

Guaranteed Error Bounds on Approximate Model Abstractions through Reachability Analysis

Luca Cardelli¹, Mirco Tribastone², Max Tschaikowski², and Andrea Vandin³

¹ Microsoft Research & University of Oxford, UK

² IMT School for Advanced Studies Lucca, Italy

³ DTU Compute, Denmark

Abstract. It is well known that exact notions of model abstraction and reduction for dynamical systems may not be robust enough in practice because they are highly sensitive to the specific choice of parameters. In this paper we consider this problem for nonlinear ordinary differential equations (ODEs) with polynomial derivatives. We introduce approximate differential equivalence as a more permissive variant of a recently developed exact counterpart, allowing ODE variables to be related even when they are governed by nearby derivatives. We develop algorithms to (i) compute the largest approximate differential equivalence; (ii) construct an approximate quotient model from the original one via an appropriate parameter perturbation; and (iii) provide a formal certificate on the quality of the approximation as an error bound, computed as an over-approximation of the reachable set of the perturbed model. Finally, we apply approximate differential equivalences to study the effect of parametric tolerances in models of symmetric electric circuits.

1 Introduction

Ordinary differential equations (ODEs) are a prominent model of dynamical systems across many branches of science and engineering, and have enjoyed increasing popularity in computer science, for instance, in computational systems biology [13,30,4], as an approximation to large-scale Markov models and as the laws of continuous motion in hybrid systems [7]. This has motivated techniques for the *comparison* and *minimization* of ODEs based on behavioral relations, along the lines of other foundational quantitative models of computation, e.g. [25]. Here we consider *differential equivalence* [11], recently developed as an equivalence over ODE variables yielding a quotient that preserves the dynamics of the original one. However, differential equivalences (reviewed in Section 2) are exact, hence highly sensitive to parameter values and initial conditions. This may hinder their practical usability in some applications domains, for instance due to parameter uncertainty arising from finite-precision measurements in biology or the tolerances of electric components in electrical engineering.

Our objective is to develop approximate variants of differential equivalence (in Section 3). We study models with derivatives given by multivariate polynomials (over the ODE variables) of any degree, thus restricting the scope of [11]. However we remark that this is still quite a generous class since it includes chemical reaction networks (CRNs)

with mass-action kinetics and linear/affine systems, thus covering, e.g., continuous-time Markov chains through their Kolmogorov equations.

Considering polynomial derivatives allows us to introduce a notion of equivalence that just concerns the ODE “syntax” (while the nonlinear class of ODEs of [11] required symbolic SMT-based checks). Our main idea is to consider a threshold parameter $\varepsilon \geq 0$, which intuitively captures perturbations in polynomial coefficients. This allows relating ODE variables that would be distinct otherwise. Like in other established approaches such as behavioral pseudometrics (e.g., [1,9]), $\varepsilon = 0$ corresponds to an exact differential equivalence of [11]. In addition to defining criteria for approximate differential equivalences, we provide an algorithm for obtaining the largest one. This is done via partition refinement, computing the coarsest refinement of a given initial partition of ODE variables for a given “structural” tolerance ε .

A quotient ODE system can be constructed from a *reference model*, obtained through a perturbation of the coefficients of the original model which makes the given approximate differential equivalence an exact one. By considering a metric (the Euclidean norm) to measure the degree of perturbation, the reference model is the one which minimizes such perturbation. This can be done efficiently by solving an optimization problem which runs polynomially with the size of the ODE system [22]. This approach is analogous to optimal approximate lumping for Markov chains (e.g., [16]), although our theory can be applied to other choices of reference models.

The bound of the error produced by the reference model with respect to the original system can be computed by studying the reachable set of the reference model from an uncertain set of initial conditions that covers the applied perturbation. Since the reference model subsumes any behavior of the quotient, the bound formally relates the quotient model to the original model. Section 4 presents a bound which relies on a linearization of the reference model (which can be efficiently computed in the case of polynomial ODE systems). First, we bound the reachable set of the linearized model using closed form solutions, similarly to [24]. Then, we provide a conservative condition (i.e., an over-approximation) that ensures that the linearized model describes the original nonlinear behavior dynamics well. Our bound is given in terms of an ε - δ argument (similar in spirit to the ones routinely used in calculus). Informally, it states the following: for any choice of the structural tolerance ε , there exist a degree of perturbation δ and an *amplifier* λ such that, for any ODE system obtained by applying a perturbation to the reference model of at most δ , at all time points the difference between the solution of the reference model and the perturbed one is at most λ times the perturbation.

