



# ARCH-COMP26 Category Report: Hybrid Systems Theorem Proving

Stefan Mitsch<sup>1</sup>, Xiangyu Jin<sup>2</sup>, Shuling Wang<sup>3</sup>, and Bohua Zhan<sup>4</sup>

<sup>1</sup> School of Computing, DePaul University, Chicago, IL, USA  
[smitsch@depaul.edu](mailto:smitsch@depaul.edu)

<sup>2</sup> National Key Laboratory of Space Integrated Information System, Institute of Software, Chinese Academy of Sciences, Beijing, China  
[jinxy@ios.ac.cn](mailto:jinxy@ios.ac.cn)

<sup>3</sup> National Key Laboratory of Space Integrated Information System, Institute of Software, Chinese Academy of Sciences, Beijing, China  
[wangsl@ios.ac.cn](mailto:wangsl@ios.ac.cn)

<sup>4</sup> Huawei Technologies Co., Ltd, Beijing, China  
[zhanbohua@huawei.com](mailto:zhanbohua@huawei.com)

## Abstract

This paper reports on the Hybrid Systems Theorem Proving (*HSTP*) category in the ARCH-COMP Friendly Competition 2026. *HSTP* focuses on flexibility of programming languages as structuring principles for hybrid systems, unambiguity and precision of program semantics, and mathematical rigor of logical reasoning principles. The benchmark set includes nonlinear and parametric continuous and hybrid systems and hybrid games, each in three modes: fully automatic verification, semi-automatic verification from proof hints, proof checking from scripted tactics. This instance of the competition focuses on presenting the differences between the provers on a subset of the benchmark examples.

## 1 Introduction

This report summarizes the experimental results of the Hybrid Systems Theorem Proving (*HSTP*) category in the ARCH-COMP26 friendly competition, focusing on a feature comparison between the participating theorem provers. Details on the benchmark sets and the evaluation modes can be found in previous editions of the *HSTP* category [[MMJ<sup>+</sup>20](#), [MJZ<sup>+</sup>21](#), [MZS<sup>+</sup>22](#), [MSZ<sup>+</sup>23](#)]. The examples in the benchmark competition are grouped into the following categories:

- Hybrid systems design shapes: small-scale examples over a large variety of model shapes to test for prover flexibility.
- Nonlinear continuous models: test for prover flexibility in terms of generating and proving properties about continuous dynamics, based on [[SMT<sup>+</sup>19](#), [SMT<sup>+</sup>21](#)].
- Hybrid games: small-scale examples with adversary dynamics in differential dynamic game logic.

- Hybrid systems case studies: hybrid systems models and specifications at scale to test for application scalability and efficiency, based on [MGVP17].
- Hybrid systems from Simulink/Stateflow models: examples translated from models expressed in Simulink/Stateflow.

In each of these categories, tools can select the degree of automation depending on their focus in the spectrum from fast proof checking to full proof automation:

- (A) Automated: hybrid systems models and specifications are the only input, proofs and counterexamples are produced fully automatically.
- (H) Hints: select proof hints (e.g., loop invariants) are provided as part of the specifications.
- (S) Scripted: significant parts of the verification is done with dedicated problem-specific scripts or tactics.

Benchmark examples in the hybrid systems design shapes, nonlinear continuous models, hybrid games and hybrid systems case study benchmarks are available at <https://github.com/LS-Lab/KeYmaeraX-projects/tree/master/benchmarks> and specified in differential dynamic logic (dL) [Pla08, Pla17]. Benchmark examples for HHLPy, including the Simulink/Stateflow models and their translations to Hybrid CSP [ZWR95], are available at <https://gitee.com/bhzhhan/mars/tree/master/hhlpy/examples/simulink>. An introduction to the problem format syntax is in Section 2. The participating tools are presented in Section 3. An overview of the examples together with the findings from the competition is given in Section 4.

## 2 Problem Format

Benchmarks in the hybrid systems design shapes, nonlinear continuous models, hybrid games and hybrid systems case study categories are written in differential dynamic logic (dL) [Pla08, Pla17] which has axioms and an unambiguous semantics available [BRV<sup>+</sup>17] in KeYmaera 3, KeYmaera X, Isabelle/HOL, and Coq. A tutorial on the modeling principles in dL can be found in [QML<sup>+</sup>16], details on the ASCII syntax are in [MMJ<sup>+</sup>20]. Libraries of pre-defined functions (e.g., `import kyx.math.abs`) [GTMP22] help expressing models more conveniently. Benchmarks in the hybrid systems design shapes and nonlinear continuous models are also translated to the HHLPy input language, along with the Simulink/Stateflow benchmarks. In the second subsection, we describe the input language for HHLPy in the competition.

