 3 ∨ drain ≥ 0) ∧ (level ≤ 10 ∨ drain ≤ 0)}
828
```

## 829 5.6 On Translation into $d\mathcal{L}$

830 HABS programs can be tested and validated, but the programmer needs to avoid writing programs  
831 that are **1.** inherently difficult to interpret and **2.** have a high degree of non-determinism. Both  
832 are good programming and software engineering practices, of course, and the fact that HABS is a  
833 *programming language* enables one to apply standard techniques for discrete programs.

834 A back-translation from  $d\mathcal{L}$  to HABS would provide meaningful validation only for deterministic  
835  $d\mathcal{L}$  models. While being possible even in the general case, two traits of  $d\mathcal{L}$  programs prohibit easy  
836 interpretation and simulation:

837 **Highly Non-Deterministic Structure** Additionally to non-deterministic assignment, branching  
838 and repetition are both non-deterministic: the—rather non-intuitive—representation of  $(s)^*$  in  
839 HABS is a loop that non-deterministically chooses to break out.

```
840   1 while(True) {
      2     Int i = random(2);
      3     if ( i == 1 ) break;
      4     s;
      5 }
```

841 This loop may never terminate, while the semantics of  $d\mathcal{L}$  loops defines an arbitrary but  
842 countable number of repetitions. A similar pattern has to be employed for branching.

843 **Tests** The test  $?φ$  discards a run based on a  $d\mathcal{L}$ -guard. Translation would require **1.** to evaluate  
844  $d\mathcal{L}$  formulas, as opposed to Boolean expressions, and **2.** a mechanism to abort the program.  
845 This can be emulated by exceptions, but it obfuscates the semantics.

## 846 6 Related & Future Work, Conclusion

### 847 6.1 Related Work

848 Previous work on hybrid programming concentrated on purely sequential languages: `HybCore` [39]  
849 is a while-language with hybrid behavior and a simulator [40], but lacks formal verification  
850 techniques. Its extensional semantics is not able to express the timed properties needed for our  
851 distributed controller. `Whiledtc` [77] is also a while-language and uses infinitesimals instead of ODEs  
852 to model continuous dynamics. It has a simple verification system based on Hoare triples [42], but  
853 is not executable.

854 Hybrid Rebeca (HR) [46] proposes to embed hybrid automata directly into the actor language  
855 Rebeca. In contrast to HABS, no simulation is available and verification is not object-modular: the  
856 whole model is translated to a single monolithic hybrid automaton. Because of this, a number

857 of boundedness constraints have to be imposed. The translation is also the semantics: HR has  
858 no semantics beyond this translation and is mainly a frontend for Hybrid Automata tools. The  
859 verification backend of HR does not support non-linear ODEs (our examples are linear, but HABS,  
860 KeYmaera X, and Maxima, support non-linear ODEs; HABS models with non-linear ODEs are  
861 found in the online supplement).

862 Recent efforts [58, 64] split the verification task in  $d\mathcal{L}$  into manageable pieces by modularizing  
863 deductive hybrid systems verification with component-based modeling and verification techniques,  
864 but impose strict structural requirements on components and communication. The Sphinx  
865 modeling tool [62] for  $d\mathcal{L}$  represents non-distributed hybrid programs with UML class and activity  
866 diagrams, but for verification purposes it translates these model artifacts into a single monolithic  
867 hybrid program.

868 The Architecture Analysis and Design Language (AADL), a language to model hardware and  
869 software components in embedded systems, has a hybrid extension [2], which uses the HHL [80]  
870 theorem prover as its verification backend [1]. HHL is based on Hoare triples over hybrid CSP  
871 programs and duration calculus formulas [57]. Hybrid AADL offers structuring elements for  
872 components and their connections on the architecture level. The semantics of hybrid AADL is  
873 given as a translation of the *synchronous* fragment of AADL into hybrid CSP, while we extend  
874 the semantics of the actor-based programming language ABS to combine reasoning about the  
875 asynchronous behavior of communicating components in ABS with reasoning about the internal  
876 combined discrete and continuous component behavior in differential dynamic logic. As a side  
877 effect, the extended semantics enables proving the correctness of the translation to differential  
878 dynamic logic, as well as translating HABS to other formal languages.

