

Chemical Case Studies in KeYmaera X

Rose Bohrer^[0000–0001–5201–9895]

Worcester Polytechnic Institute, Worcester MA 01609, USA
rose.bohrer.cs@gmail.com

Abstract. Safety-critical chemical processes are the backbone of multi-billion-dollar industries, thus society deserves the strongest possible guarantees that they are safe. To that end, models of 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 the KeYmaera X theorem-prover to verify chemical models with differential dynamic logic. Our case studies are novel in combining the following: we provide strong general-case correctness theorems, use particularly rich hybrid dynamics, and have particularly rigorous proofs. This novel combination is made possible by KeYmaera X. Simultaneously, 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

Modern industry relies critically on all kinds of chemical processes: some occur in computer-controlled reactors, some occur free of control. Chemical engineering has provided many classical insights about both: safe and optimal control [12] of reactors [32] is a field in its own right, as are reaction kinematics (dynamics) even in the absence of control [30].

Because both controlled and uncontrolled reactions are crucial, we consider both: an irreversible exothermic reaction with a model-predictive bang-bang controller (Sec. 3.1) and an uncontrolled reversible reaction (Sec. 3.2). Both have verification challenges which make for good benchmark problems. The non-reversible reaction’s nuanced dynamics entail nontrivial correctness arguments for model-predictive controllers. The reversible reaction’s long-term asymptotic behavior, though classic, tests the ability of current-generation tools to verify asymptotic properties, e.g., stability [22] or persistence [31].

Safe reactions are crucial to human safety. Properties like persistence, stability, and optimality are crucial to human productivity. Thus, formal methods for chemical reactions are extensively studied [3,28,20,14,24].

To our knowledge, however, the reaction models and proofs presented here are the first-ever in a *hybrid systems theorem prover*. Specifically, we use the KeYmaera X [11] prover for *differential dynamic logic* (dL) [26] to achieve a unique

combination of expressive dynamics, general-case guarantees, and rigor for the first time. The tradeoffs between theorem-proving and other formal methods are well-known; see Sec. 4 for detailed discussion.

Our contribution was enabled by new stability [33], variant [33], and Darboux polynomial [27] proof tools in KeYmaera X, simplifying our proof arguments. Our case studies make essential use of these features and thus demonstrate the impacts of the latest advances in proof automation.

2 Background

All our proofs are computer-checked in the KeYmaera X prover, which carefully prevents the use of unsound reasoning [5]. This rigor is crucial in practice: many techniques used here had predecessors [33, Table 1] which were found to be unsound, which is unacceptable for safety-critical systems.

In KeYmaera X, correctness properties are stated and proved in *differential dynamic logic* (dL) [26], where hybrid systems are written in *hybrid program* notation. We discuss dL, then KeYmaera X usage.

2.1 Differential Dynamic Logic

We provide a primer on dL syntax and semantics (meaning); see the literature [26] for details. Semantics are state-based: a state ω maps every variable x to a real-number value $\omega(x) : \mathbb{R}$. The syntax consists of terms (with a numeric meaning in each state), hybrid programs (which can nondeterministically change the state when run), and formulas (which are true or false in each state). Hybrid programs and formulas may both contain each other. We use standard notation to define syntax, e.g., $B ::= C \mid D$ means every B is either a C or a D .

Definition 1 (Terms). *Terms e, \tilde{e} of dL are defined by:*

$$e, \tilde{e} ::= q \mid x \mid e + \tilde{e} \mid e \cdot \tilde{e} \quad \text{where } q \in \mathbb{Q}$$

Rational-valued literal numbers are written q . Real-valued variables are written x . Sum $e + \tilde{e}$ is the sum of terms e and \tilde{e} . Product $e \cdot \tilde{e}$ is the product of e and \tilde{e} . In every state, the meaning of every term is some real number.

Definition 2 (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. 3 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. If desired, standard conditional and looping constructs are derivable (where P is a formula, $\neg P$ is its negation, and α is a hybrid program):

$$\begin{aligned} \text{if}(P)\{\alpha\}\text{else}\{\beta\} &\equiv \{?P; \alpha\} \cup \{?\neg P; \beta\} \\ \text{while}(P)\{\alpha\} &\equiv \{?P; \alpha\}^*; ?\neg P \end{aligned}$$

Definition 3 (Formulas). *There are many formulas P, Q in dL. We only use:*

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

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* of α starting from state ω ends in a state where P is true.

