Chemical Case Studies in KeYmaera X

Rose Bohrer[0000−0001−5201−9895]

Worcester Polytechnic Institute, Worcester MA 01609, USA
rbohrer@wpi.edu

Abstract. Safety-critical chemical processes are well-studied in the formal methods literature, including hybrid systems models which combine discrete and continuous dynamics. This paper is the first to use a theorem-prover to verify hybrid chemical models: the KeYmaera X prover for differential dynamic logic. KeYmaera X provides parametric results that hold for a whole range of parameter values, non-linear physical dynamics, and a small trusted computing base.
We tell a general story about KeYmaera X: recent advances in automated reasoning about safety and liveness for differential equations have enabled elegant proofs about reaction dynamics.

Keywords: Hybrid Systems · Theorem Proving · Chemical Reactor

1 Introduction

Classical results on safe and optimal control [18] of chemical reactions [40] are the conceptual foundation for industrial chemical processes. Formal methods for chemical reactors are well-studied [4, 20, 25, 30, 36], but even textbook cases [18] lack *high fidelity* models (e.g., nonlinear dynamics and wide ranges of parameter values). We study textbook cases; these inform the study of practical cases.

We study (1) model-predictive control of an irreversible reaction (Sec. 3.1) and (2) an uncontrolled reversible reaction (Sec. 3.2) in KeYmaera X [16], a theorem-prover for *differential dynamic logic* (dL) [34]. See Sec. 4 for tradeoffs.

Both reactions contain challenges suitable as verification benchmarks: (1) nonlinear dynamics interacting with model-predictive controllers, and (2) theorems that test current tools' abilities regarding asymptotic properties, e.g., stability [27] or persistence [39]. Though reaction (2) is continuous, continuous reasoning is essential to hybrid. We find that new stability [41], variant [41], and Darboux polynomial [35] tools in KeYmaera X simplify our proofs.

2 Background

In KeYmaera X, correctness properties are stated and proved in *differential dynamic logic* (dL) [34], where hybrid systems are written in *hybrid program* notation. We discuss dL, then KeYmaera X usage.

2.1 Differential Dynamic Logic

We introduce dL syntax and semantics; see literature [34] for details. Semantics are state-based: state ω maps variables x to real numbers $\omega(x) : \mathbb{R}$. The syntax consists of terms (with a numeric meaning in each state), hybrid programs (which nondeterministically change the state when run), and formulas (which are true or false in each state). Terms are real-valued polynomials. Hybrid programs and formulas may contain each other. We use standard notation, e.g., $B ::= C \mid D$ means every B is either a C or a D.

Definition 1 (Hybrid Programs). *Hybrid programs α, β are defined by:*

$$\alpha, \beta ::= \ ?P \mid x := e \mid \{x' = f(x) \& Q\} \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^*$$

Hybrid programs are defined by their *runs*: from a starting state, what final states are reachable? Hybrid programs can have one run (deterministic), many runs (nondeterministic), or zero runs (early termination). Programs $?P$ and $\{x' = f \& Q\}$ contain formulas P and Q; see Def. 2 for more about formulas.

The test program $?P$ never modifies the state; if formula P is true, then $?P$ ends in the current state, but if P is false, then $?P$ has no final states, representing execution failure. Deterministic assignment $x := e$ updates the state by storing the current value of term e in variable x. Ordinary differential equation systems (ODEs) are the defining feature of hybrid programs: ODEs composed with discrete operations model hybrid systems. ODE $\{x' = f(x) \& Q\}$ evolves in continuous time with $x' = f(x)$, where $f(x)$ is a term. The duration of evolution is nondeterministic. If an *evolution domain constraint* Q is provided, Q is tested continuously, and evolution must stop before Q ever becomes false. Choices $\alpha \cup \beta$ nondeterministically run *either* α *or* β, as opposed to running both. Composition $\alpha; \beta$ runs α, then β in the resulting state(s). Duration of loops α^* is nondeterministically-chosen but finite: zero, one, or many repetitions can occur. We also use $\mathsf{if}(P)\{\alpha\}\mathsf{else}\{\beta\}$, which reduces to choices and tests.

Definition 2 (Formulas). *There are many formulas P, Q in* dL*. We only use:*

$$P, Q ::= \cdots \mid e \geq \tilde{e} \mid \neg P \mid P \wedge Q \mid P \rightarrow Q \mid [\alpha]P \mid \langle \alpha \rangle P$$

Formulas represent true/false questions about the state ω. Comparison $e \geq \tilde{e}$ is true whenever the value of e is at least that of $\tilde{e}$ in a given state. All other comparisons $e > \tilde{e}, e = \tilde{e}, e \neq \tilde{e}, e \leq \tilde{e}, e < \tilde{e}$ are definable using $e \geq \tilde{e}$ and other logical connectives, so we use them freely. Negation $\neg P$ is true when P is false. Conjunction $P \wedge Q$ is true when both P and Q are. Implication $P \rightarrow Q$ is true when P's truth would imply Q's truth.

The defining formulas of dL, $[\alpha]P$ and $\langle \alpha \rangle P$, are respectively true in state ω if *every* or *some* run of α starting from state ω ends in a state where P is true. For many programs α, including all in this paper, *all runs* equates to *all time*.

KeYmaera X proves truth in *every state*, called *validity*.

We use standard notation for axioms and proof rules. Each rule has a horizontal line and means: if all *premise* formulas above the line are valid, so is

the *conclusion* formula below. Rules can use *schema variables* (e.g., P, α) for arbitrary programs or formulas, respectively. For example, the loop rule

$$\text{LOOP}\ \frac{P \to J \qquad J \to [\alpha]J \qquad J \to Q}{P \to [\alpha^*]Q}$$

means for all P, Q, J, α that if premises $P \to J$, $J \to [\alpha]J$, and $J \to Q$ are all valid, so is $P \to [\alpha^*]Q$. Formula J is *proved* true for all iterations, thus we call J the *loop invariant*. This *proven* loop invariant should not be confused with use of the word *invariant* in hybrid automata to mean an *assumed* constraint on ODE evolution. We call such constraints *evolution domain constraints*.

2.2 KeYmaera X

We discuss the KeYmaera X [28] user interface (Fig. 1). KeYmaera X is an interactive, tactical prover: users interactively pick proof techniques at each step. Each technique is implemented as a *tactic* [15] program. Tactics range from propositional rules (e.g., conjunction and implication) to complex search procedures. The *default* (`auto`) procedure tries many methods, solving many simple problems automatically. User effort varies much between proofs. We will discuss how automation reduces effort. Tactics help with rigor: complex methods reduce to simple, trusted steps.

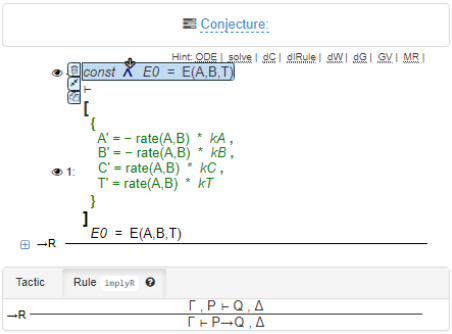

Fig. 1. KeYmaera X. Clicking highlighted symbol does a step. Last rule is shown at bottom. Top shows suggestions.