879 A similar approach based on Stateflow/Simulink is implemented in the MARS toolkit [22]. The  
880 MARS approach is orthogonal to HABS: MARS connects a verification toolkit around a simulation  
881 language (which is a daunting task given the missing formal semantics of Stateflow/Simulink),  
882 while HABS is designed specifically to enable verification and simulation through its languages  
883 features. This is reflected in the soundness proof, which is based on a *bidirectional* translation.

884 Another approach based on CSP and the duration calculus combines these formalisms with  
885 Object-Z [45]. This enables model-checking for real-time systems (clocks with resets), while  
886 we support hybrid systems theorem proving with (non-linear) differential equations. A further  
887 integration of Object-Z and (Timed) CSP was investigated by Mahony & Dong [60].

888 Hybrid Event-B [12, 13] extends Event-B refinement reasoning with continuous behavior  
889 between the usual discrete Event-B events. A more lightweight approach [76, 21] models hybrid  
890 systems in an abstract way as action systems without differential equations directly in Event-B,  
891 and complements analysis in Event-B with simulation in Matlab. Similarly, Dupont et al. [34]  
892 use Event-B for a correct-by-construction approach to hybrid systems. They embed the ODEs  
893 used for continuous modeling by declaring them as a special theory within Event-B instead of  
894 extending the core language itself.

895 Integrated tools such as Ptolemy [71], Stateflow/Simulink except the aforementioned MARS  
896 toolkit, and Modelica, all emphasize simulation, reachability analysis (e.g., CHARON [6, 7], Ariadne  
897 [15]), or testing (e.g., [30]). As supporting techniques, they provide modeling notation for timing  
898 aspects, signals, and data flow between heterogeneous models. Formal verification of hybrid  
899 systems with reachability analysis and model checking tools (SpaceEx [35], CORA [4], Flow\* [23])  
900 support modularity [33] based on hybrid I/O automata [59], assume-guarantee reasoning [17, 43],  
901 and hybridization [24]. However, they work best for finite-horizon analysis and finite regions  
902 (because over-approximations stay tight only for bounded time and from small starting regions).  
903 Similar restrictions apply to dReal/dReach [37, 55].

904 Dynamic I/O automata [9] for modeling dynamic systems introduce a notion of externally



905 visible behavior, the ability to create and destroy automata and change their signature dynamically;  
 906 those features are all naturally available in our object-oriented approach and do not need special  
 907 extension like automata-based modeling tools. Our work contrasts with all mentioned simulation  
 908 and verification approaches by providing a uniform modeling language, validation by simulation,  
 909 modular infinite-horizon and infinite-region theorem proving through translation from HABS to  $d\mathcal{L}$ .

910 Translation among hybrid system languages so far centers around hybrid automata as a  
 911 unifying concept [11, 79]. Others focus on the discrete fragment [38]. Our translation from HABS  
 912 to  $d\mathcal{L}$  translates complete hybrid system models written in a *programming language*, including  
 913 annotations (preconditions, invariants, etc.). It is sound relative to the formal semantics of HABS  
 914 and  $d\mathcal{L}$ .

915 Hybrid systems validation through simulation is addressed with translation to Stateflow/Simulink  
 916 [10]; with a combination of discrete-event and numerical methods [19]; and with co-simulation  
 917 between control software and dedicated physics simulators [26, 78, 82]. Here, we focus on safety  
 918 verification, the distributed aspect of HABS models, and take a pragmatic first step for simulating  
 919 continuous models.

920 In summary, HABS is designed for modular deductive verification (unlike simulation-centric  
 921 tools), infinite-horizon analysis on infinite regions (unlike reachability analysis and model checking  
 922 tools), without sacrificing high-level programming language features (unlike hybrid systems  
 923 modularization techniques and assume-guarantee reasoning).