When α is an ODE, *all runs* equates to *all time*, e.g., these readings apply:

- $P \rightarrow [\alpha]Q$ assumes P at first, then proves Q forever
- $P \rightarrow \langle \alpha \rangle Q$, assumes P at first, then proves Q eventually
- $P \rightarrow \langle \alpha \rangle [\alpha]Q$, assumes P at first, then proves Q eventually becomes true, then stays true forever.

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

Definition 4 (Validity). *A dL formula is valid if it is true in every state.*

We use standard notation for axioms and proof rules.

Definition 5 (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 the line. Rules can use schema variables such as P or α when the rule applies to all programs or formulas, respectively.*

For example, the loop rule

$$\text{LOOP} \frac{P \rightarrow J \quad J \rightarrow [\alpha]J \quad J \rightarrow Q}{P \rightarrow [\alpha^*]Q}$$

means for all P, Q, J, α that if premises $P \rightarrow J$, $J \rightarrow [\alpha]J$, and $J \rightarrow Q$ are all valid, so is $P \rightarrow [\alpha^*]Q$. 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 briefly discuss the user experience of KeYmaera X [23]. The user interface is displayed in Fig. 1. KeYmaera X is an interactive, tactic-based prover. This means that the user interactively tells the prover which proof technique to use, but each technique is implemented as a *tactic* [10], i.e., a program. A proof technique can be a simple, specific rule or a complex proof search procedure. For example, there is a *default* (or *auto*) proof procedure which attempts many proof techniques and can solve many simpler problems fully-automatically. In summary, the amount of user effort can vary greatly between proofs. Throughout this paper, we will discuss the level of interaction needed for each proof and discuss how new rules and automation helped keep the level of user effort manageable. The tactic-based approach also means that no matter how complex proof methods are, they are implemented using simple steps from the small trusted core of the prover, thus proofs stay rigorous.

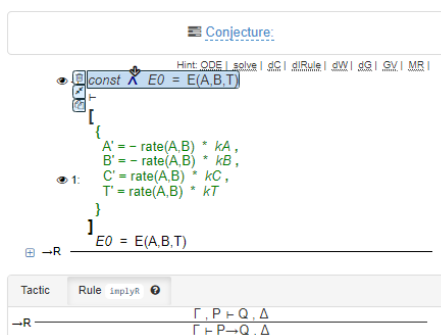



Fig. 1. KeYmaera X screenshot. Clicking the highlighted symbol performs a proof step. The last proof rule is shown at bottom. Recommended proof steps are displayed as hints.

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, which we chose because batch reactors [30, §2.10] are a foundational technology for chemical plants throughout industry. 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 reversible reactions because they too are essential to industry. Notably, ammonia synthesis is a reversible reaction that provides the backbone for modern fertilizer-based, industrial-scale agriculture [16]. Both case studies emphasize recent advances in KeYmaera X proof automation, which contributed to highly general results. Remaining limits on generality are discussed in each subsection.

3.1 Controlled Irreversible Reactions

We formalize a classic scenario: an irreversible, exothermic reaction in an adiabatic, well-mixed batch reactor. *Irreversible* [30, §2.1] means the reaction is one-way: outputs do not react to create inputs. *Adiabatic* [30, §2.14] means heat does not leave or enter the reactor. *Well-mixed* [30, §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:



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 [12], 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 depends linearly on temperature, changing exponentially with respect to 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. We discuss why, including verification challenges.

Each model approximates textbook [30, 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 $\text{rate} = kAB$ where k is an exponential given by the Arrhenius equation [30, 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 richer than prior models [36,28,20,14,24]. 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. Notwithstanding these limitations, the results are fully general in the sense that they are fully parametric, e.g., the results can be applied to *any* 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.

$$\begin{aligned}
E &\equiv KE + U & U &\equiv \min(A/k_A, B/k_B) k_T & KE &\equiv T & \text{rate} &\equiv T_s A_0 B_0 k_{ra} + k_{rb} \\
\text{const} &\equiv k_{ra} > 0 \wedge k_{rb} \geq 0 \wedge k_A > 0 \wedge k_B > 0 \wedge k_C > 0 \wedge k_T > 0 \\
\text{ode} &\equiv \{A' = -\text{rate } k_A, B' = -\text{rate } k_B, C' = \text{rate } k_C, T' = \text{rate } k_T\} \\
(P \rightarrow [\alpha]Q) &\equiv (\text{const} \wedge E_0 = E \rightarrow [\text{ode}]E_0 = E)
\end{aligned}$$