3 Results

We contribute case studies on two classic kinds of chemical reactions. The first is an irreversible reaction in a well-mixed adiabatic batch reactor. The second case study is a reversible reaction between two compounds, i.e., where the output can react again and form the input. We chose these examples because they are classic [38]. Both case studies emphasize recent advances in KeYmaera X proof automation, which simplified proofs. Where limitations remain, we discuss them.

3.1 Controlled Irreversible Reactions

We formalize a classic scenario: an irreversible, exothermic reaction in an adiabatic, well-mixed batch reactor. *Irreversible* [38, §2.1] means the reaction is one-way: outputs do not react to create inputs. *Adiabatic* [38, §2.14] means heat

does not leave or enter the reactor. *Well-mixed* [38, §2.12] means the reaction occurs evenly in space throughout the reactor. In this basic synthesis reaction, two (first-order) reactants react to form a third, plus heat ($A+B \longrightarrow C+heat$).

The case study contains four models, each with proof. The first shows conservation of energy, validating that adiabatic reactors are closed systems. The remaining three models add a model-predictive bang-bang controller [18], which predicts future behavior according to the model, then applies an all-or-nothing control action. It is proved that the control ensures a safety property: overheating is prevented. We use this standard control approach in order to focus on the continuous reaction dynamics. The driving difference between the last three models is their increasingly complex reaction dynamics, which mandate increasingly complex controls and proofs. In the second model, the reaction rate is constant. In the third model, the rate is linear in temperature, thus exponential in time. In the final model, the rate is proportional to the product of temperature and each concentration, with resulting dynamics beyond a simple exponential, yet still approximate. Approximate results are the best that can be expected for non-linear dynamics. We discuss why, including verification challenges.

Each model approximates textbook [38, Eq. 2.93] reaction dynamics, where the reaction *rate* is proportional to the product of concentrations of each reactant A and B multiplied by a coefficient. Recall that the *concentration* of a reactant in a mixture is the quantity of that reactant per unit quantity of the mixture. The rate equation is $rate = kAB$ where k is an exponential given by the Arrhenius equation [38, Eq. 5.1]. That is, $k(T) = k_0 e^{-E/RT}$ where T is temperature, R is the ideal gas constant, E is the reaction's activation energy and k_0 a constant.

Analysis of the reaction rate dynamics is nontrivial: *rate* is a product of three continuously-changing quantities, resulting in a non-linear ODE. Moreover, $k(T)$ is exponential in T, resulting in a *non-polynomial* ODE. KeYmaera X handles non-linear ODEs well, but is restricted to polynomial ODEs, as is standard. We thus reach our first limitation: to ensure a polynomial ODE, we approximate the temperature dependence as linear. This assumption is reasonable because polynomial ODEs are a standard assumption, and our nonlinear dynamics are still more precise than prior models [20, 25, 30, 36, 45]. Our second limitation is that the reactants are first-order, so their influence on rate is linear. We do so because such reactions are common and lead to elegant equations. KeYmaera X supports polynomials of any degree, so we expect the approach to work for higher-order reactions, so long as the order is fixed. Limitations aside, the results are fully parametric, e.g., the results can be applied to *any* first-order reactants in *any* amount by plugging in new coefficients and concentrations.

Energy Conservation. The basic dL model for energy conservation is presented in Fig. 2. Energy conservation is interesting in its own right, because it implies the system is closed. This helps support our claim that the model is *adiabatic*: heat energy does not leave nor enter. The variables A, B, and C stand for the current concentration of each reactant present in the reactor. Reactor temperature is written T. In our analysis, we decompose energy into

$$\mathsf{E} \equiv \mathsf{KE} + \mathsf{U} \quad \mathsf{U} \equiv \min(\mathsf{A}/\mathsf{k}_A, \mathsf{B}/\mathsf{k}_B)\,\mathsf{k}_\mathsf{T} \quad \mathsf{KE} \equiv \mathsf{T} \quad \mathsf{rate} \equiv \mathsf{T}_s\,\mathsf{A}_0\,\mathsf{B}_0\,\mathsf{k}_\mathsf{ra} + \mathsf{k}_\mathsf{rb}$$
$$\mathsf{const} \equiv \mathsf{k}_\mathsf{ra} > 0 \land \mathsf{k}_\mathsf{rb} \geq 0 \land \mathsf{k}_A > 0 \land \mathsf{k}_B > 0 \land \mathsf{k}_C > 0 \land \mathsf{k}_\mathsf{T} > 0$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = -\mathsf{rate}\,\mathsf{k}_A, \mathsf{B}' = -\mathsf{rate}\,\mathsf{k}_B, \mathsf{C}' = \mathsf{rate}\,\mathsf{k}_C, \mathsf{T}' = \mathsf{rate}\,\mathsf{k}_\mathsf{T}\}$$
$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{E}_0 = \mathsf{E} \to [\mathsf{ode}]\mathsf{E}_0 = \mathsf{E})$$

Fig. 2. Conservation-of-energy for uncontrolled irreversible reaction, constant heating

kinetic (heat) and potential (chemical) energy: $\mathsf{E} \equiv \mathsf{KE} + \mathsf{U}$. Potential energy $\mathsf{U} \equiv \min(\mathsf{A}/\mathsf{k}_A, \mathsf{B}/\mathsf{k}_B)\,\mathsf{k}_\mathsf{T}$ is the product of the amount (concentration) of C remaining to be produced (the reaction ends when either A or B is exhausted) with the heat released per unit amount (C). That is, we model C as if it possesses no potential energy, since we are interested only in energies relevant to the current reaction. We model the reaction rate as $\mathsf{T}_s\,\mathsf{A}_0\,\mathsf{B}_0\,\mathsf{k}_\mathsf{ra} + \mathsf{k}_\mathsf{rb}$, which makes two intentional simplifications. First, we use approximate *current* concentrations A, B with *initial* concentrations $\mathsf{A}_0, \mathsf{B}_0$. Secondly, we simplify the temperature factor to T_s, which is a *constant* even as temperature T changes, thus the influence of heat is *static* throughout the reaction. We determine the reaction rate as a product of the concentration factor and temperature factor. For generality, the coefficients $\mathsf{k}_\mathsf{ra}, \mathsf{k}_\mathsf{rb}$ let the rate be any *linear function of* the product. Formula const specifies signs of constants.

The ode indicates that all concentrations $\mathsf{A}, \mathsf{B}, \mathsf{C}$ and the reactor temperature T all change proportional to the reaction rate; A and B are lost as C and heat are gained. Coefficients $\mathsf{k}_A, \mathsf{k}_B, \mathsf{k}_C, \mathsf{k}_\mathsf{T}$ indicate the rates at which each changes, which may depend respectively on the stoichiometric coefficients of the reaction or how strongly exothermic it is.

Finally, the theorem statement $(P \to [\alpha]Q)$ states that under the simple constant assumptions, energy is conserved because at all times the current energy E remains equal to its initial value E_0. We now describe the KeYmaera X proof.