## 924 6.2 Future Work

925 The present work lifts the research on formal semantics of programming languages for hybrid  
 926 systems from verification-centric minimalistic languages to distributed object-oriented languages.  
 927 Carrying over techniques, ideas, and analyses from programming language research to hybrid  
 928 systems programming, presents an intriguing research direction. Our ongoing work on larger  
 929 case studies with HABS, in particular in connection with co-simulation [54], is expected to reveal  
 930 additional challenges.

931 We plan to combine the verification of CHABS presented here with the more modular approach  
 932 based on post-regions [51], which does not support timed input requirements yet. Future research  
 933 avenues include investigating how the static analyses for ABS, in particular the deadlock analysis for  
 934 boolean guards [50], can be extended for HABS, extending approximate simulation of non-solvable  
 935 differential equations, experimenting with various computer algebra systems, and supporting  
 936 guards with non-urgent semantics.

## 937 6.3 Conclusion

938 Distributed hybrid systems are not only difficult to *verify* formally, it is equally hard to *validate*  
 939 a formal model of them, especially with components using symbolic computations, such as servers.  
 940 Both activities have conflicting demands, so we propose a translation-based approach: modeling is  
 941 guided by patterns over hybrid programs and class specifications in HABS, a hybrid extension of  
 942 the concurrent active-object language ABS. These are automatically decomposed and translated  
 943 (Thm. 5) into sequential proof obligations of the verification-oriented differential dynamic logic  $d\mathcal{L}$   
 944 and discharged by the hybrid theorem prover KeYmaera X.

945 We illustrated the viability of our approach by a case study that features many complications:  
 946 concurrent behavior, possible non-termination, correctness depending on timing constants, multi-  
 947 dimensional domain, time lag in sensing, etc.

948 **Acknowledgments**

949 This work is partially supported by the FormbaR project, part of AG Signalling/DB RailLab in  
 950 the Innovation Alliance of Deutsche Bahn AG and TU Darmstadt. This material is based upon  
 951 work supported by AFOSR grant FA9550-16-1-0288.

— **References** —

- 1 Ehsan Ahmad, Yunwei Dong, Shuling Wang, Naijun Zhan, and Liang Zou. Adding formal meanings to AADL with hybrid annex. In Ivan Lanese and Eric Madelaine, editors, *Formal Aspects of Component Software - 11th International Symposium, FACS 2014, Bertinoro, Italy, September 10-12, 2014, Revised Selected Papers*, volume 8997 of *Lecture Notes in Computer Science*, pages 228–247. Springer, 2014.
- 2 Ehsan Ahmad, Brian R. Larson, Stephen C. Barrett, Naijun Zhan, and Yunwei Dong. Hybrid annex: an AADL extension for continuous behavior and cyber-physical interaction modeling. In Michael Feldman and S. Tucker Taft, editors, *Proceedings of the 2014 ACM SIGAda annual conference on High integrity language technology, HILT 2014, Portland, Oregon, USA, October 18-21, 2014*, pages 29–38. ACM, 2014.
- 3 Elvira Albert, Frank S. de Boer, Reiner Hähnle, Einar Broch Johnsen, Rudolf Schlatte, Silvia Lizeth Tapia Tarifa, and Peter Y. H. Wong. Formal modeling and analysis of resource management for cloud architectures: an industrial case study using real-time ABS. *Service Oriented Computing and Applications*, 8(4):323–339, 2014.
- 4 M. Althoff. An introduction to CORA 2015. In *Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems*, 2015.
- 5 Rajeev Alur, Costas Courcoubetis, Thomas A. Henzinger, and Pei Hsin Ho. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In *Hybrid Systems, volume 736 of LNCS*, pages 209–229, Berlin, Heidelberg, 1993. Springer.
- 6 Rajeev Alur, Thao Dang, Joel M. Esposito, Rafael B. Fierro, Yerang Hur, Franjo Ivancic, Vijay Kumar, Insup Lee, Pradyumna Mishra, George J. Pappas, and Oleg Sokolsky. Hierarchical hybrid modeling of embedded systems. In Thomas A. Henzinger and Christoph M. Kirsch, editors, *Embedded Software, First International Workshop, EM-SOFT 2001, Tahoe City, CA, USA, October, 8-10, 2001, Proceedings*, volume 2211 of *Lecture Notes in Computer Science*, pages 14–31. Springer, 2001.