Being based on a linearization, it is perhaps not surprising that the as-computed δ will account only for *small* perturbations of the parameters. Yet numerical experiments in Section 5 show that these can be enough to explain quasi-symmetric behavior due to parametric tolerances in components of real electric network designs [31]. By comparing to over-approximation techniques supported by state-of-the-art tools C2E2 [17], CORA [3,2] and Flow* [12], we show that our bounding technique can complement them in that it can scale to larger systems while being more conservative in the size of the initial uncertain set that it supports.

Further related work. Differential equivalence is promising when the ODE system is composed of several identical subsystems that depend on some common context [33].

It is related but not comparable to bisimulation for differential systems [18,20] since it partitions ODE variables rather than the state space. Likewise, it complements [8] that captures nonlinear relations between ODE variables but does not enjoy a polynomial time algorithm like [10].

A classic approximation approach relies on Lyapunov-like functions [27,15] known from stability theory of ODEs. However, for nonlinear systems the automatic computation of Lyapunov-like functions remains a challenging task. Restricting to special classes of Lyapunov-like functions (e.g., sum-of-squares polynomials [18]) leads to efficient construction algorithms which may provide tight bounds, but existence is not guaranteed. On the other hand, approximations with differential inequalities [35] can be computed efficiently, but may be loose. Abstraction, supported by CORA and Flow*, locally approximates the nonlinear model by a multivariate polynomial or an affine system, see [5,12] and references therein. Similarly in spirit we linearize across a reference trajectory. A closer approach to ours is discussed in [17] and supported by the corresponding tool C2E2. It combines local Lyapunov-like functions and techniques based on sensitivity analysis [24]. Our bound is however different because the nonlinear part is bounded analytically by restricting to polynomial derivatives.

More in general, research on approximate quotients of ODE systems spans many disciplines. In chemistry, it can be traced back to Kuo and Wei [23]. They studied monomolecular reaction networks, which give rise to affine ODE systems. The approximation consists in *nearly exact lumping*, i.e., a linear transformation of the state space that would be exact up to a perturbation of the parameters (hence we are similar in spirit). The approximation, however, only applies when the transition matrix underlying the linear system is diagonalizable. Li and Rabitz extend approximate lumping to general CRNs [26], but an explicit error bound is not given. In a similar vein, approximate quotients in ecology have been studied from the point of view of finding a reduced ODE system whose derivatives are as close as possible (in norm) to the derivatives of the original ODE system, where the 0-distance induce the exact quotient [21]. The justification that variables underlying similar ODEs have nearby solutions is grounded on Gronwall’s inequality which is also at the basis of more recent quotient constructions [19], which however are not algorithmic, unlike in this paper.

Notation and basic definitions. We denote the infinity norm $\|\cdot\|_\infty$ by $\|\cdot\|$, while $\|\cdot\|_2$ is the Euclidian norm. Whenever convenient, for a given partition of variables \mathcal{H} , we write $H = \{x_{H,1}, \dots, x_{H,|H|}\}$ for any $H \in \mathcal{H}$. We denote by $\psi[t/s]$ the term that arises by replacing each occurrence of t in ψ by s . Let S be a set, and $\mathcal{H}_1, \mathcal{H}_2$ two partitions of S . Then, \mathcal{H}_1 is a *refinement* of \mathcal{H}_2 if for any block $H_1 \in \mathcal{H}_1$ there exists a block $H_2 \in \mathcal{H}_2$ such that $H_1 \subseteq H_2$. For any partition \mathcal{H} of S , let $\sim_{\mathcal{H}}$ denote the unique equivalence relation with $\mathcal{H} = S/\sim_{\mathcal{H}}$. The transitive closure of a relation \sim is denoted by \sim^* .