Proof. The default proof procedure of KeYmaera X (Sec. 2.2) proves the theorem automatically with *differential invariants* [34, Lem. 11.3], demonstrating the capabilities of this standard dL rule. We present the (relevant case of the) *differential invariant* [34, Lem. 11.3] rule

$$\text{DI}\ \frac{Q \to [x' := f(x)](e)' = (\tilde{e})'}{e = \tilde{e} \to [\{x' = f(x) \& Q\}]e = \tilde{e}}$$

which shows $e = \tilde{e}$ is true throughout an ODE if it holds initially and differentials are equal throughout. We prove $\mathsf{E}_0 = \mathsf{E}$ thus: E_0 is constant, so proving $\mathsf{E}' = 0$ throughout suffices. Expanding the definition of E yields $(\mathsf{E})' = (\mathsf{T} + \min(\mathsf{A}/\mathsf{k}_A, \mathsf{B}/\mathsf{k}_B)\,\mathsf{k}_\mathsf{T})' = \mathsf{rate}\,\mathsf{k}_\mathsf{T} + \min((\mathsf{A})'/\mathsf{k}_A, (\mathsf{B})'/\mathsf{k}_B)\,\mathsf{k}_\mathsf{T} = \mathsf{rate}\,\mathsf{k}_\mathsf{T} + \min(-\mathsf{rate}\,\mathsf{k}_A/\mathsf{k}_A, -\mathsf{rate}\,\mathsf{k}_B/\mathsf{k}_B)\,\mathsf{k}_\mathsf{T} = \mathsf{rate}\,\mathsf{k}_\mathsf{T} + \min(-\mathsf{rate}, -\mathsf{rate})\,\mathsf{k}_\mathsf{T} = (\mathsf{rate} - \mathsf{rate})\,\mathsf{k}_\mathsf{T} = 0$. Due to KeYmaera X's automation, the entire proof is automatic.

On-Off Reactions. This model keeps the basic heating dynamics but adds bang-bang control. Fig. 3 describes the model in full. Parts unchanged from Fig. 2 are grayed out to aid comparison. The impact of this theorem is that the reactor is provably safe under idealistic assumptions, i.e., when concentrations and temperatures change very little or have little impact on reaction rate.

$$\mathsf{rate} \equiv \mathsf{T}_s\,\mathsf{A}_0\,\mathsf{B}_0\,\mathsf{k}_{\mathsf{ra}} + \mathsf{k}_{\mathsf{rb}}$$
$$\mathsf{const} \equiv \mathsf{k}_{\mathsf{ra}} > 0 \land \mathsf{k}_{\mathsf{rb}} \geq 0 \land \mathsf{k}_A > 0 \land \mathsf{k}_B > 0 \land \mathsf{k}_C > 0 \land \mathsf{k}_{\mathsf{T}} > 0 \land \mathsf{T} > 0 \land \epsilon > 0$$
$$\mathsf{ctrl} \equiv \{\mathsf{if}(\mathsf{T}_{max} - \mathsf{T} \leq \epsilon\,\mathsf{rate}\,\mathsf{k}_{\mathsf{R}})\{\mathsf{isOn} := 0\}\mathsf{else}\{\mathsf{isOn} := 1\}\}; t := 0$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_A, \mathsf{B}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_B, \mathsf{C}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k}_C,$$
$$\mathsf{T}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k}_{\mathsf{T}}, t' = 1 \land t \leq \epsilon \land \mathsf{A} \geq 0 \land \mathsf{B} \geq 0 \land \mathsf{C} \geq 0\}$$

$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{T} \leq \mathsf{T}_{max} \to [\{\mathsf{ctrl}; \mathsf{ode}\}^*]\mathsf{T} \leq \mathsf{T}_{max})$$

Fig. 3. Safety for irreversible reaction with bang-bang control, constant heating

The greatest change is the addition of a *time-triggered* controller: the system now repeats in a loop, with the controller guaranteed to run at least every $\epsilon > 0$ time units. The controller (ctrl) is *model-predictive* because it *predicts* whether it would be dangerous to keep the reaction running for ϵ time: if the remaining temperature buffer $\mathsf{T}_{max} - \mathsf{T}$ is no more than the temperature change that could occur after reacting for time ϵ, it would be unsafe to keep reacting. If so, the reaction shuts off ($\mathsf{isOn} := 0$), else it turns on ($\mathsf{isOn} := 1$). Note isOn is an *indicator variable*; its only possible values are 0 and 1. Specifically, the controller linearly predicts the maximum temperature change as $\epsilon\,\mathsf{rate}\,\mathsf{k}_{\mathsf{R}}$ and shuts off if the safe temperature would be exceeded. Importantly, this approach predicts unsafe events before they occur and shuts off before the damage is done. Either way, the timer t is reset to 0.

The ode is updated so that each reaction equation is multiplied by isOn, causing no physical changes to occur when the reactor is turned off. This model is best-suited for situations where it is possible to quickly halt a reaction. The ode gains an *evolution domain constraint*, which serves to restrict its duration of evolution: an ODE may evolve only while the constraint remains true. Our constraint serves two purposes. Firstly, $t \leq \epsilon$ implements time-triggering: if each iteration takes at most ϵ time, there is at most ϵ delay between control cycles. Secondly, the constraints $\mathsf{A} \geq 0 \land \mathsf{B} \geq 0 \land \mathsf{C} \geq 0$ model the assumption of nonnegative concentrations. For example, the reaction ends if A or B reach zero.

Finally, the updated theorem statement $(P \to [\alpha]Q)$ is now a safety statement, stating that the reactor never exceeds its maximum safe temperature.

Proof. As the model now contains a loop, the proof uses *loop invariant* reasoning in addition to *differential invariant reasoning*, both distinct concepts from *evolution domain constraints*. We prove that the safety condition $\mathsf{T} \leq \mathsf{T}_{max}$ is a *loop invariant*, meaning it holds before and after every loop repetition. We use the standard loop rule from Sec. 2.1.

Already, a lemma arises in the ODE proof. Certain *differential invariant* proofs can only succeed by first proving lemmas, called *differential cut* formulas, which are then available as assumptions in the invariant proof. Specifically, we prove the cut $\mathsf{T}_{max} - \mathsf{T} > (\epsilon - t)\,\mathsf{rate}\,\mathsf{k_T}$, meaning the remaining safe temperature gap exceeds the projected temperature change during the remaining time. The cut proves automatically by the differential invariant rule, from which the loop invariant, then safety condition, follow by automatic proof.