**Problem Format Example.** The KeYmaera X ASCII syntax is illustrated in the example below, with tactics using position identifiers to refer to formulas and terms in a sequent.

```

1  ArchiveEntry "Benchmark Example 1"
2
3  Definitions                                /* definitions cannot change their value */
4  import kyx.math.abs;                       /* import absolute value function */
5  Real A = 5;                                /* real-valued maximum acceleration defined to be 5 */
6  Real b;                                     /* real-valued braking, undefined so unknown value */
7  Bool geq(Real x, Real y) <-> x>=y;        /* predicate geq defined to be formula x>=y */
8  HP drive ::= {                               /* program drive defined to choose either */
9      ?v<=5; a:=A;                             /* maximum acceleration if slow enough */
10     ++ a:=-b;                                /* or braking, nondeterministically */
11 };
12 End.
13
14 ProgramVariables                            /* program variables may change their value over time */
15 Real x;                                     /* real-valued position */
16 Real v;                                     /* real-valued velocity */
17 Real a;                                     /* current acceleration chosen by controller */
18 End.
19
20 Problem                                     /* conjecture in differential dynamic logic */

```

```

21 | v>=0 & b>0          /* initial condition */
22 | ->                  /* implies */
23 | [                   /* all runs of this hybrid program */
24 | {                   /* braces {} group programs */
25 |   drive;            /* expand program drive here as defined above */
26 |   { x'=v, v'=a & v>=0 } /* differential equation system */
27 | } * @invariant(v>=0) /* loop repeats, with @invariant contract */
28 | ] v>=0              /* safety/postcondition after hybrid program */
29 | End.
30 |
31 | Tactic "Automated proof in KeYmaera X"
32 | auto
33 | End.
34 |
35 | Tactic "Scripted proof in extended Bellerophon tactic language"
36 | implyR('R=="[?v<=5;a:=5;+a:=-b();]{x'=v,v'=a&v>=0}]v>=0*");
37 | loop(*v>=0*, 'R=="[?v<=5;a:=5;+a:=-b();]{x'=v,v'=a&v>=0}]v>=0*"); < ( /* < splits branches */
38 | "Init":
39 |   id, /* initial case: shown with close by identity */
40 | "Post":
41 |   QE, /* postcondition: prove by real arithmetic QE */
42 | "Step":
43 |   compose('R=="[?v<=5;a:=5;+a:=-b();]{x'=v,v'=a&v>=0}]v>=0*");
44 |   solve('R=="[?v<=5;a:=5;+a:=-b();]#[{x'=v,v'=a&v>=0}]v>=0*");
45 |   choiceb('R=="[?v<=5;a:=5;+a:=-b();]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v
46 |     \hookrightarrow >=0)->a*t_+v>=0)");
47 |   /* separate controller branches */
48 |   andR('R=="[?v<=5;a:=5;]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*t_+v
49 |     \hookrightarrow >=0)&[a:=-b();]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*t_+v
50 |     \hookrightarrow >=0)"); < (
51 |     "[?v<=5;a:=5;]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*t_+v>=0)*":
52 |     /* decompose some steps then ask auto */
53 |     compose('R=="[?v<=5;a:=5;]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->
54 |       \hookrightarrow a*t_+v>=0)");
55 |     testb('R=="[?v<=5;][a:=5;]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*
56 |       \hookrightarrow t_+v>=0)");
57 |     auto,
58 |     "[a:=-b();]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*t_+v>=0)*":
59 |     /* assignment, then real arithmetic */
60 |     assignb('R=="[a:=-b();]\forall t_ (t_>=0->\forall s_ (0<=s_&s_<=t_->a*s_+v>=0)->a*t_
61 |       \hookrightarrow +v>=0)");
62 |     QE
63 |   )
64 | )
65 | End.
66 | End. /* end of ArchiveEntry */

```

**Input Language for HHLPy** Benchmarks in the Hybrid Systems from Simulink/Stateflow category are modeled using Hybrid CSP, with properties specified as Hoare triples in Hybrid Hoare Logic. Both Hybrid CSP programs and properties as Hoare triples can be written using an ASCII syntax, as illustrated below. They are composed of pre-conditions, programs and post-conditions. The program is annotated with invariants and rules for proof.

```

1 | # ArchiveEntry Benchmark Example 2
2 |
3 | pre; /* pre-condition */
4 | t := 0; /* Assignment command */
5 | x := 0;
6 | {
7 |   Chart_A_done := 0;
8 |   if (t >= 1) { /* If command */
9 |     t := 0;
10 |    x := 0;
11 |    Chart_A_done := 1;
12 |   }
13 |   Chart_ret := Chart_A_done;
14 |   {x_dot = 1, t_dot = 1 & t < 1} /* differential equation systems */

```

```

15 |   invariant [x == t]{di};           # differential equation invariant, with proof rule
16 | }*                                # loop repeats
17 |   invariant [x == t] [0 <= x] [x <= 1]; # loop invariant
18 | post [0 <= x] [x <= 1];          # post-condition

```

The pre-condition is true and the post-condition is  $0 \leq x \wedge x \leq 1$  in the above example. In the program,  $t$  and  $x$  are assigned as 0, followed by a loop command. The invariants of the loop are  $x == t$ ,  $0 \leq x$  and  $x \leq 1$ . A differential equation is in the loop with invariant  $x == t$  proved by the rule dI (differential invariant).

### 3 Participating Tools

**KeYmaera X.** KeYmaera X [FMQ<sup>+</sup>15] is a theorem prover for the hybrid systems logic differential dynamic logic (dL). It implements the uniform substitution calculus of dL [Pla17]. A comparison of the internal reasoning principles in the KeYmaera family of provers with a discussion of their relative benefits and drawbacks is in [MP20], and model structuring and proof management on top of uniform substitution is discussed in [Mit21]. KeYmaera X supports systems with nondeterministic discrete jumps, nonlinear differential equations, nondeterministic inputs, and allows defining functions implicitly through their characterizing differential equations [GTMP22]. It provides invariant construction and proving techniques for differential equations [SMT<sup>+</sup>21, PT20], and stability verification techniques for switched systems [TMP22]. To discharge proof obligations in real arithmetic, KeYmaera X interacts with trusted backend procedures for quantifier elimination (Z3, Wolfram Mathematica, Wolfram Engine); a verified backend procedure based on virtual term substitution is under development [SCMP21]. Proofs in KeYmaera X can be conducted interactively [MP16], steered with tactics [FMBP17], or attempted fully automatic. KeYmaera X provides a definitions mechanism for functions [GTMP22]. This year’s tool improvements include porting to Scala 3, and an improved benchmarking setup that allows per-example backend selection and runs both locally and via Docker.