- 7 Rajeev Alur, Radu Grosu, Yerang Hur, Vijay Kumar, and Insup Lee. Modular specification of hybrid systems in CHARON. In Nancy A. Lynch and Bruce H. Krogh, editors, *Hybrid Systems: Computation and Control, Third International Workshop, HSCC 2000, Pittsburgh, PA, USA, March 23-25, 2000, Proceedings*, volume 1790 of *Lecture Notes in Computer Science*, pages 6–19. Springer, 2000.
- 8 Sven Apel, Don S. Batory, Christian Kästner, and Gunter Saake. *Feature-Oriented Software Product Lines - Concepts and Implementation*. Springer, 2013.
- 9 Paul C. Attie and Nancy A. Lynch. Dynamic input/output automata: A formal and compositional model for dynamic systems. *Inf. Comput.*, 249:28–75, 2016.
- 10 Stanley Bak, Omar Ali Beg, Sergiy Bogomolov, Taylor T. Johnson, Luan Viet Nguyen, and Christian Schilling. Hybrid automata: from verification to implementation. *STTT*, 21(1):87–104, 2019.
- 11 Stanley Bak, Sergiy Bogomolov, and Taylor T. Johnson. HYST: a source transformation and translation tool for hybrid automaton models. In Antoine Girard and Sriram Sankaranarayanan, editors, *HSCC’15*, pages 128–133. ACM, 2015.
- 12 Richard Banach, Michael J. Butler, Shengchao Qin, Nitika Verma, and Huibiao Zhu. Core hybrid Event-B I: single hybrid Event-B machines. *Sci. Comput. Program.*, 105:92–123, 2015.
- 13 Richard Banach, Michael J. Butler, Shengchao Qin, and Huibiao Zhu. Core hybrid Event-B II: multiple cooperating hybrid Event-B machines. *Sci. Comput. Program.*, 139:1–35, 2017.
- 14 Don S. Batory, Jacob Neal Sarvela, and Axel Rauschmayer. Scaling step-wise refinement. *IEEE Trans. Software Eng.*, 30(6):355–371, 2004.
- 15 Luca Benvenuti, Davide Bresolin, Pieter Collins, Alberto Ferrari, Luca Geretti, and Tiziano Villa. Assume-guarantee verification of nonlinear hybrid systems with Ariadne. *International Journal of Robust and Nonlinear Control*, 24(4):699–724, 2014.
- 16 Joakim Björk, Frank S. de Boer, Einar Broch Johnsen, Rudolf Schlatte, and Silvia Lizeth Tapia Tarifa. User-defined schedulers for real-time concurrent objects. *Innovations in Systems and Software Engineering*, 9(1):29–43, 2013.
- 17 Sergiy Bogomolov, Goran Frehse, Marius Greitschus, Radu Grosu, Corina S. Pasareanu, Andreas Podelski, and Thomas Strump. Assume-guarantee abstraction refinement meets hybrid systems. In Eran Yahav, editor, *Hardware and Software: Verification and Testing - 10th International Haifa Verification Conference, HVC 2014, Haifa, Israel, November 18-20, 2014. Proceedings*, volume 8855 of *Lecture Notes in Computer Science*, pages 116–131. Springer, 2014.
- 18 Brandon Bohrer, Yong Kiam Tan, Stefan Mitsch, Magnus O. Myreen, and André Platzer. VeriPhy: Verified controller executables from verified cyber-physical system models. In Dan Grossman, editor, *PLDI*, pages 617–630. ACM, 2018.
- 19 Christopher X. Brooks, Edward A. Lee, David Lorenzetti, Thierry S. Noudui, and Michael Wetter. CyPhySim: a cyber-physical systems simulator. In Antoine Girard and Sriram Sankaranarayanan, editors, *HSCC’15*, pages 301–302. ACM, 2015.