Fixed Exponents. For the next model, the first fundamental change is that we update the definition of rate to use the current temperature, so that the reaction rate evolves exponentially over time. Because dynamic reaction rates are an increase in complexity, we simplify other aspects of the reaction rate formula by dropping $\mathsf{k_{ra}}$ and $\mathsf{k_{rb}}$. The remaining changes follow from that one: amts is a helper definition for definitions such as $\mathsf{taylor}^+(x,t)$, which is an upper bound on temperature over time, constructed as a Taylor series approximation. Taylor series bring a fundamentally new proof approach for more complicated dynamics: exponential dynamics need approximations in dL. Taylor series are a flexible approximation: if precision were unsatisfactory, the degree could be increased. However, this Taylor bound is only provably an upper bound on a limited time interval which happens to be $1/(2\,\mathsf{amts})$, which we thus take as our upper limit on ϵ. In practice, we hypothesize that the time limit is artificial: time could be expressed in any desired units, increasing the interval. The constants are updated to include assumptions on initial values of amounts and the controller is updated to use the Taylor approximation. The ode is updated to explicitly assume nonnegative temperature, which is a safe assumption since our goal is to avoid high, not low, temperatures. This new result shows safety with idealized modeling of concentrations under more realistic *heating* assumptions.

$$\mathsf{rate} \equiv \mathsf{T}\,\mathsf{A}_0\,\mathsf{B}_0 \quad \epsilon \equiv 1/(2\,\mathsf{amts}) \quad \mathsf{amts} \equiv \mathsf{k_T}\,\mathsf{A}_0\,\mathsf{B}_0 \quad \mathsf{taylor}^+(x,t) \equiv (1 + 2\,t\,\mathsf{amts})\,x$$
$$\mathsf{const} \equiv \mathsf{k}_A > 0 \land \mathsf{k}_B > 0 \land \mathsf{k}_C > 0 \land \mathsf{k_T} > 0 \land \epsilon > 0 \land \mathsf{A}_0 \geq 0 \land \mathsf{B}_0 \geq 0$$
$$\mathsf{ctrl} \equiv \{\mathsf{if}(\mathsf{T}_{max} \leq \mathsf{taylor}^+(\mathsf{T},\epsilon))\{\mathsf{isOn} := 0\}\mathsf{else}\{\mathsf{isOn} := 1\}\}; t := 0$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_A, \mathsf{B}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_B, \mathsf{C}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k}_C,$$
$$\mathsf{T}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k_T}, t' = 1 \land t \leq \epsilon \land \mathsf{A} \geq 0 \land \mathsf{B} \geq 0 \land \mathsf{C} \geq 0 \land \mathsf{T} \geq 0\}$$
$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{T} > 0 \land \mathsf{T} \leq \mathsf{T}_{max} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = \mathsf{B}_0 \to [\{\mathsf{ctrl}; \mathsf{ode}\}^*]\mathsf{T} \leq \mathsf{T}_{max})$$

Fig. 4. Safety for irreversible reaction with bang-bang control, fixed-exponent heating

Proof. The loop invariant is unchanged. We add several differential cuts; order matters since each one can serve as an assumption in following proofs: i) $t \geq 0$ just means time is nonnegative, ii) $\mathsf{A}_0\,\mathsf{B}_0\,\mathsf{T}\,\mathsf{k_T} \geq 0$ ensures forward (or 0) reaction rate, and iii) $\mathsf{taylor}^+(\mathsf{T}_{old}, t) - \mathsf{T} \geq 0$ bounds temperature T above with $\mathsf{taylor}^+()$ in terms of old temperature T_{old}. The final cut requires advanced proof techniques because term $\mathsf{taylor}^+(\mathsf{T}_{old}, t) - \mathsf{T}$ decreases; differential invariants alone are provably [32, Thm 6.1] insufficient for such terms. KeYmaera X can

solve this goal with the following high-level rule that uses Darboux polynomial (inequality) reasoning [35, Corr. 3.2]:

$$\text{DBX}_{\succcurlyeq}\ \frac{Q \rightarrow (p)' \geq g\,p}{p \succcurlyeq 0 \rightarrow [\{x' = f(x) \& Q\}]p \succcurlyeq 0}$$

Here, both instances of $\succcurlyeq$ are replaced uniformly with one of $>$ or $\geq$. Note $(e)'$ is shorthand for the *Lie derivative* of p, with all variables of form x' replaced by their corresponding $f(x)$. The polynomial p is called a *Darboux* polynomial if the premise holds, then polynomial g is called its *cofactor*. It is natural to ask what power is gained by the addition of this proof rule. Certainly, it is stronger than differential invariant reasoning (which would require $Q \rightarrow (p)' \geq 0$) because $g\,p$ is allowed to be negative. Yet its full usefulness goes deeper, as the rule serves as a basis for differential radical invariant reasoning which is provably complete for semianalytic invariants [35, Thm. 4.5], a large class of invariants.

KeYmaera X's built-in invariant generator can search for Darboux polynomials, but it did not find a suitable polynomial for our example, so we found one manually by algebra. Using the definition of the ODE, we solved for a polynomial that satisfies the proof goal, in this case: $g \equiv \mathsf{A}_0\,\mathsf{B}_0\,\mathsf{k}_\mathsf{T}$. After choosing a suitable Darboux polynomial, the remaining proof goals completed using KeYmaera X's default proof method. Further applications of Taylor approximations are discussed in Sec. 4.

Dynamic Exponents. Even our final controlled model (Fig. 5) makes some important simplifying assumptions. Note that our model makes the impact of temperature on reaction rate a linear one, whereas the true Arrhenius equation [38, Eq. 5.1] implies an exponential effect on reaction rate. Linear functions can locally approximate exponential ones, but exponentials remain of future interest. Despite these limitations, the final model is important because it shows safety with both nontrivial heating *and* concentration dynamics.

The core change in the final model is a more advanced reaction rate dynamics, where the reaction rate dynamically changes in response to the concentration of each reactant. Definitions amts and ϵ are updated for the same reason. The timestep ϵ now changes dynamically: as the reaction proceeds, the acceptable delay *increases*, thus becoming easier to satisfy. It simplifies the analysis to have ϵ change only at each loop iteration rather than continuously, so we introduce variables $\mathsf{A}_1, \mathsf{B}_1$ to stand for the values of A, B at the *start* of each ODE evolution. The changes to the model are modest, but the dynamic changes are notable: the reaction rate is now a product of three changing variables, no longer an exponential with a fixed base. Likewise, additional proof steps will be required to account for changing concentrations, but the core proof approach is unchanged.

Proof. In this proof, the reaction rate changes as the concentration of each reactant changes, so we strengthen the loop invariant to capture the status of the reactant concentrations: $0 \leq \mathsf{T} \wedge \mathsf{T} \leq \mathsf{T}_{max} \wedge \mathsf{A} \leq \mathsf{A}_0 \wedge \mathsf{B} \leq \mathsf{B}_0$. The differential cuts are similar to before, with an additional lemma that the concentrations of