**HHL Prover/HHLPy.** HHL Prover is a verification tool for hybrid systems modeled by Hybrid CSP (HCSP) [He94, ZWR96], implemented in Isabelle/HOL. HCSP is an extension of CSP by introducing differential equations for modeling continuous evolution and interrupts for modeling interaction between continuous and discrete dynamics. The proof system of HHL Prover is Hybrid Hoare Logic (HHL) [LLQ<sup>+</sup>10].

HHLPy is a new verification tool for HCSP, that provides a friendlier web-based user interface. Currently it handles only the sequential fragment of HCSP, with reasoning rules similar to those in dL. We briefly introduce each of the two tools in the following paragraphs.

**HHL Prover** HHL Prover [WZZ15] is an interactive theorem prover for verifying hybrid systems modeled by Hybrid CSP (HCSP). We use the trace-based hybrid Hoare logic for reasoning about HCSP processes as in last year. Traces for both sequential and parallel HCSP processes are represented as lists of *trace blocks*. There are two types of trace blocks: ODE blocks and communication blocks. ODE blocks specify evolution of the process over an interval of time, consisting of duration of the interval, the state of the process as a function of time, and a set of communications that are ready during the interval. Communication blocks are of three types: input, output, and IO. Input and output blocks specify an unmatched communication event, while IO blocks specify a matched communication event. All three types of events also specify the value that is communicated. The input and output blocks are synchronized during parallel composition of HCSP processes.

**HHLPy** HHLPy is a theorem prover with a friendlier user interface, currently for verifying sequential HCSP programs only. It is implemented using Python and JavaScript. The sequential fragment of HCSP contains ODEs with domain boundary, but not communication, interrupts, and parallel processes. Extending HHLPy to handle the full HCSP language is left for future work. Given a sequential HCSP process  $P$ , a specification takes the form of Hoare triple,  $\{Pre\}P\{Post\}$ , where  $Pre$  and  $Post$  are pre-/post-conditions in first-order logic.

To reason about differential equations, HHLPy makes use of a set of proof rules that are inspired by  $d\mathcal{L}$  [Pla10, Pla11, PT18], but adapted to the semantics of sequential HCSP. The differential weakening (dW) rule reduces a Hoare triple concerning ODEs to an invariant triple of the ODE and some verification conditions. Invariant triple is in the form of  $\llbracket P \rrbracket \langle \dot{x} = e \rangle \llbracket Q \rrbracket$ , whose semantics is roughly stated as follows: for any solution to the differential equation  $\dot{x} = e$ , if  $Q$  is satisfied at beginning and  $P$  is satisfied throughout, then  $Q$  is satisfied throughout. Rules such as differential invariant (dI), differential cut (dC), Darboux’s rule (dbx) and barrier certificates (barrier), many of which borrowed from differential dynamic logic, are then used to prove invariant triples.

HHLPy stores proof information after the corresponding assertion (post-condition, invariant), so that the user can still reuse proofs when they modify the program or the assertions slightly. Specifically, proof rules are stored after the corresponding invariants, and proof methods for proving verification conditions, for example, Z3 or Wolfram Engine, are stored after the assertion that generates the corresponding verification condition. Sometimes one assertion corresponds to several verification conditions. HHLPy also proposed a labeling system to distinguish these verification conditions.

## 4 Benchmarks

One of the strengths of hybrid systems theorem proving as a verification technique is its support for combined automated and interactive verification steps as well as its applicability to proof search and proof checking. The benchmark examples were analyzed in three modes:

**Automated** The specification is the only input to the theorem prover. Proofs and counterexamples are obtained fully automated to highlight the capabilities of theorem provers in terms of invariant generation, proof search, and proof checking.

**Hints** Known design properties of the system, such as loop invariants and invariants of differential equations, are annotated in the model and allowed to be exploited during an otherwise fully automated proof to highlight the capabilities of theorem provers in terms of proof search and proof checking.

**Scripted** User guidance with proof scripts is allowed to highlight the capabilities of theorem provers in terms of proof checking.

The benchmark examples are structured into 5 categories: hybrid systems design shape examples to test for system design variations at a small scale, nonlinear continuous models to test for continuous invariant construction and proving capabilities, hybrid game examples to test adversarial dynamics, hybrid systems case studies to test for prover scalability, and a category for hybrid systems from Simulink/Stateflow models.

**Experimental Setup.** HHLPy participated in three benchmark sets, which are hybrid systems design shapes, nonlinear continuous models and hybrid systems from Simulink/Stateflow

models. The performance results were obtained on Windows 11, with Z3-solver 4.8.12.0 and Wolfram Engine 13.0, on the machine with 8-core Intel(R) Core(TM) i5-1035G4 CPU @ 1.10GHz and 16 GB memory.

HHL Prover participated in hybrid systems design shapes. The results were obtained on Windows 10, with Isabelle2020 and afp-2020-12-22 on a machine with Intel(R) Core(TM) i7-8550U CPU @ 1.80GHz.