2 Background

Throughout the paper we consider a polynomial initial value problem (PIVP) over the set of ODE variables $S = \{x_1, \dots, x_n\}$. It is defined by the ODEs $\dot{x}_i = q_i$, $1 \leq i \leq n$, where q_i is a multivariate polynomial over S . The initial condition of the PIVP is given

by $\sigma : \mathcal{S} \rightarrow \mathbb{R}$; $x_i(t)$ denotes the unique solution for variable x_i at time point t starting from $x_i(0) = \sigma(x_i)$. We consider PIVPs that do not exhibit finite explosion times, i.e., whose solutions do not have a singularity at any finite point in time. This property is shared by the vast majority of practical models and can be efficiently checked via numerical ODE solvers.

A polynomial q_i is given in the normal form if each monomial $x^\alpha \equiv \prod_{x_i \in \mathcal{S}} x_i^{\alpha_{x_i}}$, where $\alpha \in \mathbb{N}_0^{\mathcal{S}}$ is a multi-index, appears in q_i at most once. The normal form of a polynomial q_i is denoted by $\mathcal{N}(q_i)$. Without loss of generality, we assume that the polynomials q_i of a PIVP $\dot{x} = q(x)$, where $q = (q_i)_{x_i \in \mathcal{S}}$, are given in normal form. For a polynomial q_i in normal form with variables in \mathcal{S} , let $c(q_i, x^\alpha)$ denote the coefficient of the monomial x^α , where $\alpha \in \mathbb{N}_0^{\mathcal{S}}$.

Example 1. We use the following ODE system, with variables $\mathcal{S} = \{x_1, x_2, x_3\}$, as a running example.

$$\dot{x}_1 = -4.00x_1 + x_2 + x_3 \quad \dot{x}_2 = 1.99x_1 - x_2 \quad \dot{x}_3 = 2.01x_1 - x_3 \quad (1)$$

In [11] two variants of differential equivalence were introduced for IDOL, a class of nonlinear ODE systems covering derivatives more general than polynomials. Here we find it convenient to restate them for a PIVP. (The proofs for this correspondence are straightforward hence we omit them.)

We begin with *backward differential equivalence* (BDE), which relates variables that have the same solutions at all time points. The definition of BDE for PIVP makes pairwise comparisons between the coefficients of any two variables in the same equivalence class.

Definition 1 (BDE). Fix a PIVP, a partition \mathcal{H} of \mathcal{S} and write $x_i \sim_{\mathcal{H}}^B x_j$ if all coefficients of the following polynomial are zero,

$$\wp_{i,j}^{\mathcal{H}} := (q_i - q_j)[x_{H',1}/x_{H'}, \dots, x_{H',|H'|}/x_{H'} : H' \in \mathcal{H}]$$

i.e., when

$$\sum_{\alpha \in \mathbb{N}_0^{\mathcal{S}}} |c(\wp_{i,j}^{\mathcal{H}}, x^\alpha)| = 0. \quad (2)$$

A partition \mathcal{H} is a BDE if $\mathcal{H} = \mathcal{S} / (\sim_{\mathcal{H}}^B \cap \sim_{\mathcal{H}})$.

Essentially, establishing that a candidate partition is a BDE consists in comparing the coefficients of the monomials of the ODEs of related variables, up to the natural equivalence class induced on monomials by the equivalence relation through $\wp_{i,j}^{\mathcal{H}}$ — for instance, the partition $\{\{x_1\}, \{x_2, x_3\}\}$ will equate the monomials x_1x_2 and x_1x_3 . Then, for any two variables in the same block it must hold that the differences between the coefficients of the same monomials (modulo the induced equivalence class) are zero.