$$\mathsf{rate} \equiv \mathsf{T}\,\mathsf{A}\,\mathsf{B} \quad \epsilon \equiv 1/(2\,\mathsf{A}_1\,\mathsf{B}_1\,\mathsf{k}_\mathsf{T}) \quad \mathsf{amts} \equiv \mathsf{A}\,\mathsf{B}\,\mathsf{k}_\mathsf{T}$$
$$\mathsf{const} \equiv \mathsf{k}_A > 0 \land \mathsf{k}_B > 0 \land \mathsf{k}_C > 0 \land \mathsf{k}_\mathsf{T} > 0 \land \epsilon > 0 \land \mathsf{A}_0 \geq 0 \land \mathsf{B}_0 \geq 0$$
$$\mathsf{ctrl} \equiv \{\mathsf{if}(\mathsf{T}_{max} \leq \mathsf{taylor}^{+}(\mathsf{T}, \epsilon))\{\mathsf{isOn} := 0\}\mathsf{else}\{\mathsf{isOn} := 1\}\}; t := 0; A_1 := A; B_1 := B$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_A, \mathsf{B}' = \mathsf{isOn} \cdot -\mathsf{rate}\,\mathsf{k}_B, \mathsf{C}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k}_C,$$
$$\mathsf{T}' = \mathsf{isOn} \cdot \mathsf{rate}\,\mathsf{k}_\mathsf{T}, t' = 1 \land t \leq \epsilon \land \mathsf{A} \geq 0 \land \mathsf{B} \geq 0 \land \mathsf{C} \geq 0 \land \mathsf{T} \geq 0\}$$
$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{T} > 0 \land \mathsf{T} \leq \mathsf{T}_{max} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = \mathsf{B}_0 \to [\{\mathsf{ctrl}; \mathsf{ode}\}^*]\mathsf{T} \leq \mathsf{T}_{max})$$

Fig. 5. Safety for irreversible reaction with bang-bang control, advanced heating

the first two reactants do not increase: $\mathsf{A} \leq \mathsf{A}_1 \land \mathsf{A}_1 \leq \mathsf{A}_0 \land \mathsf{B} \leq \mathsf{B}_1 \land \mathsf{B}_1 \leq \mathsf{B}_0$. The differential cut for the Taylor series is unchanged, and the same Darboux polynomial $g \equiv \mathsf{A}_0\,\mathsf{B}_0\,\mathsf{k}_\mathsf{T}$ suffices.

3.2 Uncontrolled Reversible Reactions

We study reversible reactions. We consider a textbook scenario where two reactants A and B can each react to form the other ($A \rightleftharpoons B$). To our knowledge, we provide the first computer-checked proofs for the asymptotic behavior of this classic, widely-used textbook scenario. Specifically, our final model shows *persistence* [39], a relative of stability: the system eventually gets arbitrarily close to its equilibrium state, then stays close forever. We build up to this result with lemmas: the system is always moving (nonstrictly) toward equilibrium and can arbitrarily approach equilibrium in finite, bounded time. To complete the story, we show that although the equilibrium can always be arbitrarily approximated, it can never be reached exactly.

Pure Reactant Decreases. We consider a scenario starting with pure reactant A, which then becomes a mixture. We show the current amount of A never exceeds the initial amount, which is intuitive by conservation of mass. The lemma might be of practical use directly, e.g., to verify that a container never overflows, but we mainly use it as a lemma for persistence. Here, the two reactants are

$$\mathsf{ode} \equiv \{\mathsf{A}' = -\mathsf{A}\,\mathsf{k}_\mathsf{F} + \mathsf{B}\,\mathsf{k}_\mathsf{R}, \mathsf{B}' = \mathsf{A}\,\mathsf{k}_\mathsf{F} - \mathsf{B}\,\mathsf{k}_\mathsf{R}\}$$
$$\mathsf{const} \equiv \mathsf{A}_0 > 0 \land \mathsf{k}_\mathsf{R} > 0 \land \mathsf{k}_\mathsf{F} > 0$$
$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = 0 \to [\mathsf{ode}]\mathsf{A} \leq \mathsf{A}_0)$$

Fig. 6. Concentration of A is nonincreasing during reversible reaction.

named A and B, with initial values $\mathsf{A} = \mathsf{A}_0 > 0$ and $\mathsf{B} = 0$. Reactants A and B are engaged in a *reversible reaction* where A converts to B at forward rate k_F and B converts to A at reverse rate k_R. It is well-known [38, Ch. 3] that the system asymptotically approaches an equilibrium state, called a *dynamic equilibrium*,

in which the forward and reverse reactions perfectly cancel out. We define ode using a classic textbook model of a reversible reaction, which does not model heat: the reaction rates are based solely on concentrations and constants.

Proof. This proof completes automatically: the automatic prover successfully reasons by differential invariant.

Equilibrium Avoidance. We show that the amounts of the reactants never exactly reach the equilibrium. Though not directly used in the persistence proof, we prove this because it is a fundamental property in its own right which tacitly influences how a chemical plant is designed and operated. An operator would never wait for perfect equilibrium to occur, only for the system to get *close* to equilibrium, because perfect equilibrium (provably) never occurs.

The initial condition and ODE are unchanged, only the postcondition changes, which mandates a new proof approach. To state the new postcondition, we define the amounts of A present at the equilibrium ($\tilde{A}$). The above definition of

$$\mathsf{ode} \equiv \{\mathsf{A}' = -\mathsf{A}\,\mathsf{k_F} + \mathsf{B}\,\mathsf{k_R}, \mathsf{B}' = \mathsf{A}\,\mathsf{k_F} - \mathsf{B}\,\mathsf{k_R}\}$$
$$\mathsf{const} \equiv \mathsf{A}_0 > 0 \land \mathsf{k_R} > 0 \land \mathsf{k_F} > 0 \qquad \tilde{A} \equiv \mathsf{A}_0\,(\mathsf{k_R}/(\mathsf{k_F} + \mathsf{k_R}))$$
$$(P \to [\alpha]Q) \equiv (\mathsf{const} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = 0 \to [\mathsf{ode}]\mathsf{A} \neq \tilde{A})$$

Fig. 7. Equilibrium is never reached during reversible reaction.

$\tilde{A}$ can be found by solving for equilibrium ($\mathsf{A}' = 0 \land \mathsf{B}' = 0$) in ode subject to conservation of mass ($\mathsf{A} + \mathsf{B} = \mathsf{A}_0$).

Proof. A simple change in postcondition creates a major increase in proof complexity, because we now wish to show a lower bound instead of an upper bound. We use multiple differential cuts, one of which uses Darboux reasoning.

- $\mathsf{A} - \mathsf{A}_0\,(\mathsf{k_R}/(\mathsf{k_F} + \mathsf{k_R})) > 0$ means A's rate of change is always in the direction of the equilibrium
- $\mathsf{A} + \mathsf{B} = \mathsf{A}_0$ is conservation of mass
- $\mathsf{A} > 0 \land \mathsf{B} \geq 0$ means we never have a negative amount of either reactant, the first being positive. This requires a Darboux argument with polynomial $-(\mathsf{k_F} + \mathsf{k_R})$ because the amount of the first reactant does decrease with time.

Once these cuts have been proven, automation suffices to finish the proof.