Fig. 2. Basic irreversible model conserves energy

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 kinetic (heat) and potential (chemical) energy: $E \equiv KE + U$. Potential energy $U \equiv \min(A/k_A, B/k_B) k_T$ 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 $T_s A_0 B_0 k_{ra} + k_{rb}$, which makes two intentional simplifications. First, we use approximate *current* concentrations A, B with *initial* concentrations A_0, 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 k_{ra}, k_{rb} let the rate be any *linear function of* the product. Formula *const* simply specifies the signs of constants.

The *ode* indicates that all concentrations A, B, 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 k_A, k_B, k_C, k_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 $(P \rightarrow [\alpha]Q)$ 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 proof of the theorem in KeYmaera X.

Proof. The default proof procedure of KeYmaera X (Sec. 2.2) proves the theorem automatically with *differential invariants* [26, Lem. 11.3], demonstrating the capabilities of this standard dL rule. We present (the relevant case of) *differential invariant* [26, Lem. 11.3] rule

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

which shows $e = \tilde{e}$ is true throughout an ODE if it holds initially and differentials are equal throughout. We prove $E_0 = E$ thus: E_0 is constant, so proving $E' = 0$ throughout suffices. Expanding the definition of E yields $(E)' = (T + \min(A/k_A, B/k_B) k_T)' = \text{rate} k_T + \min((A)'/k_A, (B)'/k_B) k_T = \text{rate} k_T + \min(-\text{rate} k_A/k_A, -\text{rate} k_B/k_B) k_T = \text{rate} k_T + \min(-\text{rate}, -\text{rate}) k_T = (\text{rate} - \text{rate}) k_T = 0$. 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.

$$\begin{aligned}
 \text{rate} &\equiv T_s A_0 B_0 k_{ra} + k_{rb} \\
 \text{const} &\equiv k_{ra} > 0 \wedge k_{rb} \geq 0 \wedge k_A > 0 \wedge k_B > 0 \wedge k_C > 0 \wedge k_T > 0 \wedge T > 0 \wedge \epsilon > 0 \\
 \text{ctrl} &\equiv \{\text{if}(T_{max} - T \leq \epsilon \text{rate} k_R)\{\text{isOn} := 0\}\text{else}\{\text{isOn} := 1\}\}; t := 0 \\
 \text{ode} &\equiv \{A' = \text{isOn} \cdot -\text{rate} k_A, B' = \text{isOn} \cdot -\text{rate} k_B, C' = \text{isOn} \cdot \text{rate} k_C, \\
 &\quad T' = \text{isOn} \cdot \text{rate} k_T, t' = 1 \wedge t \leq \epsilon \wedge A \geq 0 \wedge B \geq 0 \wedge C \geq 0\} \\
 (P \rightarrow [\alpha]Q) &\equiv (\text{const} \wedge T \leq T_{max} \rightarrow [\{\text{ctrl}; \text{ode}\}^*] T \leq T_{max})
 \end{aligned}$$

Fig. 3. Bang-bang irreversible model safe

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 so, the reaction shuts off (`isOn := 0`), else it turns on (`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 $\epsilon \text{rate} k_R$ 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 $A \geq 0 \wedge B \geq 0 \wedge C \geq 0$ model the physical assumption that concentrations cannot be negative. For example, the reaction would end if A or B reach zero.

Finally, the updated theorem statement $(P \rightarrow [\alpha]Q)$ 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 $T \leq T_{max}$ 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 following cut:

- $T_{max} - T > (\epsilon - t) \text{rate } k_T$, meaning the remaining safe temperature gap exceeds the projected temperature change during the remaining time.

The cut proves automatically by differential invariant, 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 simply other aspects of the reaction rate formula by dropping k_{ra} and k_{rb} . The remaining changes follow from that one: *amts* is a helper definition to definitions such as $\text{taylor}^+(x, t)$, which is an upper bound on temperature over time, constructed as a Taylor series approximation. This use of a Taylor series approximation represents a fundamental change in proof approach for a fundamentally more complicated dynamics: for exponential dynamics, polynomial approximations are a crucial tool to simplify reasoning. However, this Taylor bound is only provably an upper bound on a limited time interval which happens to be $1/(2 \text{amts})$, 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.