Example 2. In our running example let us consider the partition of variables $\mathcal{H} = \{H_1, H_2\}$, with $H_1 = \{x_1\}$ and $H_2 = \{x_2, x_3\}$. Then \mathcal{H} is not a BDE because

$$\wp_{2,3}^{\mathcal{H}} = -0.02x_1 \quad \text{and} \quad c(\wp_{2,3}^{\mathcal{H}}, x_1) = -0.02 \neq 0.$$

Forward differential equivalence (FDE) identifies a partition that induces a quotient ODE that tracks sums of variables in each equivalence class. For instance, for any initial condition we have that $\{\{x_1\}, \{x_2, x_3\}\}$ is an FDE for (1) because we can find an ODE system for $x_2 + x_3$:

$$\dot{x}_1 = -4.00x_1 + (x_2 + x_3) \quad (x_2 + x_3) = 4.00x_1 - (x_2 + x_3)$$

The change of variable $x_{23} = x_2 + x_3$ gives us the quotient ODE $\dot{x}_1 = -4.00x_1 + x_{23}$, $\dot{x}_{23} = 4.00x_1 - x_{23}$. From this we conclude that the solution satisfies $x_{23}(t) = x_2(t) + x_3(t)$ for all time t if this holds for the initial condition, i.e., $x_{23}(0) = x_2(0) + x_3(0)$.

For a PIVP, FDE can be checked by requiring that the evaluation of the polynomial that represents the quotient derivative for an equivalence class is invariant with respect to a redistribution of the values of any two variables within that equivalence class.

Definition 2 (FDE). Fix a PIVP, a partition \mathcal{H} of \mathcal{S} and write $x_i \sim_{\mathcal{H}}^F x_j$ if all coefficients of the polynomial $\sum_{H \in \mathcal{H}} \wp_{i,j}^H$ are zero, where

$$\wp_{i,j}^H := \sum_{x_k \in H} q_k - \sum_{x_k \in H} q_k [x_i/s(x_i + x_j), x_j/(1-s)(x_i + x_j)]$$

That is, when

$$\sum_{k=1}^m \sum_{\alpha \in \mathbb{N}_0^{S \cup \{s\}}} |c(\wp_{i,j}^{H_k}, x^\alpha)| = 0. \quad (3)$$

$\mathcal{H} = \{H_1, \dots, H_m\}$ is an FDE when $\mathcal{H} = \mathcal{S} / (\sim_{\mathcal{H}}^{F*} \cap \sim_{\mathcal{H}})$.

3 Approximate Differential Equivalences

Definitions. Approximate differential equivalence relaxes the equality conditions (2)-(3) of Definition 1 and 2 to inequalities with respect to a tolerance level ε .

Definition 3 (Approximate BDE). Fix a PIVP, a partition $\mathcal{H} = \{H_1, \dots, H_m\}$ of \mathcal{S} , and $\varepsilon \geq 0$. We write $x_i \sim_{\mathcal{H}, \varepsilon}^B x_j$ if $\sum_{\alpha \in \mathbb{N}_0^S} |c(\wp_{i,j}^{\mathcal{H}}, x^\alpha)| \leq \varepsilon$, where $\wp_{i,j}^{\mathcal{H}}$ is as in Definition 1. A partition \mathcal{H} is an ε -BDE if $\mathcal{H} = \mathcal{S} / (\sim_{\mathcal{H}, \varepsilon}^{B*} \cap \sim_{\mathcal{H}})$.

Definition 4 (Approximate FDE). Fix a PIVP, a partition $\mathcal{H} = \{H_1, \dots, H_m\}$ of \mathcal{S} , and $\varepsilon \geq 0$. We write $x_i \sim_{\mathcal{H}, \varepsilon}^F x_j$ if $\sum_{k=1}^m \sum_{\alpha \in \mathbb{N}_0^{S \cup \{s\}}} |c(\wp_{i,j}^{H_k}, x^\alpha)| \leq \varepsilon$, where $\wp_{i,j}^H$ is as in Definition 2. A partition \mathcal{H} is an ε -FDE when $\mathcal{H} = \mathcal{S} / (\sim_{\mathcal{H}, \varepsilon}^{F*} \cap \sim_{\mathcal{H}})$.

Setting $\varepsilon = 0$ recovers the exact counterparts in both cases. That is, \mathcal{H} is an BDE (resp., FDE) partition if and only if \mathcal{H} is a 0-BDE (resp., 0-FDE) partition. The two approximate differential equivalences are not comparable since their exact counterparts are not [11]. Since these two notions have similar structure in the rest of this paper we will illustrate only approximate BDE using small examples. Instead, both notions will be discussed in more detail for the numerical evaluation of Section 5.