Equilibrium Approach. We show that we get arbitrarily close to the equilibrium, given sufficient time. For every positive epsilon ($\epsilon > 0$), there exists a time when we get that close to the equilibrium. The assumption changes slightly; the

theorem statement changes more: we prove a *diamond* modality $\langle \mathsf{ode} \rangle \mathsf{A} \leq \tilde{A} + \epsilon$ because we want to show we *eventually* approach the equilibrium. The practical impact of this result is that if an engineer desires an almost-perfect equilibrium, that can be attained, but the cost is time.

$$\mathsf{const} \equiv \mathsf{A}_0 > 0 \land \mathsf{k_R} > 0 \land \mathsf{k_F} > 0 \land \epsilon > 0$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = -\mathsf{A}\,\mathsf{k_F} + \mathsf{B}\,\mathsf{k_R}, \mathsf{B}' = \mathsf{A}\,\mathsf{k_F} - \mathsf{B}\,\mathsf{k_R}\} \qquad \tilde{A} \equiv \mathsf{A}_0\,(\mathsf{k_R}/(\mathsf{k_F} + \mathsf{k_R}))$$
$$(P \to \langle \alpha \rangle Q) \equiv (\mathsf{const} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = 0 \to \langle \mathsf{ode} \rangle \mathsf{A} \leq \tilde{A} + \epsilon)$$

Fig. 8. Equilibrium is approached during reversible reaction.

Proof. Previous proofs highlighted advances in proof automation for box properties of ODEs; this proof relies on advances in proof automation for diamond properties of ODEs. The *differential variant* rule is the diamond counterpart to *differential invariant* reasoning for box properties. The *differential variant* principle [41, Corr. 24] says: if a progress bound $d > 0$ on derivative $(p)'$ holds everywhere outside the goal region $(\neg(p \succcurlyeq 0))$, then we reach the goal eventually:

$$\mathrm{DV}^{\exists}_{\succcurlyeq} \frac{\exists d > 0\, \forall x (\neg(p \succcurlyeq 0) \to (p)' \geq d)}{\langle \{x' = f(x)\} \rangle p \succcurlyeq 0}$$

where $\succcurlyeq$ stands uniformly for either $>$ or $\geq$, where d is a fresh variable, and where $x' = f(x)$ provably has a global solution (i.e., for all time). In the premise, $(p)'$ is shorthand for the *Lie derivative* of p, with all variables of form x' replaced by their corresponding $f(x)$.

The key insight behind our proof is that the rate of progress is proportional to our current displacement from the equilibrium. Since we seek to get the displacement within some ϵ, we can assume without loss of generality that the current displacement is at least ϵ, giving a bound d on the progress rate: $d = \epsilon\,(\mathsf{k_F} + \mathsf{k_R})$. This progress rate also confirms standard intuitions about the system dynamics: higher rates of progress are made when far away from the equilibrium and when reaction rates are high.

Persistence. Persistence means there exists a point after which we forever remain within ϵ of the equilibrium. Persistence is of practical importance because it shows both that the system can get arbitrarily close to equilibrium *and* that the system stays that way *indefinitely*. In short, this result is important from a control perspective because it shows the system is well-controlled, even without a controller. As a theorem-proving case study, persistence is an excellent comprehensive test case because it combines boxes and diamonds. Only the theorem statement need be updated; all other definitions are unchanged:

$$\mathsf{const} \equiv \mathsf{A}_0 > 0 \land \mathsf{k_R} > 0 \land \mathsf{k_F} > 0 \land \epsilon > 0$$
$$\mathsf{ode} \equiv \{\mathsf{A}' = -\mathsf{A}\,\mathsf{k_F} + \mathsf{B}\,\mathsf{k_R}, \mathsf{B}' = \mathsf{A}\,\mathsf{k_F} - \mathsf{B}\,\mathsf{k_R}\} \qquad \tilde{A} \equiv \mathsf{A}_0\,(\mathsf{k_R}/(\mathsf{k_F} + \mathsf{k_R}))$$
$$(P \to \langle\alpha\rangle Q) \equiv (\mathsf{const} \land \mathsf{A} = \mathsf{A}_0 \land \mathsf{B} = 0 \to \langle \mathsf{ode}\rangle[\mathsf{ode}]\mathsf{A} \leq \tilde{A} + \epsilon)$$

Fig. 9. Reversible reaction is persistent.

Proof. We combine proof techniques, first showing we eventually approach the equilibrium (variant reasoning, as in Fig. 8), then showing the concentration of A stays near the equilibrium (invariant reasoning, as in Fig. 6).

A major strength of logic is *compositionality*: complex proofs are but combinations of simple parts. For example, our dL proof of form $pre \to \langle\alpha\rangle[\alpha]P$ (call this formula D for short) can be divided into a variant proof and invariant proof, respectively proofs of some formulas of form $B \equiv pre \to \langle\alpha\rangle A$ and $C \equiv \mathsf{const} \land A \to [\alpha]P$ for some A. At a high level, KeYmaera X lived up to its compositionality promise, but at a low level, there is always room for improvement. The differential variant rule only allows inequalities as postconditions, but C expects $\mathsf{const} \land A$. We bridge this gap using the `mond` rule and `Kd2` axiom:

$$\textsc{Kd2}\ [\alpha]P \to \langle\alpha\rangle Q \to \langle\alpha\rangle(P \land Q) \qquad\qquad \textsc{mond}\ \frac{P \vdash Q}{\langle\alpha\rangle P \vdash \langle\alpha\rangle Q}$$

Invariants prove C. Applying `mond` on C yields $\langle\alpha\rangle(\mathsf{const} \land A) \to \langle\alpha\rangle[\alpha]P$. Prove the left side by `Kd2`. Its first premise holds by vacuity because const is constant; its second is by lemma B, which holds by a variant argument. The result is D, as desired.

Lessons for KeYmaera X Development. To our knowledge, the limitation to inequalities in differential variants is not fundamental, but incidental to KeYmaera X's implementation. We recommend that the developers relax this limitation. More generally, we found ourselves manually proving properties of the form $const \to \langle\alpha\rangle const$ where α does not modify free variables of *const*. Such formulas only hold when α has a run (i.e., $\langle\alpha\rangle\mathsf{true}$ holds), thus are nontrivial to automate, yet still deserve attention because they are common. The `mond` rule and `Kd2` axiom were key to our proof, but are only visible on the UI when the user searches for them by name. We recommend that the developers provide visibility, either through the UI or through example proofs.

Tactics in KeYmaera X seek to enable concise proof scripts, so it is desirable to automate counting the size of proof scripts and underlying proof terms. To our knowledge, KeYmaera X's current support for size counting is experimental. We recommend that the developers promote size counting to a stable feature. Our proof scripts ranged from 3 to 41 proof steps, and experience suggests that a tactic-free proof would likely be much longer. This is consistent with results from the literature, where tens of lines of tactics may correspond to $>200,000$ steps [8, §4.1]. Our slowest proof completed in 8 seconds on a modern workstation. Model complexity and proof-checking time were not directly related: some

simple models ran slower than complex ones because simple models support the highest level of automation, but highly-automated proofs check more slowly than highly-interactive proofs.

In short, theorem-proving case studies are not only important because they demonstrate the benefits of new automation, but because they discover directions for future development.

4 Related Work

Related work includes hybrid systems verification, reactor design, and reaction kinetics. We begin with theorem-proving approaches to verification, specifically.

Hybrid Systems Theorem Proving. Specialized *hybrid systems* provers [16, 44] provide a high degree of generality (parametricity, nonlinearity, unbounded time) and rigor, while making efforts to mitigate the high degree of user effort typical of theorem-proving. For example, generality in our case study means many different reactions and reactors are supported by modifying parameter values, with no new proof effort. Rigor is not merely of theoretical interest: in many hybrid systems reasoning techniques which do not share our rigorous logical foundations, many soundness edge cases have recently been identified [41, Tab. 1]. Soundness violations are unacceptable in verification.