In this edition, KeYmaera X participated in all three modes in the design shapes, games, nonlinear continuous models, and case study benchmarks categories. The performance results reported here are obtained with Z3 4.15.4 and Wolfram Engine 14.3.0 as a backend on an Apple M2 Pro 32GB.

## 4.1 Hybrid Systems Design Shapes

This category is designed to test for basic verification features on simple examples. The benchmark examples are grouped as follows:

**Static semantics correctness** 9 examples with various sequential orders and nested structures of assignments, differential equations, and loops.

**Dynamics** 30 examples with differential equations ranging from solvable to nonlinear.

**LICS Tutorial 9 dL** tutorial examples [Pla12] ranging from basic time-triggered motion control to model-predictive control.

**STTT Tutorial 12 dL** modeling tutorial examples [QML<sup>+</sup>16] ranging from basic discrete event-triggered and time-triggered control for straight-line motion to speed control with a trajectory generator and lane-keeping with two-dimensional curved motion.

**KeYmaera X** KeYmaera X participated in scripted, hints, and automated mode (proof attempts were aborted after 60s, every proof attempt was made in a fresh prover instance with all caches cleared):

Table 1: KeYmaera X results in the hybrid systems design shapes category

	Examples	2025			2026		
		Proved	Average [s]	Max [s]	Proved	Average [s]	Max [s]
<b>Script</b>		60	0.67	5.3	60	0.86	4.7
<b>Hint</b>	61	52	0.54	1.8	49	0.59	2.1
<b>Auto</b>		50	0.44	2.1	50	0.63	2.2

Porting to Scala 3 largely maintained the prover performance. As in previous competitions ([MZS<sup>+</sup>22]) the remaining unsolved example is a progress proof.

**HHL Prover/HHLPy.** The HHL Prover successfully proved 49 of the 61 examples in Isabelle/HOL using our proof system. Since the benchmarks are originally formulated in terms of dynamic logic, some modifications are made to adapt it to a Hoare-logic style system.

The level of automation of HHLPy is between that of hints mode and scripted mode, requiring the users to annotate both the invariants and the rules of differential equations. Proof attempts were aborted after 300s. HHLPy verified 50 out of 61 examples in this category. For the 11 unverified examples, eight of them could not be translated into Hoare triples of HCSP programs, due to the semantics difference between dL and Hoare triples of HCSP programs; one of them is

non-polynomial; and we are still not clear about how to prove the last two. All of the verification conditions generated of the verified ones could be proved by Z3.

## 4.2 Nonlinear Continuous Models

The examples in this category remained unchanged from [MMJ<sup>+</sup>20] for direct comparison of the verification performance with previous results; the examples test for pure continuous verification performance. Future competitions may additionally utilize the extended benchmark set of [SMT<sup>+</sup>21].

**KeYmaera X.** KeYmaera X participated in the scripted, hints, and automated format (proof attempts were aborted after 60s, every proof attempt was made in a fresh prover instance with all caches cleared):

Table 2: KeYmaera X results in the nonlinear continuous category

	Examples	2025			2026		
		Proved	Average [s]	Max [s]	Proved	Average [s]	Max [s]
<b>Script</b>		107	1.2	36.5	101	1.1	14.2
<b>Hint</b>	141	96	1.2	19.6	<b>56</b>	1.6	35.1
<b>Auto</b>		59	1.7	10.5	<b>85</b>	1.7	34.9

The drop in the hints category is due to a change in the command-line interface of KeYmaera X: full automation is supported as a primary tactic, but no longer as a fallback tactic when a user-provided tactic fails. Full automation, by itself, improved significantly over last year’s results.

**HHLPy.** The level of automation of HHLPy is between that of hints mode and scripted mode, requiring the users to annotate both the invariants and the rules of differential equations. Proof attempts were aborted after 300s. HHLPy verified 103 out of 141 examples in this category. For most of the unverified ones, we have not found appropriate invariants. Most of the generated verification conditions could be proved both by Z3 and Wolfram Engine, while some could only be proved by Z3 or Wolfram Engine. We found that Z3 has advantages of handling complex boolean expressions, while Wolfram Engine is better at handling decimals and quantifiers.

## 4.3 Hybrid Games

The hybrid games benchmark tests basic games reasoning over 3 examples with adversarial dynamics. Future editions of the competition may utilize extended games case studies, such as [CMP23].

**KeYmaera X.** KeYmaera X participated in the scripted, hints, and automated format (proof attempts were aborted after 60s, every proof attempt was made in a fresh prover instance with all caches cleared):

Porting to Scala 3 largely maintained prover performance.

Table 3: KeYmaera X results in the hybrid games category

	Examples	2025			2026		
		Proved	Average [s]	Max [s]	Proved	Average [s]	Max [s]
<b>Script</b>		3	0.5	1.2	3	0.4	0.8
<b>Hint</b>	3	2	0.3	0.5	2	0.35	0.61
<b>Auto</b>		2	0.26	0.46	2	0.36	0.65

#### 4.4 Hybrid Systems Case Study Benchmarks

**Category overview.** The benchmark examples in this category are selected to test theorem provers for scalability and efficiency on examples of a significant size and interest in applications and remained unchanged from [MST<sup>+</sup>19]. The benchmark examples<sup>1</sup> are inspired from prior case studies on train control [PQ09, ZLW<sup>+</sup>13], flight collision avoidance [PC09], robot collision avoidance [MGVP17], a lunar lander descent guidance protocol [ZYZ<sup>+</sup>14], and rollercoaster safety [BLCP18].