Algorithm 1 Template partition refinement algorithm for the computation of the coarsest ε -FDE/ ε -BDE partition that refines a given initial partition \mathcal{G} .

Require: A PIVP over variables \mathcal{S} , a partition \mathcal{G} of \mathcal{S} , a threshold $\varepsilon \geq 0$, and a mode $\chi \in \{F, B\}$.

```

 $\mathcal{H} \leftarrow \mathcal{G}$ 
while true do
   $\mathcal{H}' \leftarrow S / (\sim_{\mathcal{H}, \varepsilon}^{\chi} \cap \sim_{\mathcal{H}})$ 
  if  $\mathcal{H}' = \mathcal{H}$  then
    return  $\mathcal{H}$ 
  else
     $\mathcal{H} \leftarrow \mathcal{H}'$ 
  end if
end while

```

Example 3. Let us consider our running example (1). Then, the partition $\{\{x_1\}, \{x_2, x_3\}\}$ is a 0.02-BDE partition, as can be easily seen from Example 2.

The next two results show the existence of largest approximate differential equivalences and of a partition-refinement algorithm to compute it.

Theorem 1. Fix a PIVP, a partition \mathcal{G} of \mathcal{S} , and $\varepsilon \geq 0$. Then, there exists a unique coarsest ε -FDE (ε -BDE) partition refining \mathcal{G} .

Theorem 2. Fix a PIVP, a partition \mathcal{G} of \mathcal{S} , and $\varepsilon \geq 0$. Then, Algorithm 1 computes the coarsest ε -FDE (ε -BDE) that refines \mathcal{G} if $\chi = F$ ($\chi = B$).

We now study how efficiently the conditions for approximate differential equivalence can be computed. The notions are defined with respect to the coefficients of the polynomials $\wp_{i,j}^H$ and $\wp_{i,j}^{\mathcal{H}}$, and thus require the computation of their normalization. In the case of ε -FDE, this yields exponential complexity due to term replacement. To see this, consider for instance the PIVP $\dot{x}_1 = x_2^k, \dot{x}_2 = x_1^k$, for some $k > 0$. Then, for $\mathcal{H} = \{\{x_1, x_2\}\}$, the term $q_1[x_1/s(x_1 + x_2), x_2/(1-s)(x_1 + x_2)]$ will be of size $\mathcal{O}(2^k)$. This stands in stark contrast to ε -BDE, where the conditions involve a difference between polynomials terms with no term rewritings. This discussion can be formalized as follows.

Theorem 3. There exists a polynomial Π such that, under the assumptions of Theorem 2, the number of steps done by Algorithm 1 is $\mathcal{O}(\Pi(2^d \cdot p))$ if $\chi = F$ and $\mathcal{O}(\Pi(p))$ if $\chi = B$, respectively, where d is the maximum degree of the polynomial and p is the number of monomials present in the PIVP.

In practice, d is usually not large. For example, mass-action CRNs feature ODEs with degree-two polynomials because in nature at most two species interact in every reaction. An experimental comparison between the reduction runtimes of ε -FDE and ε -BDE will be presented in Section 5. We also remark that since 0-FDE/BDE coincides with FDE/BDE, the above result provides a complexity bound for a subclass of ODE systems considered in [11].

Other computational considerations motivate the choice of the definitions of approximate differential equivalence given in this paper. Another natural definition could have

involved the computation of the maximal distance between derivatives “semantically”, i.e., under all possible evaluations within a given domain of interest. For example, consider the PIVP $\dot{x}_1 = x_1^3 - x_2$, $\dot{x}_2 = x_1 - x_2^3$. Establishing that $\{\{x_1, x_2\}\}$ is an ε -BDE would require checking that the difference between the derivatives satisfies

$$|\dot{x}_1 - \dot{x}_2| = |x_1^3 - x_1 + x_2^3 - x_2| \leq \varepsilon, \text{ for all } 0 \leq x_1, x_2 \leq C \quad (4)$$

for some finite $C > 0$ that represents some bounded domain where the trajectories are assumed to live. Since this example shows that this question is in general equivalent to solving a nonconvex optimization problem, we infer that the problem is NP-hard [29].