- 20 Lilian Burdy, Yoonsik Cheon, David R. Cok, Michael D. Ernst, Joseph R. Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer*, 7(3):212–232, 2005.
- 21 Michael J. Butler, Jean-Raymond Abrial, and Richard Banach. Modelling and refining hybrid systems in Event-B and Rodin. In Luigia Petre and Emil Sekerinski, editors, *From Action Systems to Distributed Systems - The Refinement Approach*, pages 29–42. Chapman and Hall/CRC, 2016.
- 22 Mingshuai Chen, Xiao Han, Tao Tang, Shuling Wang, Mengfei Yang, Naijun Zhan, Hengjun Zhao, and Liang Zou. MARS: A toolchain for modelling, analysis and verification of hybrid systems. In Michael G. Hinchey, Jonathan P. Bowen, and Ernst-Rüdiger Olderog, editors, *Provably Correct Systems, NASA Monographs in Systems and Software Engineering*, pages 39–58. Springer, 2017.
- 23 Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. Flow\*: An analyzer for non-linear hybrid systems. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 258–263. Springer, 2013.
- 24 Xin Chen and Sriram Sankaranarayanan. Decomposed reachability analysis for nonlinear systems. In *2016 IEEE Real-Time Systems Symposium, RTSS 2016, Porto, Portugal, November 29 - December 2, 2016*, pages 13–24. IEEE Computer Society, 2016.
- 25 Dave Clarke, Radu Muschevici, José Proença, Ina Schaefer, and Rudolf Schlatte. Variability modelling in the ABS language. In *FMCO*, volume 6957 of *Lecture Notes in Computer Science*, pages 204–224. Springer, 2010.
- 26 Fabio Cremona, Marten Lohstroh, David Broman, Edward A. Lee, Michael Masin, and Stavros Tripakis. Hybrid co-simulation: it’s about time. *Software and Systems Modeling*, 18(3):1655–1679, 2019.
- 27 P.J.L. Cuijpers and M.A. Reniers. Hybrid process algebra. *J. of Logic and Algebraic Programming*, 62(2):191–245, 2005.
- 28 Frank S. de Boer, Dave Clarke, and Einar Broch Johnsen. A complete guide to the future. In *ESOP*, volume 4421 of *Lecture Notes in Computer Science*, pages 316–330. Springer, 2007.
- 29 Frank S. de Boer, Vlad Serbanescu, Reiner Hähnle, Ludovic Henrio, Justine Rochas, Crystal Chang Din, Einar Broch Johnsen, Marjan Sirjani, Ehsan Khamespanah, Kiko Fernandez-Reyes, and Albert Mingkun Yang. A survey of active object languages. *ACM Computing Surveys*, 50(5):1–39, 2017.
- 30 Ankush Desai, Amar Phanishayee, Shaz Qadeer, and Sanjit A. Seshia. Compositional programming and testing of dynamic distributed systems. *Proc. ACM Program. Lang.*, 2(OOPSLA):159:1–159:30, 2018.
- 31 Crystal Chang Din, Reiner Hähnle, Einar Broch Johnsen, Ka I Pun, and Silvia Lizeth Tapia Tarifa. Locally abstract, globally concrete semantics of concurrent programming languages. In Renate A. Schmidt and Cláudia Nalon, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 26th International Conference, TABLEAUX 2017, Brasília, Brazil, September 25-28, 2017, Proceedings*, volume 10501 of *Lecture Notes in Computer Science*, pages 22–43. Springer, 2017.
- 32 Crystal Chang Din and Olaf Owe. Compositional reasoning about active objects with shared futures. *Formal Asp. Comput.*, 27(3):551–572, 2015.
- 33 Alexandre Donzé and Goran Frehse. Modular, hierarchical models of control systems in SpaceX. In *European Control Conference, ECC 2013, Zurich, Switzerland, July 17-19, 2013*, pages 4244–4251. IEEE, 2013.
- 34 Guillaume Dupont, Yamine Aït Ameur, Neeraj Kumar Singh, and Marc Pantel. Event-B hybridization: A proof and refinement-based framework for modelling hybrid systems. *ACM Trans. Embed. Comput. Syst.*, 20(4):35:1–35:37, 2021.