**KeYmaera X.** KeYmaera X participated in the scripted, hints, and automated format (proof attempts were aborted after 300s, every proof attempt was made in a fresh prover instance with all caches cleared), and attempted 8 examples (3 ETCS train control, 3 flight collision avoidance, 2 robot collision avoidance).

Table 4: KeYmaera X results in the hybrid systems case study category

	Examples	2025			2026		
		Proved	Average [s]	Max [s]	Proved	Average [s]	Max [s]
<b>Script</b>		8	7.5	30.6	8	9.2	31.2
<b>Hint</b>	8	6	1.8	2.9	5	2.8	4.1
<b>Auto</b>		5	1.8	3.4	5	2.8	4.1

The difference in average/maximum duration between hints and scripted format is due to the three additional solved examples, which are more complex arithmetically and result in longer checking times in the external arithmetic solver.

#### 4.5 Hybrid Systems from Simulink/Stateflow Models

**Category overview.** This category contains hybrid systems modeled in Simulink/Stateflow. These models are first translated into the modeling language used for verification, and then its properties are verified using the appropriate tools. For now, only 5 benchmark problems are included, which illustrates the basic semantics of Simulink and Stateflow<sup>2</sup>. They include Stateflow charts with one or two states, ODEs within each state, delay blocks in Simulink diagrams, and a cruise control system modeled by Simulink diagrams.

<sup>1</sup><https://github.com/LS-Lab/KeYmaeraX-projects/blob/master/benchmarks/advanced.kyx>

<sup>2</sup><https://gitee.com/bhzhhan/mars/tree/master/hhlpy/examples/simulink>

**HHLPy.** HHLPy successfully verified all 5 examples. All of the generated verification conditions were proved by Z3. The Simulink/Stateflow diagrams were translated into HCSP programs automatically using the methods in [XZW<sup>+</sup>23, GZX<sup>+</sup>22]. The resulting HCSP programs were annotated manually with pre- and post-conditions specifying the desired properties, invariants and necessary proof rules for differential equations. The annotated Hoare triples were then verified automatically by HHLPy. For example, the Simulink diagram of the cruise control system consisted of subsystems for the PI controller and the vehicle. The translation process combined the controller and vehicle dynamics into a single differential equation. The initial and desired values for speed and control signal were annotated as pre- and post-conditions. The invariants were derived following the standard theory for analyzing linear dynamical systems. HHLPy generated 7 verification conditions for this example and all of them were proved by Z3.

As a more complex case study, we consider the Water Tank Control System from the official Simulink documentation [The25]. This hybrid system adjusts tank water levels through dynamic intake valve switching based on real-time sensor feedback. Its continuous plant behavior is governed by an ODE parameterised by the valve control variable, while the discrete controller implements dual-threshold regulation with timing and edge-triggered mechanisms.

A key challenge in verifying this system is that the monolithic parameterised ODE does not admit a unified differential invariant over the entire state space. To overcome this, the ODE is decomposed into three distinct evolution modes, each corresponding to a different level variation rate (the valve opening state, low-rate closing state and high-rate closing state). Each mode is augmented with customized differential invariants, while global invariants and boundary constraints are supplemented to bound the continuous evolution of the water level variable. With these invariants annotated, HHLPy produces over 690 proof obligations, all of which are discharged fully automatically by HHLPy’s backend solvers. The formal proof rigorously verifies that the water level remains safely within [190,1100] throughout system execution. Below is a snippet of the ODE part for this example in HHLPy.

```

1  pre [t==0][tt==0][_tick==0][pre_x1==0][pre_x2==0][x6==2]
2  [x11==500][Unit_Delay_state == 0][Check_if_Unit_Delay_state == 0];
3  {
4  .....
5  if(x6==20){
6    {tt_dot = 1, x11_dot = 10 & tt < 1}
7    invariant [pre_x1>=0 && pre_x2>=0]
8              [x11 <= 1090 + 10*tt]
9              [x11 >= 190];
10 }else if(x6==0){
11   {tt_dot = 1, x11_dot = -10 & tt < 1}
12   invariant [pre_x1>=0 && pre_x2>=0]
13             [x11 <= 1100]
14             [x11 >= 200 -10*tt];
15 }else if(x6==2){
16   {tt_dot = 1, x11_dot = -8 & tt < 1}
17   invariant [pre_x1>=0 && pre_x2>=0]
18             [x11 <= 1100]
19             [x11 >= 200 -8*tt];
20 }
21   t := t + tt;
22   _tick := _tick + 1;
23 }
24 invariant [pre_x1>=0 && pre_x2>=0]
25           [x11>=190][x11<=1100];
26 post[x11>=190][x11<=1100];

```

## 4.6 Scheduler

This case study is adapted from [XWZ<sup>+</sup>22], which describes a scheduler controlling two task modules under a preemptive priority-based scheduling policy. At any given time, at most one task is allowed to execute, and the scheduler always selects the highest-priority ready task for execution. Each task module consists of three states: WAIT, READY, and RUN. In the READY state, the task is prepared for execution and waits for permission from the scheduler. In the RUN state, the task is executing and may be preempted by the scheduler whenever a higher-priority task becomes available. In the WAIT state, the task has completed its current execution and waits for the beginning of the next period.