However it can be easily shown that our approximate differential equivalence, defined through the coefficients of the polynomials, corresponds to checks such as (4) in the following sense: If a partition \mathcal{H} satisfies constraints similar to (4) with respect to some $\varepsilon > 0$, then there exists an $\varepsilon' > 0$ such that \mathcal{H} is an ε' -FDE/BDE, and vice versa. The basic idea is to observe that a polynomial is the zero function if and only if its coefficients are all zero.

Finally, we remark that our structural/syntactic criteria can be used for PIVPs only. It is the lack of analogous conditions in the case of more general functions like minima or roots which prevents our approximate differential equivalences to be extended in a straightforward way to the full class of nonlinear ODEs of [11].

Reference PIVP. Given a partition of variables that represents an approximate differential equivalence, we construct a *reference PIVP* by finding a “perturbation” of the original PIVP — i.e., a modification of the initial condition σ and the coefficients present in q_1, \dots, q_n — which ensures that that very partition becomes an exact differential equivalence. On this reference PIVP one can use the quotienting algorithms for FDE/BDE developed in [11] (and not restated here formally for brevity). Therefore, the as-obtained quotient represents an approximate reduction of the original PIVP.

We obtain the desired perturbation by treating the original initial conditions and polynomial coefficients uniformly as initial conditions on an *extended* PIVP where every coefficient is parameterized and turned into a new ODE variable.

Definition 5. *The parameterization of a polynomial q_i in normal form with variables \mathcal{S} is denoted by \hat{q}_i and arises from q_i by replacing, for each $\alpha \in \mathbb{N}_0^{\mathcal{S}}$, the constant $c(q_i, x^\alpha)$ with the parameter $\mathfrak{c}(\hat{q}_i, x^\alpha)$.*

Example 4. The polynomials $q_2 = 1.99x_1 - x_2$ and $q_3 = 2.01x_1 - x_3$ from Example 1 give rise to the parameterized polynomials $\hat{q}_2 = \mathfrak{c}(\hat{q}_2, x_1)x_1 + \mathfrak{c}(\hat{q}_2, x_2)x_2$ and $\hat{q}_3 = \mathfrak{c}(\hat{q}_3, x_1)x_1 + \mathfrak{c}(\hat{q}_3, x_3)x_3$, respectively.

Definition 6 (Extended PIVP). *For a PIVP \mathcal{P} with variables \mathcal{S} , set $\Theta = \{\mathfrak{c}(\hat{q}_i, x^\alpha) \mid 1 \leq i \leq n, \alpha \in \mathbb{N}_0^{\mathcal{S}}\}$. Its extended version $\hat{\mathcal{P}}$ has variables $\mathcal{S} \cup \Theta$ and is given by $\dot{x}_i = \hat{q}_i$ and $\dot{\mathfrak{c}}(\hat{q}_i, x^\alpha) = 0$, where $x_i \in \mathcal{S}$ and $\alpha \in \mathbb{N}_0^{\mathcal{S}}$. For a given $\hat{\sigma} \in \mathbb{R}^{\mathcal{S} \cup \Theta}$, let $\hat{\mathcal{P}}(\hat{\sigma})$ denote the PIVP which arises from $\hat{\mathcal{P}}$ by replacing each $v \in \mathcal{S} \cup \Theta$ by the corresponding real value $\sigma(v) \in \mathbb{R}$ in $\hat{\mathcal{P}}$. In particular, let $\hat{\sigma}_0 \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ be such that $\mathcal{P}(\sigma) = \hat{\mathcal{P}}(\hat{\sigma}_0)$.*