Proof. The loop invariant is unchanged. We add several differential cuts; order matters since each one can serve as an assumption in following proofs:

- $t \geq 0$ just means time moves forward,
- $A_0 B_0 T k_T \geq 0$ ensures forward reaction rate, and
- $\text{taylor}^+(T_{old}, t) - T \geq 0$ bounds temperature T above with $\text{taylor}^+(\cdot)$ in terms of old temperature T_{old} .

The final cut requires advanced proof techniques because term $\text{taylor}^+(T_{old}, t) - T$ decreases; differential invariants alone are provably [25, Thm 6.1] insufficient for such terms. The earliest suitable techniques required defining new (*ghost*,

$$\begin{aligned}
\text{rate} &\equiv \mathsf{T} \mathsf{A}_0 \mathsf{B}_0 \quad \epsilon \equiv 1/(2 \text{amts}) \quad \text{amts} \equiv k_{\mathsf{T}} \mathsf{A}_0 \mathsf{B}_0 \quad \text{taylor}^+(x, t) \equiv (1 + 2 t \text{amts}) x \\
\text{const} &\equiv k_A > 0 \wedge k_B > 0 \wedge k_C > 0 \wedge k_{\mathsf{T}} > 0 \wedge \epsilon > 0 \wedge \mathsf{A}_0 \geq 0 \wedge \mathsf{B}_0 \geq 0 \\
\text{ctrl} &\equiv \{\text{if}(\mathsf{T}_{max} \leq \text{taylor}^+(\mathsf{T}, \epsilon))\{\text{isOn} := 0\}\text{else}\{\text{isOn} := 1\}\}; t := 0 \\
\text{ode} &\equiv \{A' = \text{isOn} \cdot -\text{rate } k_A, B' = \text{isOn} \cdot -\text{rate } k_B, C' = \text{isOn} \cdot \text{rate } k_C, \\
&\quad \mathsf{T}' = \text{isOn} \cdot \text{rate } k_{\mathsf{T}}, t' = 1 \wedge t \leq \epsilon \wedge A \geq 0 \wedge B \geq 0 \wedge C \geq 0 \wedge \mathsf{T} \geq 0\} \\
(P \rightarrow [\alpha]Q) &\equiv (\text{const} \wedge \mathsf{T} > 0 \wedge \mathsf{T} \leq \mathsf{T}_{max} \wedge A = \mathsf{A}_0 \wedge B = \mathsf{B}_0 \rightarrow [\{\text{ctrl}; \text{ode}\}^*] \mathsf{T} \leq \mathsf{T}_{max})
\end{aligned}$$

Fig. 4. Bang-bang irreversible model safe with fixed exponent

or *auxiliary*) variables in each proof, but constructing suitable definitions can be non-obvious in practice. Fortunately, KeYmaera X has provided proof rules based on Darboux polynomial (inequality) reasoning [27, Corr. 3.2] which can prove the same problems, but are higher-level:

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

Here, both instances of \succsim are replaced uniformly with one of $>$ or \geq , where $(e)'$ is the differential of e , for polynomials p, g where p is called a *Darboux* polynomial if the premise holds and 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 gp are 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 [27, Thm. 4.5], a large class of invariants.

Darboux-based rules are complete for large classes of theorems, yet it is challenging to automatically find suitable polynomials in every case. For our example model, KeYmaera X did not find a suitable polynomial, but performing algebra by hand did result in a suitable polynomial: using the definition of the ODE, solve for a polynomial that satisfies the proof goal, in this case: $g \equiv \mathsf{A}_0 \mathsf{B}_0 k_{\mathsf{T}}$. 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, below, 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 [30, 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 non-trivial heating dynamics *and* nontrivial 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 A_1, 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.

```

rate  $\equiv T A B \quad \epsilon \equiv 1/(2 A_1 B_1 k_T) \quad \text{amts} \equiv A B k_T \quad \text{ctrl}$ 
const  $\equiv k_A > 0 \wedge k_B > 0 \wedge k_C > 0 \wedge k_T > 0 \wedge \epsilon > 0 \wedge A_0 \geq 0 \wedge B_0 \geq 0$ 
ctrl  $\equiv \{\text{if}(T_{max} \leq \text{taylor}^+(T, \epsilon))\{\text{isOn} := 0\}\text{else}\{\text{isOn} := 1\}\}; t := 0; A_1 := A; B_1 := B$ 
ode  $\equiv \{A' = \text{isOn} \cdot -\text{rate } k_A, B' = \text{isOn} \cdot -\text{rate } k_B, C' = \text{isOn} \cdot \text{rate } k_C,$ 
 $T' = \text{isOn} \cdot \text{rate } k_T, t' = 1 \wedge t \leq \epsilon \wedge A \geq 0 \wedge B \geq 0 \wedge C \geq 0 \wedge T \geq 0\}$ 
( $P \rightarrow [\alpha]Q$ )  $\equiv (\text{const} \wedge T > 0 \wedge T \leq T_{max} \wedge A = A_0 \wedge B = B_0 \rightarrow [\{\text{ctrl}; \text{ode}\}^*]T \leq T_{max})$ 