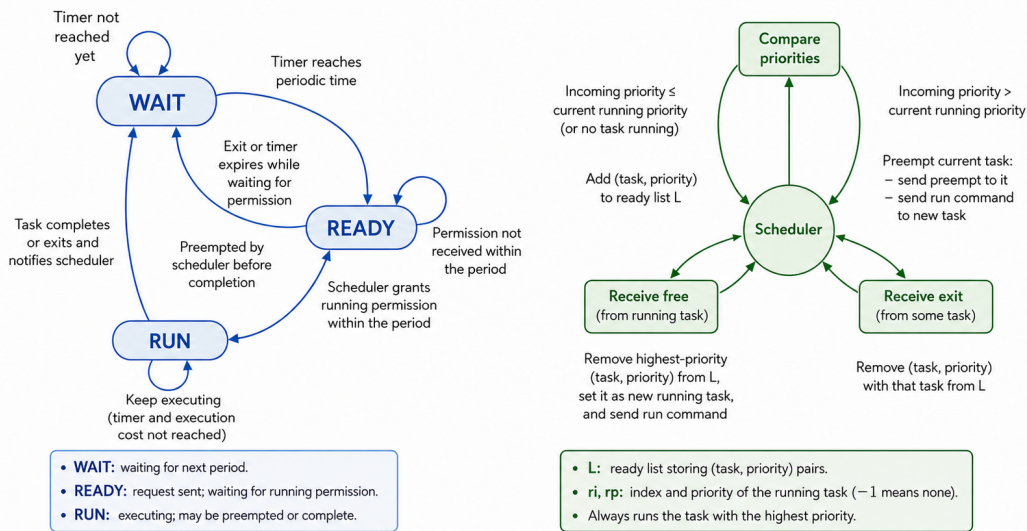


Figure 1: The Scheduler Case Study

**HHLProver.** Using HHLProver, we verify a simplified version of this case study. To focus on the essential scheduling behavior, we abstract away the internal computation performed by each task and model only its execution duration through a continuous-time evolution described by an ordinary differential equation. This abstraction preserves the timing and control aspects relevant to scheduling while significantly reducing the complexity of the model. The resulting system consists of three concurrent processes—two task modules and one scheduler—executing in parallel. The scheduler enforces mutual exclusion by guaranteeing that at most one task is in the RUN state at any time. At each scheduling point, it selects the highest-priority task among those in the READY state and may preempt the currently running task if a higher-priority task becomes ready. This case study demonstrates the capability of HHLProver to reason formally about nontrivial concurrent hybrid systems involving synchronization, preemption, and scheduling decisions. To complete the verification, we extend HHLProver with support for additional data structures, such as lists, thereby enabling the formal verification of practical scheduling algorithms that rely on dynamically maintained ready queues. The proof details can be seen in [ZJZ<sup>+</sup>23].

## 5 Conclusion and Outlook

The hybrid systems theorem proving friendly competition focuses on the characteristic features of hybrid systems theorem proving: flexibility of programming language principles for hybrid systems, unambiguous program semantics, and mathematically rigorous logical reasoning principles.

The automation tactic simplifications, nonlinear invariant generator improvements, and concurrent arithmetic backend utilization make a difference on some examples and especially in pure continuous systems verification performance, but their potential is not yet truly realized in case study verification performance. Future competitions are planned to extend the case study sub-category with game examples [CMP23] and stability examples [TMP22] to provide better assessment of verification performance on realistic examples, and to gain insight into potential proof automation to generalize the current specialized tactics and proof scripts from single example applicability to general-purpose proof automation using LLM-powered automated theorem proving [KLM<sup>+</sup>26]. A related challenge for proof repeatability and transferability are timeouts used in proof automation to decide how long to explore specific proof alternatives, and overall proof timeouts as used in this competition.

**Acknowledgments.** We thank the entire Logical Systems Lab for their many contributions to KeYmaera X and its associated tools. This material is based upon work supported by the National Science Foundation (NSF) under Grant No. CCF2427581. Xiangyu Jin and Shuling Wang are funded partly by the National Key R&D Program of China under grant No. 2022YFA1005104, and the NSFC under grant No. 62572459 and 62432005.

## References

- [BLCP18] Rose Bohrer, Adriel Luo, Xue An Chuang, and André Platzer. CoasterX: A case study in component-driven hybrid systems proof automation. *IFAC-PapersOnLine*, 2018. Analysis and Design of Hybrid Systems ADHS.
- [BRV<sup>+</sup>17] Rose Bohrer, Vincent Rahli, Ivana Vukotic, Marcus Völpl, and André Platzer. Formally verified differential dynamic logic. In Yves Bertot and Viktor Vafeiadis, editors, *Certified Programs and Proofs - 6th ACM SIGPLAN Conference, CPP 2017, Paris, France, January 16-17, 2017*, pages 208–221, New York, 2017. ACM.
- [CMP23] Rachel Cleaveland, Stefan Mitsch, and André Platzer. Formally verified next-generation airborne collision avoidance games in ACAS X. *ACM Trans. Embed. Comput. Syst.*, 22(1):10:1–10:30, 2023.
- [FMBP17] Nathan Fulton, Stefan Mitsch, Rose Bohrer, and André Platzer. Bellerophon: Tactical theorem proving for hybrid systems. In Mauricio Ayala-Rincón and César A. Muñoz, editors, *ITP*, volume 10499 of *LNCS*, pages 207–224. Springer, 2017.
- [FMQ<sup>+</sup>15] Nathan Fulton, Stefan Mitsch, Jan-David Quesel, Marcus Völpl, and André Platzer. KeYmaera X: An axiomatic tactical theorem prover for hybrid systems. In Amy Felty and Aart Middeldorp, editors, *CADE*, volume 9195 of *LNCS*, pages 527–538, Berlin, 2015. Springer.
- [GTMP22] James Gallicchio, Yong Kiam Tan, Stefan Mitsch, and André Platzer. Implicit definitions with differential equations for keymaera X - (system description). In *Automated Reasoning - 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022, Proceedings*, pages 723–733, 2022.
- [GZX<sup>+</sup>22] Panhua Guo, Bohua Zhan, Xiong Xu, Shuling Wang, and Wenhui Sun. Translating a large subset of stateflow to hybrid CSP with code optimization. *J. Syst. Archit.*, 130:102665, 2022.