Example 5. If \mathcal{P} is the PIVP from Example 1, its extended version $\hat{\mathcal{P}}$ is

$$\begin{aligned} \dot{x}_1 &= \mathbf{c}(\hat{q}_1, x_1)x_1 + \mathbf{c}(\hat{q}_1, x_2)x_2 + \mathbf{c}(\hat{q}_1, x_3)x_3, & \dot{\mathbf{c}}(\hat{q}_1, x_i) &= 0, & i &= 1, 2, 3, \\ \dot{x}_2 &= \mathbf{c}(\hat{q}_2, x_1)x_1 + \mathbf{c}(\hat{q}_2, x_2)x_2, & \dot{\mathbf{c}}(\hat{q}_2, x_i) &= 0, & i &= 1, 2, 3, \\ \dot{x}_3 &= \mathbf{c}(\hat{q}_3, x_1)x_1 + \mathbf{c}(\hat{q}_3, x_2)x_2, & \dot{\mathbf{c}}(\hat{q}_3, x_i) &= 0, & i &= 1, 2, 3. \end{aligned}$$

The corresponding $\hat{\sigma}_0$ satisfies $\hat{\sigma}_0(x_i) = \sigma(x_i)$ for $1 \leq i \leq 3$ and

$$\begin{aligned} \hat{\sigma}_0(\mathbf{c}(\hat{q}_1, x_1)) &= -4.00, & \hat{\sigma}_0(\mathbf{c}(\hat{q}_1, x_2)) &= 1.00, & \hat{\sigma}_0(\mathbf{c}(\hat{q}_1, x_3)) &= 1.00, \\ \hat{\sigma}_0(\mathbf{c}(\hat{q}_2, x_1)) &= 1.99, & \hat{\sigma}_0(\mathbf{c}(\hat{q}_2, x_2)) &= -1.00, \\ \hat{\sigma}_0(\mathbf{c}(\hat{q}_3, x_1)) &= 2.01, & \hat{\sigma}_0(\mathbf{c}(\hat{q}_3, x_2)) &= -1.00. \end{aligned}$$

The following is needed for the definition of the reference PIVP.

Definition 7. Given constant free polynomial $\hat{\phi}$ (i.e., such that $\hat{\phi}(0) = 0$) and $\Xi \subseteq \mathcal{S} \cup \Theta \cup \{s\}$, let $\mathbf{t}(\hat{\phi}, x^\alpha, \Xi)$ denote the coefficient term of x^α in $\mathcal{N}(\hat{\phi}, \Xi)$, where $\alpha \in \mathbb{N}_0^\Xi$ and $\mathcal{N}(\hat{\phi}, \Xi)$ is the normal form of $\hat{\phi}$ where variables outside Ξ are treated as parameters.

Example 6. With \hat{q}_2 and \hat{q}_3 as in Example 4 and $\Xi = \{x_1, x_2, x_3\}$, the normal form $\mathcal{N}(\hat{q}_2 - \hat{q}_3, \Xi)$ is given by $(\mathbf{c}(\hat{q}_2, x_1) - \mathbf{c}(\hat{q}_3, x_1))x_1 + (\mathbf{c}(\hat{q}_2, x_2) - \mathbf{c}(\hat{q}_3, x_2))x_2$, while $\mathbf{t}(\hat{q}_2 - \hat{q}_3, x_1, \Xi) = \mathbf{c}(\hat{q}_2, x_1) - \mathbf{c}(\hat{q}_3, x_1)$.

Definition 8. Given a PIVP with variables \mathcal{S} and an ε -FDE partition \mathcal{H} of \mathcal{S} , the set of linear constraints of \mathcal{H} is given by

$$\{\mathbf{t}(\tilde{\phi}_{i,j}^H, x^\alpha, \mathcal{S} \cup \{s\}) = 0 \mid \alpha \in \mathbb{N}_0^{\mathcal{S} \cup \{s\}}, H \in \mathcal{H} \text{ and } x_i \sim_{\mathcal{H}} x_j\} \quad (5)$$

with $\tilde{\phi}_{i,j}^H = \sum_{x_k \in H} \hat{q}_k - \sum_{x_k \in H} \hat{q}_k [x_i/s(x_i + x_j), x_j/(1-s)(x_i + x_j)]$.