```

Fig. 5. Bang-bang irreversible model safe with dynamic exponent

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: $0 \leq T \wedge T \leq T_{max} \wedge A \leq A_0 \wedge B \leq B_0$. The differential cuts are similar to before, with an additional lemma that the concentrations of the first two reactants decrease: $A \leq A_1 \wedge A_1 \leq A_0 \wedge B \leq B_1 \wedge B_1 \leq B_0$. The differential cut for the Taylor series is unchanged, and the same Darboux polynomial $g \equiv A_0 B_0 k_T$ suffices.

3.2 Uncontrolled Reversible Reactions

We study reversible reactions, which are crucial to society. For example, ammonia synthesis is critical to modern agriculture [16]. We consider a textbook scenario where two reactants A and B can each react to form the other:



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* [31], 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 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 where we start 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 in its own right, e.g., to verify that a container never overflows, but we mainly use the lemma as a building block for persistence. Here, the two reactants are named A and B , with initial values $A = A_0 > 0$ and

$$\begin{aligned} \text{ode} &\equiv \{A' = -A k_F + B k_R, B' = A k_F - B k_R\} \\ \text{const} &\equiv A_0 > 0 \wedge k_R > 0 \wedge k_F > 0 \\ (P \rightarrow [\alpha]Q) &\equiv (\text{const} \wedge A = A_0 \wedge B = 0 \rightarrow [\text{ode}]A \leq A_0) \end{aligned}$$

Fig. 6. Reversible model decreases A

$B = 0$. 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 [30, 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 \tilde{A} and \tilde{B} of A and B present at the equilibrium. The above definitions of \tilde{A} and \tilde{B} can be found by solving for equilibrium ($A' = 0 \wedge B' = 0$) in `ode` subject to conservation of mass ($A + B = A_0$).

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.

- $A - A_0 (k_R / (k_F + k_R)) > 0$ means A 's rate of change is always in the direction of the equilibrium
- $A + B = A_0$ is conservation of mass

$$\begin{aligned}
\text{ode} &\equiv \{A' = -A k_F + B k_R, B' = A k_F - B k_R\} \\
\text{const} &\equiv A_0 > 0 \wedge k_R > 0 \wedge k_F > 0 \\
\tilde{A} &\equiv A_0 (k_R / (k_F + k_R)) \quad \tilde{B} \equiv A_0 (k_F / (k_F + k_R)) \\
(P \rightarrow [\alpha]Q) &\equiv (\text{const} \wedge A = A_0 \wedge B = 0 \rightarrow [\text{ode}]A \neq \tilde{A})
\end{aligned}$$

Fig. 7. Reversible model never at equilibrium

- $A > 0 \wedge B \geq 0$ means we never have a negative amount of either reactant, the first being positive. This requires a Darboux argument with polynomial $-(k_F + k_R)$ because the amount of the first reactant does decrease with time.