- [He94] J. He. From CSP to hybrid systems. In *A Classical Mind, Essays in Honour of C.A.R. Hoare*, pages 171–189. Prentice Hall International (UK) Ltd., 1994.
- [KLM<sup>+</sup>26] Aditi Kabra, Jonathan Laurent, Ruben Martins, Stefan Mitsch, and André Platzer. Llm-powered automatic theorem proving and synthesis for hybrid systems and game. *CoRR*, abs/2603.00737, 2026.
- [LLQ<sup>+</sup>10] Jiang Liu, Jidong Lv, Zhao Quan, Naijun 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 *LNCS*, pages 1–15. Springer, 2010.
- [MGVP17] Stefan Mitsch, Khalil Ghorbal, David Vogelbacher, and André Platzer. Formal verification of obstacle avoidance and navigation of ground robots. *I. J. Robotics Res.*, 36(12):1312–1340, 2017.
- [Mit21] Stefan Mitsch. Implicit and explicit proof management in KeYmaera X. In *6th Workshop on Formal Integrated Development Environment, Proceedings*, 2021.
- [MJZ<sup>+</sup>21] Stefan Mitsch, Xiangyu Jin, Bohua Zhan, Shuling Wang, and Naijun Zhan. ARCH-COMP21 category report: Hybrid systems theorem proving. In *8th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH21), Brussels, Belgium, July 9, 2021*, pages 120–132, 2021.
- [MMJ<sup>+</sup>20] Stefan Mitsch, Jonathan Julián Huerta Y Munive, Xiangyu Jin, Bohua Zhan, Shuling Wang, and Naijun Zhan. Arch-comp20 category report: Hybrid systems theorem proving. In Goran Frehse and Matthias Althoff, editors, *ARCH20. 7th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH20)*, volume 74 of *EPiC Series in Computing*, pages 153–174. EasyChair, 2020.
- [MP16] Stefan Mitsch and André Platzer. The KeYmaera X proof IDE - concepts on usability in hybrid systems theorem proving. In *Proceedings of the Third Workshop on Formal Integrated Development Environment, F-IDE@FM 2016, Limassol, Cyprus, November 8, 2016*, pages 67–81, 2016.
- [MP20] Stefan Mitsch and André Platzer. A retrospective on developing hybrid system provers in the KeYmaera family - A tale of three provers. In Wolfgang Ahrendt, Bernhard Beckert, Richard Bubel, Reiner Hähnle, and Mattias Ulbrich, editors, *Deductive Software Verification: Future Perspectives - Reflections on the Occasion of 20 Years of KeY*, volume 12345 of *LNCS*, pages 21–64. Springer, 2020.
- [MST<sup>+</sup>19] Stefan Mitsch, Andrew Sogokon, Yong Kiam Tan, Xiangyu Jin, Bohua Zhan, Shuling Wang, and Naijun Zhan. ARCH-COMP19 category report: Hybrid systems theorem proving. In Goran Frehse and Matthias Althoff, editors, *ARCH19. 6th International Workshop on Applied Verification of Continuous and Hybrid Systems, part of CPS-IoT Week 2019, Montreal, QC, Canada, April 15, 2019*, volume 61 of *EPiC Series in Computing*, pages 141–161. EasyChair, 2019.
- [MSZ<sup>+</sup>23] Stefan Mitsch, Huanhuan Sheng, Bohua Zhan, Shuling Wang, Simon Foster, and Jonathan Julian Huerta Y Munive. ARCH-COMP23 category report: Hybrid systems theorem proving. In Goran Frehse and Matthias Althoff, editors, *Proceedings of 10th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH23)*, volume 96 of *EPiC Series in Computing*, pages 170–188. EasyChair, 2023.
- [MZS<sup>+</sup>22] Stefan Mitsch, Bohua Zhan, Huanhuan Sheng, Alexander Bentkamp, Xiangyu Jin, Shuling Wang, Simon Foster, Christian Pardillo Laursen, and Jonathan Julián Huerta Y Munive. ARCH-COMP22 category report: Hybrid systems theorem proving. In Goran Frehse, Matthias Althoff, Erwin Schoitsch, and Jeremie Guiochet, editors, *Proceedings of 9th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH22)*, volume 90 of *EPiC Series in Computing*, pages 185–203. EasyChair, 2022.
- [PC09] André Platzer and Edmund M. Clarke. Formal verification of curved flight collision avoid-