- 35 Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. SpaceX: Scalable verification of hybrid systems. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 379–395. Springer, 2011.
- 36 Nathan Fulton, Stefan Mitsch, Jan-David Quesel, Marcus Völp, and André Platzer. KeYmaera X: an axiomatic tactical theorem prover for hybrid systems. In Amy P. Felty and Aart Middeldorp, editors, *CADE’25*, volume 9195 of *LNCS*, pages 527–538. Springer, 2015.
- 37 Sicun Gao, Soonho Kong, and Edmund M. Clarke. dReal: An SMT solver for nonlinear theories over the reals. In Maria Paola Bonacina, editor, *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, volume 7898 of *Lecture Notes in Computer Science*, pages 208–214. Springer, 2013.
- 38 Luis Garcia, Stefan Mitsch, and André Platzer. Hy-PLC: hybrid programmable logic controller program translation for verification. In *ICCP’19*, pages 47–56, 2019.
- 39 Sergey Goncharov and Renato Neves. An adequate while-language for hybrid computation. *CoRR*, abs/1902.07684, 2019.
- 40 Sergey Goncharov, Renato Neves, and José Proença. Implementing hybrid semantics: From functional to imperative. In Violet Ka I Pun, Volker Stolz, and Adenilso Simão, editors, *Theoretical Aspects of Computing - ICTAC 2020 - 17th International Colloquium, Macau, China, November 30 - December 4, 2020, Proceedings*, volume 12545 of *Lecture Notes in Computer Science*, pages 262–282. Springer, 2020.
- 41 Daniel Grahl, Richard Bubel, Wojciech Mostowski, Peter H. Schmitt, Mattias Ulbrich, and Benjamin Weiß. Modular specification and verification. In Wolfgang Ahrendt, Bernhard Beckert, Richard

- Bubel, Reiner Hähnle, Peter H. Schmitt, and Matthias Ulbrich, editors, *Deductive Software Verification - The KeY Book - From Theory to Practice*, volume 10001 of *Lecture Notes in Computer Science*, pages 289–351. Springer, 2016.
- 42 Ichiro Hasuo and Kohei Suenaga. Exercises in nonstandard static analysis of hybrid systems. In P. Madhusudan and Sanjit A. Seshia, editors, *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*, volume 7358 of *Lecture Notes in Computer Science*, pages 462–478. Springer, 2012.
  - 43 Thomas A. Henzinger, Marius Minea, and Vinayak S. Prabhu. Assume-guarantee reasoning for hierarchical hybrid systems. In Maria Domenica Di Benedetto and Alberto L. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control, 4th International Workshop, HSCC 2001, Rome, Italy, March 28-30, 2001, Proceedings*, volume 2034 of *Lecture Notes in Computer Science*, pages 275–290. Springer, 2001.
  - 44 Carl Hewitt, Peter Bishop, and Richard Steiger. A universal modular ACTOR formalism for artificial intelligence. In *IJCAI'73*, pages 235–245, 1973.
  - 45 Jochen Hoenicke and Ernst-Rüdiger Olderog. Combining specification techniques for processes, data and time. In Michael Butler, Luigia Petre, and Kaisa Sere, editors, *Integrated Formal Methods*, pages 245–266, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
  - 46 Iman Jahandideh, Fatemeh Ghassemi, and Marjan Sirjani. Hybrid Rebeca: modeling and analyzing of cyber-physical systems. *CoRR*, abs/1901.02597, 2019.
  - 47 Jean-Baptiste Jeannin, Khalil Ghorbal, Yanni Kouskoulas, Aurora Schmidt, Ryan Gardner, Stefan Mitsch, and André Platzer. A formally verified hybrid system for safe advisories in the next-generation airborne collision avoidance system. *STTT*, 19(6):717–741, 2017.
  - 48 Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In Bernhard K. Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *FMCO'10*, volume 6957 of *LNCS*, pages 142–164. Springer, 2010.