Once these cuts are proved, 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 $\langle \text{ode} \rangle A \leq \tilde{A} + \epsilon$ 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.

$$\begin{aligned}
\text{const} &\equiv A_0 > 0 \wedge k_R > 0 \wedge k_F > 0 \wedge \epsilon > 0 \\
\text{ode} &\equiv \{A' = -A k_F + B k_R, B' = A k_F - B k_R\} \\
\tilde{A} &\equiv A_0 (k_R / (k_F + k_R)) \quad \tilde{B} \equiv A_0 (k_F / (k_F + k_R)) \\
(P \rightarrow \langle \alpha \rangle Q) &\equiv (\text{const} \wedge A = A_0 \wedge B = 0 \rightarrow \langle \text{ode} \rangle A \leq \tilde{A} + \epsilon)
\end{aligned}$$

Fig. 8. Reversible model approaches equilibrium

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. A *differential variant* proof is the diamond counterpart to *differential invariant* reasoning for box properties. The *differential variant* principle [33, Corr. 24] says: if there is a lower bound on the rate of progress we make toward our goal at all times, we will get there eventually.

$$\text{dV} \stackrel{\exists}{\succ} \frac{\exists d > 0 \forall x (\neg(p \geq 0) \rightarrow (p)' \geq d)}{\langle \{x' = f(x)\} \rangle p \succ 0}$$

where \succ stands 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).

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: $d = \epsilon (k_F + k_R)$. 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 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:

$$\begin{aligned} \text{const} &\equiv A_0 > 0 \wedge k_R > 0 \wedge k_F > 0 \wedge \epsilon > 0 \\ \text{ode} &\equiv \{A' = -A k_F + B k_R, B' = A k_F - B k_R\} \\ \tilde{A} &\equiv A_0 (k_R / (k_F + k_R)) \quad \tilde{B} \equiv A_0 (k_F / (k_F + k_R)) \\ (P \rightarrow \langle \alpha \rangle Q) &\equiv (\text{const} \wedge A = A_0 \wedge B = 0 \rightarrow \langle \text{ode} \rangle [\text{ode}] A \leq \tilde{A} + \epsilon) \end{aligned}$$

Fig. 9. Reversible model is stable.

Proof. We combine proof techniques, first showing we eventually approach the equilibrium (variant reasoning), then showing the concentration of A never increases again (invariant reasoning).

A major strength of logic is *compositionality*: complex proofs are but combinations of simple parts. A dL proof of $\langle \alpha \rangle [\alpha] P$ can be divided into a variant proof and invariant proof, for example. At a high level, KeYmaera X lived up to this compositionality promise. At a low level, there is always room for improvement: the $[\alpha] P$ proof assumes *const*, i.e., it assumes constants never change. Due to limitations of the differential variant rule, we had to prove the constants never change, albeit with a simple proof. The limitation appears incidental to KeYmaera X's implementation, not fundamental. It speaks well of the implementation used in these case studies that this was the only instance where the automation added new proof challenges. This serves as a reminder that theorem-proving case studies are dually important, showing both the gains from new automation and which features deserve future optimization.

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* theorem-provers [11,35] provide a high degree of generality 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 [33, Tab. 1]. Soundness violations are unacceptable in verification.

We use the KeYmaera X [11] prover for its exceptional rigor: its axioms have been proved sound in a theorem-prover [5] and it soundly derives its advanced proof methods [33,27, Tab. 1] from sound axioms.

Hybrid Hoare Logic (HHL) [17,35] is another notable hybrid prover; 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 [18], so this aspect of automation is likely comparable in both.

Other Logical Approaches We are aware of only one prior logical proof [36] 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 [1,8,21,29] have formalized hybrid systems, including stability [29,21], but not applied them to reactions.

Reachability Model-checkers based on reachability analysis [6,2,7,9] are the primary competitors to hybrid systems theorem-provers. They provide greater automation at the cost of accepting restrictions in generality. Details vary, but common restrictions include special-case guarantees (is a *specific* reaction safe?), time-bounded analyses (am I safe *for a time*?) or conservative approximations of dynamics. Their trusting computing base is typically larger than a theorem-prover’s, complicating rigor.

Taylor approximations, particularly Taylor models [4], are broadly useful in reachability analysis, e.g., in Flow* [6] and CORA [2]. We have shown that Taylor approximations are equally useful in KeYmaera X, where they come with proofs.

Stability and Persistence Hybrid system stability is well-studied both inside [34,21,29] and outside [15,22,19] theorem-provers, with persistence also studied [31]. Lyapunov functions have shown stability of a chemical reaction on paper, but not in a prover [13]. Stability and its relatives in KeYmaera X specifically are a new topic [34]; we contribute the first worked KeYmaera X case study for an application of industrial interest.

Chemical Engineering. The chemical engineering results we formalized are classical; our innovation is the generality and rigor with which we formalize them in KeYmaera X. Standard textbooks provided kinetics for well-mixed adiabatic batch reactors [30, Eq. 2.93], uncontrolled reversible reactions [30, Ch. 3], and the Arrhenius equation [30, Eq. 5.1]. Standard control theory textbooks introduce model-predictive control and bang-bang control [12].