- ance maneuvers: A case study. In Ana Cavalcanti and Dennis Dams, editors, *FM*, volume 5850 of *LNCS*, pages 547–562, Berlin, 2009. Springer.
- [Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008.
- [Pla10] André Platzer. Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.*, 20(1):309–352, 2010.
- [Pla11] André Platzer. The structure of differential invariants and differential cut elimination. *Log. Methods Comput. Sci.*, 8(4), 2011.
- [Pla12] André Platzer. Logics of dynamical systems. In *LICS*, pages 13–24. IEEE, 2012.
- [Pla17] André Platzer. A complete uniform substitution calculus for differential dynamic logic. *J. Autom. Reas.*, 59(2):219–265, 2017.
- [PQ09] André Platzer and Jan-David Quesel. European Train Control System: A case study in formal verification. In Karin Breitman and Ana Cavalcanti, editors, *ICFEM*, volume 5885 of *LNCS*, pages 246–265, Berlin, 2009. Springer.
- [PT18] André Platzer and Yong Kiam Tan. Differential equation axiomatization: The impressive power of differential ghosts. In Anuj Dawar and Erich Grädel, editors, *LICS*, New York, 2018. ACM.
- [PT20] André Platzer and Yong Kiam Tan. Differential equation invariance axiomatization. *J. ACM*, 67(1):6:1–6:66, 2020.
- [QML<sup>+</sup>16] 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.
- [SCMP21] Matias Scharager, Katherine Cordwell, Stefan Mitsch, and André Platzer. Verified quadratic virtual substitution for real arithmetic. In *Formal Methods - 24th International Symposium, FM 2021, Virtual Event, November 20-26, 2021, Proceedings*, pages 200–217, 2021.
- [SMT<sup>+</sup>19] Andrew Sogokon, Stefan Mitsch, Yong Kiam Tan, Katherine Cordwell, and André Platzer. Pegasus: A framework for sound continuous invariant generation. In Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira, editors, *Formal Methods - The Next 30 Years - Third World Congress, FM 2019, Porto, Portugal, October 7-11, 2019, Proceedings*, volume 11800 of *LNCS*, pages 138–157. Springer, 2019.
- [SMT<sup>+</sup>21] Andrew Sogokon, Stefan Mitsch, Yong Kiam Tan, Katherine Cordwell, and André Platzer. Pegasus: sound continuous invariant generation. *Formal Methods Syst. Des.*, 58(1-2):5–41, 2021. Special issue for selected papers from FM’19.
- [The25] The MathWorks, Inc. *Counters Using Conditionally Executed Subsystems*, 2025.
- [TMP22] Yong Kiam Tan, Stefan Mitsch, and André Platzer. Verifying switched system stability with logic. In *HSCC ’22: 25th ACM International Conference on Hybrid Systems: Computation and Control, Milan, Italy, May 4 - 6, 2022*, pages 2:1–2:11, 2022.
- [WZZ15] S. Wang, N. Zhan, and L. Zou. An improved HHL prover: an interactive theorem prover for hybrid systems. In *ICFEM 2015*, volume 9407 of *LNCS*, pages 382–399. Springer, 2015.
- [XWZ<sup>+</sup>22] X. Xu, S. Wang, B. Zhan, X. Jin, J.-P. Talpin, and N. Zhan. Unified graphical co-modeling, analysis and verification of cyber-physical systems by combining AADL and simulink/stateflow. *Theor. Comput. Sci.*, 903:1–25, 2022.
- [XZW<sup>+</sup>23] Xiong Xu, Bohua Zhan, Shuling Wang, Jean-Pierre Talpin, and Naijun Zhan. A denotational semantics of simulink with higher-order UTP. *J. Log. Algebraic Methods Program.*, 130:100809, 2023.
- [ZJZ<sup>+</sup>23] N. Zhan, X. Jin, B. Zhan, S. Wang, and D. P. Guelev. A generalized hybrid hoare logic. *CoRR*, abs/2303.15020, 2023.
- [ZLW<sup>+</sup>13] Liang Zou, Jidong Lv, Shuling Wang, Naijun Zhan, Tao Tang, Lei Yuan, and Yu Liu. Veri-

- fyng chinese train control system under a combined scenario by theorem proving. In Ernie Cohen and Andrey Rybalchenko, editors, *Verified Software: Theories, Tools, Experiments - 5th International Conf., VSTTE 2013, Menlo Park, CA, USA, May 17-19, 2013, Revised Selected Papers*, volume 8164 of *LNCS*, pages 262–280. Springer, 2013.
- [ZWR95] Chaochen Zhou, Ji Wang, and Anders P. Ravn. A formal description of hybrid systems. In Rajeev Alur, Thomas A. Henzinger, and Eduardo D. Sontag, editors, *Hybrid Systems III: Verification and Control, Proceedings of the DIMACS/SYCON Workshop on Verification and Control of Hybrid Systems, October 22-25, 1995, Rutgers University, New Brunswick, NJ, USA*, volume 1066 of *LNCS*, pages 511–530. Springer, 1995.
- [ZWR96] Chaochen Zhou, Ji Wang, and Anders P. Ravn. A formal description of hybrid systems. In Rajeev Alur, Thomas A. Henzinger, and Eduardo D. Sontag, editors, *Hybrid Systems III*, volume 1066 of *LNCS*, pages 511–530. Springer, 1996.
- [ZYZ<sup>+</sup>14] Hengjun Zhao, Mengfei Yang, Naijun Zhan, Bin Gu, Liang Zou, and Yao Chen. Formal verification of a descent guidance control program of a lunar lander. In Cliff B. Jones, Pekka Pihlajasaari, and Jun Sun, editors, *FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings*, volume 8442 of *LNCS*, pages 733–748. Springer, 2014.