  - 49 Einar Broch Johnsen, Ka I Pun, and Silvia Lizeth Tapia Tarifa. A formal model of cloud-deployed software and its application to workflow processing. In Dinko Begusic, Nikola Rozic, Josko Radic, and Matko Saric, editors, *SoftCOM'17*, pages 1–6. IEEE, 2017.
  - 50 Eduard Kamburjan. Detecting deadlocks in formal system models with condition synchronization. *ECEASST*, 76, 2018.
  - 51 Eduard Kamburjan. From post-conditions to post-region invariants: Deductive verification of hybrid objects. In *HSCC*. ACM, 2021.
  - 52 Eduard Kamburjan, Crystal Chang Din, Reiner Hähnle, and Einar Broch Johnsen. Behavioral contracts for cooperative scheduling. In *20 Years of KeY*, volume 12345 of *Lecture Notes in Computer Science*, pages 85–121. Springer, 2020.
  - 53 Eduard Kamburjan, Reiner Hähnle, and Sebastian Schön. Formal modeling and analysis of railway operations with Active Objects. *Science of Computer Programming*, 166:167–193, November 2018.
  - 54 Eduard Kamburjan, Rudolf Schlatte, Einar Broch Johnsen, and S. Lizeth Tapia Tarifa. Designing distributed control with hybrid active objects. In *ISoLA*, volume 12479 of *LNCS*. Springer, 2020.
  - 55 Soonho Kong, Sicun Gao, Wei Chen, and Edmund M. Clarke. dReach:  $\delta$ -reachability analysis for hybrid systems. In Christel Baier and Cesare Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9035 of *Lecture Notes in Computer Science*, pages 200–205. Springer, 2015.
  - 56 Jia-Chun Lin, Ingrid Chieh Yu, Einar Broch Johnsen, and Ming-Chang Lee. ABS-YARN: A formal framework for modeling hadoop YARN clusters. In Perdita Stevens and Andrzej Wasowski, editors, *FASE'16*, volume 9633 of *LNCS*, pages 49–65. Springer, 2016.
  - 57 Jiang Liu, Jidong Lv, Zhao Quan, Najjun Zhan, Hengjun Zhao, Chaochen Zhou, and Liang Zou. A calculus for hybrid CSP. In Kazunori Ueda, editor, *Programming Languages and Systems - 8th Asian Symposium, APLAS 2010, Shanghai, China, November 28 - December 1, 2010. Proceedings*, volume 6461 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2010.
  - 58 Simon Lunel, Stefan Mitsch, Benoît Boyer, and Jean-Pierre Talpin. Parallel composition and modular verification of computer controlled systems in differential dynamic logic. In Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira, editors, *Formal Methods, The Next 30 Years, Third World Congress, FM, Porto, Portugal*, volume 11800 of *LNCS*, pages 354–370. Springer, 2019.
  - 59 Nancy A. Lynch, Roberto Segala, and Frits W. Vaandrager. Hybrid I/O automata. *Inf. Comput.*, 185(1):105–157, 2003.
  - 60 Brendan P. Mahony and Jin Song Dong. Deep semantic links of TCSP and Object-Z: TCOZ approach. *Formal Aspects Comput.*, 13(2):142–160, 2002.
  - 61 *Maxima Manual*, 5.43.0 edition, 2019. [maxima.sourceforge.net](http://maxima.sourceforge.net).
  - 62 Stefan Mitsch, Grant Olney Passmore, and André Platzer. Collaborative verification-driven engineering of hybrid systems. *Math. Comput. Sci.*, 8(1):71–97, 2014.
  - 63 Stefan Mitsch and André Platzer. ModelPlex: Verified runtime validation of verified cyber-physical system models. *Form. Methods Syst. Des.*, 49(1):33–74, 2016.
  - 64 Andreas Müller, Stefan Mitsch, Werner Retschitzegger, Wieland Schwinger, and André Platzer. Tactical contract composition for hybrid system component verification. *STTT*, 20(6):615–643, 2018. Special issue for selected papers from FASE'17.