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 [36]
- Optimal scheduling [28] and safety arguments [20] have used simplistic finite state machines
- A verified plant design used simple piecewise-constant dynamics [14]
- CEGAR verification of tanks [24] ignored reactors

Though we build on such broad related work, our contribution of general-purpose proofs about chemical reactors and reactions in a theorem-prover fills a significant gap in the verification literature.

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 provide the first proof in a theorem prover of these foundational chemical engineering results
- We demonstrate how recent advances in KeYmaera X's automation, such as its implementation invariant checking, Darboux reasoning, and differential variants, contribute to the proofs

One direction for future work is verifying reactors with more advanced controllers such as PID (proportional-integral-derivative) controllers [32, Ch. 13]. However, potential future work is broad in nature, reaching well beyond chemical reactor design. Techniques such as invariant checking and Taylor series are of general applicability using various tools, though KeYmaera X provides a rigorous implementation of both. Differential variants are widely useful for proving ODE properties that are true eventually, but not at every moment. We have shown one significant application for all these proof techniques; their are certainly others because the applications of hybrid systems models are diverse.

References

1. Ábrahám-Mumm, E., Steffen, M., Hannemann, U.: Verification of hybrid systems: Formalization and proof rules in PVS. In: ICECCS. IEEE (2001). <https://doi.org/10.1109/ICECCS.2001.930163>

2. Althoff, M., Grebenyuk, D., Kochdumper, N.: Implementation of Taylor models in CORA 2018. In: ARCH. EPiC Series in Computing, vol. 54. EasyChair (2018). <https://doi.org/10.29007/zzc7>
3. 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)
4. Berz, M., Hoffstätter, G.: Computation and application of Taylor polynomials with interval remainder bounds. *Reliab. Comput.* **4**(1) (1998). <https://doi.org/10.1023/A:1009958918582>
5. Bohrer, R., Rahli, V., Vukotic, I., Völpl, M., Platzer, A.: Formally verified differential dynamic logic. In: CPP. ACM (2017). <https://doi.org/10.1145/3018610.3018616>
6. Chen, X., Abraham, E., Sankaranarayanan, S.: Flow*: An analyzer for non-linear hybrid systems. In: Sharygina, N., Veith, H. (eds.) CAV. LNCS, vol. 8044. Springer (2013). https://doi.org/10.1007/978-3-642-39799-8_18
7. Duggirala, P.S., Potok, M., Mitra, S., Viswanathan, M.: C2E2: a tool for verifying annotated hybrid systems. In: HSCC. ACM (2015). <https://doi.org/10.1145/2728606.2728646>
8. Dupont, G., Ameer, Y.A., Singh, N.K., Pantel, M.: Event-b hybridization: A proof and refinement-based framework for modelling hybrid systems. *ACM Trans. Embed. Comput. Syst.* **20**(4) (2021). <https://doi.org/10.1145/3448270>
9. 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). https://doi.org/10.1007/978-3-642-22110-1_30
10. Fulton, N., Mitsch, S., Bohrer, R., Platzer, A.: Bellerophon: Tactical theorem proving for hybrid systems. In: Ayala-Rincón, M., Muñoz, C.A. (eds.) ITP. LNCS, vol. 10499. Springer (2017). https://doi.org/10.1007/978-3-319-66107-0_14
11. Fulton, N., Mitsch, S., Quesel, J., Völpl, M., Platzer, A.: Keymaera X: an axiomatic tactical theorem prover for hybrid systems. In: CADE. LNCS, vol. 9195. Springer (2015). https://doi.org/10.1007/978-3-319-21401-6_36
12. Glad, T., Ljung, L.: Control theory. CRC Press (2018)
13. Hangos, K.M.: Engineering model reduction and entropy-based Lyapunov functions in chemical reaction kinetics. *Entropy* **12**(4) (2010). <https://doi.org/10.3390/e12040772>
14. Hassapis, G., Kotini, I., Douglgeri, Z.: Validation of a SFC software specification by using hybrid automata. *IFAC Proc.* **31**(15) (1998)
15. 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.* **8**(2) (1998). <https://doi.org/10.1023/A:1008293802713>
16. Liu, H.: Ammonia synthesis catalyst 100 years: Practice, enlightenment and challenge. *Chinese Journal of Catalysis* **35**(10), 1619–1640 (2014)
17. 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). https://doi.org/10.1007/978-3-642-17164-2_1
18. Liu, J., Zhan, N., Zhao, H.: Computing semi-algebraic invariants for polynomial dynamical systems. In: Chakraborty, S., Jerraya, A., Baruah, S.K., Fischmeister, S. (eds.) EMSOFT. ACM (2011). <https://doi.org/10.1145/2038642.2038659>
19. Lozano, R., Fantoni, I., Block, D.J.: Stabilization of the inverted pendulum around its homoclinic orbit. *Systems & Control Letters* **40**(3) (2000)
20. Lukoschus, B.: Compositional verification of industrial control systems: methods and case studies. Ph.D. thesis, Christian-Albrechts Universität Kiel (2004)