We use the KeYmaera X [16] prover for its exceptional rigor: its axioms have been proved sound in a theorem-prover [7] and it soundly derives its advanced proof methods [35][41, Tab. 1] from sound axioms.

Hybrid Hoare Logic (HHL) [22, 44] is another notable hybrid systems logic; an HHL case study similar to ours could be interesting future work. HHL Prover and KeYmaera X both base their ODE invariant automation on the same core algorithm [23], so this aspect of automation is likely comparable in both.

Other Logical Approaches. We are aware of only one prior logical proof [45] of a chemical process with nontrivial hybrid dynamics. Unlike ours, it is not in a theorem-prover and does not address persistence nor reactions, but rather a mixing process. General-purpose theorem-provers [1, 12, 26, 37] have formalized hybrid systems, including stability [26, 37], but not applied them to reactions.

Reachability. Model-checkers using reachability analysis [2, 9, 11, 14] are hybrid theorem-provers' main competitors. They increase automation, but have restrictions in generality. We discuss this tradeoff, which led us to use theorem-proving.

Foremost, KeYmaera X supports persistence. To our knowledge, persistence is not among the specific classes (e.g., safety and reach-avoid correctness) of properties supported in any model-checker. Logic allows mixing existential and universal properties freely, supporting broad classes of properties.

Secondly, model-checkers use *compact* regions, i.e., variables have finite bounds. In contrast, KeYmaera X allows non-compact *parametric* results. This enables arbitrarily large reaction and heating rates, timesteps, and tank capacities.

Thirdly, we use multi-affine ODEs. Many model-checkers support multi-affine ODEs [2, 5, 9, 11], but struggle with scalability, compared to affine systems [3, 14]. Our small-scale results potentially enable future scalability: multi-affine component-based proofs scale to hundreds of variables [6], an order of magnitude beyond nonlinear ODE benchmarks [5, 9, 11].

Theorem-provers benefit from strong correctness arguments. KeYmaera X's trusted computing base is an order of magnitude smaller than self-reported counts of popular model-checkers [16] and its axioms have a machine-checked soundness proof [7]. Correctness is not merely a theoretical concern. Soundness bugs in Flow* and dReach have been identified post-release [31]. Predecessors of techniques used in this paper had known soundness bugs [41, Table 1]. The model-checking community has acknowledged these concerns. Ariadne developers [10] have specifically cited the correctness benefits of theorem-proving for reachability. Developers of SpaceEx, PHAVer, HyTech, Lyse, and VNODE-LP [13, 29] have cited implementation correctness concerns for reachability analysis. KeYmaera X is typically preferred over paper proofs, because paper proofs would employ invariant and variant techniques with comparable complexity to our own, but sacrifice automatic detection of proof errors, which are common.

Theorem-proving's downside is the requirement for interactive proofs by users. Automation discussed herein only assists, not eliminates, interaction. In contrast, push-button automation is common for model-checkers. Due to these nontrivial tradeoffs, both theorem-proving and model-checking remain essential.

Stability and Persistence. Hybrid system stability is well-studied both inside [26, 37, 42] and outside [21, 24, 27] theorem-provers, with persistence also studied [39]. Lyapunov functions have shown stability of a chemical reaction on paper, but not in a prover [19]. Stability and its relatives in KeYmaera X specifically are a new topic [42]; ours is the first application-focused study in KeYmaera X.

Chemical Engineering. The chemical engineering results we formalized are classical; our innovation is the rigor and generality (parametricity, non-linearity) with which we formalize them in KeYmaera X. Standard textbooks provided kinetics for well-mixed adiabatic batch reactors [38, Eq. 2.93], uncontrolled reversible reactions [38, Ch. 3], and the Arrhenius equation [38, Eq. 5.1]. Standard control theory textbooks introduce model-predictive control and bang-bang control [18].

Although basic models of reactors are widely-used in formal methods, ours is the first in a theorem-prover. It additionally overcomes others' limitations:

- Previous chemical proofs ignored persistence and reactors [45]
- Optimal scheduling [36] and safety proofs [25] only used state machines
- A verified plant design used simple piecewise-constant dynamics [20]
- CEGAR verification of tanks [30] ignored reactors

Industrial usage of formal methods typically prioritizes optimal control and optimization of plant configuration, accepts approximations as a tradeoff for nonlinearity, and cites scalability to networks of reactions and changes in parameter values as common issues [43]. This paper provides a parametric model

that supports nonlinear dynamics through approximation, and formally proves the approximation correct against nonlinear dynamics. Because dL is amenable to constrained optimization for control [17] and efficient verification of compound systems by decomposition into reusable components [6], it is expected that the dL-based approach can be extended to overcome the aforementioned industrial challenges in future work. If successful, the benefits to the chemical industry would include increased confidence in software correctness and potential improvements in scalability and efficiency of parameter changes, when designing plants and controllers. Maximal realism would require direct access to industrial designs, but our proofs already demonstrate that improvements in ODE realism can often be accommodated with modest changes to invariants. For models beyond ODEs, such as PDE models of non-uniform heat transfer, differential games can be explored because they can express Hamilton-Jacobi-like PDEs [33]. Though industrial users do not frequently cite concerns regarding formalization of correctness proofs [43], they still stand to benefit from such guarantees because constructing chemical plants is expensive, making design mistakes costly.

5 Conclusion

We used the KeYmaera X theorem prover for differential dynamic logic to formalize two case studies: a batch reactor and a reversible reaction, each of which consisted of four models and their proofs. This work served two purposes:

- To our knowledge, we provided the first proof in a theorem prover of these classic chemical engineering results.
- We demonstrated how recent advances in KeYmaera X's automation, such as its implementation of invariant checking, Darboux reasoning, and differential variants, contribute to the proofs.

There are two directions of future work which could promote industrial impact. A component-based approach could compose the models and proofs for individual reactions into complete reaction networks or chemical plants. Previous proofs suggest a component-based approach could scale to hundreds of variables [6], indicating potential to improve upon the scalability of competing approaches [43]. Secondly, a black-box approach [8] incorporating constrained optimization [17] could make our work useful for realistic industrial controllers, which may involve components too complex for current white-box verification techniques. Our model could be used at runtime to check whether a complex controller's control decision is within a safe range; if not, our simple controller can be used as a safe fallback.

Acknowledgements. We thank the reviewers and Yong Kiam Tan for careful readings and feedback. We thank Therese Smith, Andrew Teixeira, and Grier Wallace for helpful discussions.