  - 65 André Platzer. Differential-algebraic dynamic logic for differential-algebraic programs. *J. of Logic and Computation*, 20(1):309–352, 2010.

- 66 André Platzer. The complete proof theory of hybrid systems. In *LICS*, pages 541–550. IEEE, 2012.
- 67 André Platzer. The structure of differential invariants and differential cut elimination. *Logical Methods in Computer Science*, 8(4):1–38, 2012.
- 68 André Platzer. A complete uniform substitution calculus for differential dynamic logic. *J. Automated Reasoning*, 59(2):219–265, 2017.
- 69 André Platzer. *Logical Foundations of Cyber-Physical Systems*. Springer, 2018.
- 70 André Platzer and Yong Kiam Tan. Differential equation invariance axiomatization. *J. ACM*, 67(1):6:1–6:66, 2020.
- 71 Claudius Ptolemaeus, editor. *System Design, Modeling, and Simulation using Ptolemy II*. Ptolemy.org, 2014.
- 72 Jan-David Quesel, Stefan Mitsch, Sarah Loos, Nikos Aréchiga, and André Platzer. How to model and prove hybrid systems with KeYmaera: A tutorial on safety. *STTT*, 18(1):67–91, 2016.
- 73 Ina Schaefer, Lorenzo Bettini, Viviana Bono, Ferruccio Damiani, and Nico Tanzarella. Delta-oriented programming of software product lines. In *SPLC*, volume 6287 of *Lecture Notes in Computer Science*, pages 77–91. Springer, 2010.
- 74 Ina Schaefer, Rick Rabiser, Dave Clarke, Lorenzo Bettini, David Benavides, Goetz Botterweck, Animesh Pathak, Salvador Trujillo, and Karina Villela. Software diversity: state of the art and perspectives. *STTT*, 14(5):477–495, 2012.
- 75 Rudolf Schlatte, Einar Broch Johnsen, Jacopo Mauro, Silvia Lizeth Tapia Tarifa, and Ingrid Chieh Yu. Release the beasts: When formal methods meet real world data. In *It's All About Coordination*, volume 10865 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2018.
- 76 Wen Su, Jean-Raymond Abrial, and Huibiao Zhu. Formalizing hybrid systems with Event-B and the Rodin platform. *Sci. Comput. Program.*, 94:164–202, 2014.
- 77 Kohei Suenaga and Ichiro Hasuo. Programming with infinitesimals: A while-language for hybrid system modeling. In *ICALP (2)*, volume 6756 of *Lecture Notes in Computer Science*, pages 392–403. Springer, 2011.
- 78 Casper Thule, Kenneth Lausdahl, Cláudio Gomes, Gerd Meisl, and Peter Gorm Larsen. Maestro: The INTO-CPS co-simulation framework. *Simulation Modelling Practice and Theory*, 92:45–61, 2019.
- 79 D. A. van Beek, Michel A. Reniers, Ramon R. H. Schiffelers, and J. E. Rooda. Foundations of a compositional interchange format for hybrid systems. In Alberto Bemporad, Antonio Bicchi, and Giorgio C. Buttazzo, editors, *HSCC'07*, volume 4416 of *LNCS*, pages 587–600. Springer, 2007.
- 80 Shuling Wang, Naijun Zhan, and Liang Zou. An improved HHL prover: An interactive theorem prover for hybrid systems. In Michael Butler, Sylvain Conchon, and Fatiha Zaïdi, editors, *Formal Methods and Software Engineering*, pages 382–399. Cham, 2015. Springer International Publishing.
- 81 Peter Y. H. Wong, Elvira Albert, Radu Muschevici, José Proença, Jan Schäfer, and Rudolf Schlatte. The ABS tool suite: modelling, executing and analysing distributed adaptable object-oriented systems. *STTT*, 14(5):567–588, 2012.
- 82 Zhenkai Zhang, Emeka Eyisi, Xenofon D. Koutsoukos, Joseph Porter, Gabor Karsai, and Janos Sztipanovits. A co-simulation framework for design of time-triggered automotive cyber physical systems. *Simulation Modelling Practice and Theory*, 43:16–33, 2014.