21. Mitra, S., Chandy, K.M.: A formalized theory for verifying stability and convergence of automata in PVS. In: TPHOLS. LNCS, vol. 5170. Springer (2008). https://doi.org/10.1007/978-3-540-71067-7_20
22. Mitra, S., Liberzon, D.: Stability of hybrid automata with average dwell time: an invariant approach. In: CDC. IEEE (2004). <https://doi.org/10.1109/CDC.2004.1430238>
23. Mitsch, S., Platzer, A.: The KeYmaera X proof IDE: Concepts on usability in hybrid systems theorem proving. In: Dubois, C., Masci, P., Méry, D. (eds.) FIDE. EPTCS, vol. 240, pp. 67–81 (2016). <https://doi.org/10.4204/EPTCS.240.5>
24. Nellen, J., Ábrahám, E., Wolters, B.: A CEGAR tool for the reachability analysis of plc-controlled plants using hybrid automata. In: Bouabana-Tebibel, T., Rubin, S.H. (eds.) FMI, AISC, vol. 346. Springer (2015). https://doi.org/10.1007/978-3-319-16577-6_3
25. Platzer, A.: The structure of differential invariants and differential cut elimination. *Logical Methods in Computer Science* **8**(4) (2012). [https://doi.org/10.2168/LMCS-8\(4:16\)2012](https://doi.org/10.2168/LMCS-8(4:16)2012)
26. Platzer, A.: *Logical Foundations of Cyber-Physical Systems*. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-63588-0>
27. Platzer, A., Tan, Y.K.: Differential equation invariance axiomatization. *J. ACM* **67**(1) (2020). <https://doi.org/10.1145/3380825>
28. Potočník, B., Bemporad, A., Torrisi, F.D., Mušič, G., Zupančič, B.: Hybrid modelling and optimal control of a multiproduct batch plant. *Control Engineering Practice* **12**(9) (2004)
29. Rouhling, D.: A formal proof in Coq of a control function for the inverted pendulum. In: CPP. ACM (2018). <https://doi.org/10.1145/3167101>
30. Schmidt, L.D.: *The engineering of chemical reactions* (1998)
31. Sogokon, A., Jackson, P.B., Johnson, T.T.: Verifying safety and persistence properties of hybrid systems using flowpipes and continuous invariants. In: NFM. pp. 194–211. Springer (2017)
32. Stephanopoulos, G.: *Chemical Process Control: An Introduction to Theory and Practice*. Prentice-Hall (1984)
33. Tan, Y.K., Platzer, A.: An axiomatic approach to existence and liveness for differential equations. *Formal Aspects Comput.* **33**(4) (2021). <https://doi.org/10.1007/s00165-020-00525-0>
34. Tan, Y.K., Platzer, A.: Deductive stability proofs for ordinary differential equations. In: Groote, J.F., Larsen, K.G. (eds.) TACAS. LNCS, vol. 12652, p. 181–199. Springer (2021). https://doi.org/10.1007/978-3-030-72013-1_10
35. Wang, S., Zhan, N., Zou, L.: An improved HHL prover: An interactive theorem prover for hybrid systems. In: ICFEM. LNCS, vol. 9407. Springer (2015). https://doi.org/10.1007/978-3-319-25423-4_25
36. Xu, Q., He, W.: Hierarchical design of a chemical concentration control system. In: *Hybrid Systems*. LNCS, vol. 1066. Springer (1995). <https://doi.org/10.1007/BFb0020952>