References

1. Ábrahám-Mumm, E., Steffen, M., Hannemann, U.: Verification of hybrid systems: Formalization and proof rules in PVS. In: ICECCS. IEEE (2001)
2. Althoff, M., Grebenyuk, D., Kochdumper, N.: Implementation of Taylor models in CORA 2018. In: ARCH. EPiC Series in Computing, vol. 54. EasyChair (2018)
3. Bak, S., Tran, H., Johnson, T.T.: Numerical verification of affine systems with up to a billion dimensions. In: HSCC. ACM (2019)
4. Bauer, N., Kowalewski, S., Sand, G., Löhl, T.: A case study: Multi product batch plant for the demonstration of control and scheduling problems. In: ADPM (2000)
5. Benvenuti, L., Bresolin, D., Collins, P., Ferrari, A., Geretti, L., Villa, T.: Assume–guarantee verification of nonlinear hybrid systems with Ariadne. Intl. J. Robust Nonlin. Control (2014)
6. Bohrer, R., Luo, A., Chuang, X.A., Platzer, A.: CoasterX: A case study in component-driven hybrid systems proof automation. In: ADHS. Elsevier (2018)
7. Bohrer, R., Rahli, V., Vukotic, I., Völp, M., Platzer, A.: Formally verified differential dynamic logic. In: CPP. ACM (2017)
8. Bohrer, R., Tan, Y.K., Mitsch, S., Myreen, M.O., Platzer, A.: VeriPhy: verified controller executables from verified cyber-physical system models. In: PLDI. ACM (2018)
9. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Flow*: An analyzer for non-linear hybrid systems. In: CAV. LNCS, vol. 8044. Springer (2013)
10. Collins, P., Niqui, M., Revol, N.: A Taylor function calculus for hybrid system analysis: Validation in Coq. In: NSV (2010)
11. Duggirala, P.S., Potok, M., Mitra, S., Viswanathan, M.: C2E2: a tool for verifying annotated hybrid systems. In: HSCC. ACM (2015)
12. Dupont, G., Ameur, Y.A., Singh, N.K., Pantel, M.: Event-B hybridation: A proof and refinement-based framework for modelling hybrid systems. ACM Trans. Embed. Comput. Syst. (2021)
13. Frehse, G., Giacobbe, M., Henzinger, T.A.: Space-time interpolants. In: CAV. LNCS, vol. 10981. Springer (2018)
14. Frehse, G., Guernic, C.L., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: SpaceEx: Scalable verification of hybrid systems. In: CAV. LNCS, vol. 6806. Springer (2011)
15. Fulton, N., Mitsch, S., Bohrer, R., Platzer, A.: Bellerophon: Tactical theorem proving for hybrid systems. In: ITP. LNCS, vol. 10499. Springer (2017)
16. Fulton, N., Mitsch, S., Quesel, J., Völp, M., Platzer, A.: Keymaera X: an axiomatic tactical theorem prover for hybrid systems. In: CADE. LNCS, vol. 9195. Springer (2015)
17. Fulton, N., Platzer, A.: Verifiably safe off-model reinforcement learning. In: TACAS. LNCS, vol. 11427. Springer (2019)
18. Glad, T., Ljung, L.: Control theory. CRC Press (2018)
19. Hangos, K.M.: Engineering model reduction and entropy-based Lyapunov functions in chemical reaction kinetics. Entropy (2010)
20. Hassapis, G., Kotini, I., Doulgeri, Z.: Validation of a SFC software specification by using hybrid automata. IFAC Proc. (1998)
21. Koutsoukos, X.D., He, K.X., Lemmon, M.D., Antsaklis, P.J.: Timed Petri nets in hybrid systems: Stability and supervisory control. Discret. Event Dyn. Syst. (1998)
22. Liu, J., Lv, J., Quan, Z., Zhan, N., Zhao, H., Zhou, C., Zou, L.: A calculus for hybrid CSP. In: APLAS. LNCS, vol. 6461. Springer (2010)

23. Liu, J., Zhan, N., Zhao, H.: Computing semi-algebraic invariants for polynomial dynamical systems. In: EMSOFT. ACM (2011)
24. Lozano, R., Fantoni, I., Block, D.J.: Stabilization of the inverted pendulum around its homoclinic orbit. Systems & Control Letters (2000)
25. Lukoschus, B.: Compositional verification of industrial control systems: methods and case studies. Ph.D. thesis, Christian-Albrechts Universität Kiel (2004)
26. Mitra, S., Chandy, K.M.: A formalized theory for verifying stability and convergence of automata in PVS. In: TPHOLs. LNCS, vol. 5170. Springer (2008)
27. Mitra, S., Liberzon, D.: Stability of hybrid automata with average dwell time: an invariant approach. In: CDC. IEEE (2004)
28. Mitsch, S., Platzer, A.: The KeYmaera X proof IDE: Concepts on usability in hybrid systems theorem proving. In: FIDE. EPTCS, vol. 240 (2016)
29. Nedialkov, N.S.: Implementing a rigorous ODE solver through literate programming. In: Modeling, Design, and Simulation of Systems with Uncertainties. Springer (2011)
30. Nellen, J., Ábrahám, E., Wolters, B.: A CEGAR tool for the reachability analysis of PLC-controlled plants using hybrid automata. In: FMI, AISC, vol. 346. Springer (2015)
31. Nguyen, L.V., Schilling, C., Bogomolov, S., Johnson, T.T.: Runtime verification for hybrid analysis tools. In: RV. LNCS, vol. 9333. Springer (2015)
32. Platzer, A.: The structure of differential invariants and differential cut elimination. Log. Meth. Comput. Sci. (2012)
33. Platzer, A.: Differential hybrid games. ACM Trans. Comput. Log. (2017)
34. Platzer, A.: Logical Foundations of Cyber-Physical Systems. Springer, Cham (2018)
35. Platzer, A., Tan, Y.K.: Differential equation invariance axiomatization. J. ACM (2020)
36. Potočnik, B., Bemporad, A., Torrisi, F.D., Mušič, G., Zupančič, B.: Hybrid modelling and optimal control of a multiproduct batch plant. Control Engineering Practice (2004)
37. Rouhling, D.: A formal proof in Coq of a control function for the inverted pendulum. In: CPP. ACM (2018)
38. Schmidt, L.D.: The Engineering of Chemical Reactions. Oxford University Press (1998)
39. Sogokon, A., Jackson, P.B., Johnson, T.T.: Verifying safety and persistence properties of hybrid systems using flowpipes and continuous invariants. In: NFM. Springer (2017)
40. Stephanopoulos, G.: Chemical Process Control: An Introduction to Theory and Practice. Prentice-Hall (1984)
41. Tan, Y.K., Platzer, A.: An axiomatic approach to existence and liveness for differential equations. Formal Aspects Comput. (2021)
42. Tan, Y.K., Platzer, A.: Deductive stability proofs for ordinary differential equations. In: TACAS. LNCS, vol. 12652. Springer (2021)
43. Tsay, C., Pattison, R.C., Piana, M.R., Baldea, M.: A survey of optimal process design capabilities and practices in the chemical and petrochemical industries. Computers & Chemical Engineering (2018)
44. Wang, S., Zhan, N., Zou, L.: An improved HHL prover: An interactive theorem prover for hybrid systems. In: ICFEM. LNCS, vol. 9407. Springer (2015)
45. Xu, Q., He, W.: Hierarchical design of a chemical concentration control system. In: Hybrid Systems. LNCS, vol. 1066. Springer (1995)