If \mathcal{H} is an ε -BDE partition of \mathcal{S} , the corresponding set of linear constraints is

$$\begin{aligned} \{\mathbf{t}(\tilde{\phi}_{i,j}^{\mathcal{H}}, x^\alpha, \mathcal{S}) = 0 \mid \alpha \in \mathbb{N}_0^{\mathcal{S}}, x_i \sim_{\mathcal{H}} x_j\} \\ \cup \{x_{i_j} - x_{i_{j+1}} = 0 \mid 1 \leq j \leq k-1 \text{ and } \{x_{i_1}, \dots, x_{i_k}\} \in \mathcal{S}/\sim_{\mathcal{H}}\}, \end{aligned} \quad (6)$$

where $\tilde{\phi}_{i,j}^{\mathcal{H}} = (\hat{q}_i - \hat{q}_j)[x_{H',1}/x_{H'}, \dots, x_{H',|H'|}/x_{H'} : H' \in \mathcal{H}]$.

Example 7. From Example 2, we know that $\mathcal{H} = \{\{x_1\}, \{x_2, x_3\}\}$ is a 0.02-BDE partition of the PIVP (1). The set of linear constraints underlying \mathcal{H} is given by $\mathbf{c}(\hat{q}_2, x_1) - \mathbf{c}(\hat{q}_3, x_1) = 0$ and $x_2 - x_3 = 0$.

Remark 1. In line with its exact counterpart, an ε -BDE is “useful” under the further constraint that related variables have the same initial conditions in the reference model, as a necessary condition for having equal solutions at all time points. This translates into adding the constraints in (6) that perturbed initial conditions of related variables are equal. This leads, for instance, to the constraint $x_2 - x_3 = 0$ in the running example. For ε -FDE, instead, only constraints on the parameters Θ are made.

Theorem 4. Given a PIVP \mathcal{P} with variables \mathcal{S} , an ε -FDE/BDE partition \mathcal{H} and a configuration $\hat{\sigma} \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ that satisfies (5)/(6), it holds that \mathcal{H} is an FDE/BDE of $\hat{\mathcal{P}}(\hat{\sigma})$.

The linear system of constraints from Theorem 4 can be shown to be underdetermined, hence there are infinitely many perturbations that lead to an exact differential equivalence. This observation is an instance of the well-known fact that, in general, an approximate quotient is not unique. Here, we fix one candidate perturbation by assuming that nearby initial conditions yield nearby trajectories. This fact is asymptotically true due to Gronwall's inequality, as mentioned in Section 1.

We are interested in finding a configuration $\hat{\sigma}$ which satisfies the constraints of Theorem 4 and minimizes the distance $\|\hat{\sigma} - \hat{\sigma}_0\|_2$. Mathematically, this corresponds to the optimization problem

$$\hat{\sigma}_* = \underset{\hat{\sigma}: \text{Eq. (5)/(6) holds}}{\operatorname{argmin}} \|\hat{\sigma} - \hat{\sigma}_0\|_2 \quad (7)$$

Since the solution space of a linear system is convex, the Euclidian norm yields a convex quadratic program that can be solved in polynomial time [22].

Example 8. Let us continue Example 7 and assume that $\sigma(x_2) = \sigma(x_3)$. In such a case, it can be easily seen that $\hat{\sigma}_*$ and $\hat{\sigma}_0$ satisfy $\hat{\sigma}_*(\mathbf{c}(\hat{q}_2, x_1)) = \hat{\sigma}_*(\mathbf{c}(\hat{q}_3, x_1)) = (\hat{\sigma}_0(\mathbf{c}(\hat{q}_2, x_1)) + \hat{\sigma}_0(\mathbf{c}(\hat{q}_3, x_1)))/2 = 2.00$ and coincide on all other parameters. In other words, the closest PIVP that enjoys an exact BDE relating x_2 and x_3 is given, as expected, by perturbing the coefficients 1.99 and 2.01 of (1) to their average value, yielding:

$$\dot{x}_1 = -4.00x_1 + x_2 + x_3 \quad \dot{x}_2 = 2.00x_1 - x_2 \quad \dot{x}_3 = 2.00x_1 - x_3$$

The above discussions are summarized in the following.

Theorem 5. Given a PIVP, $\varepsilon \geq 0$, and an ε -FDE/BDE partition \mathcal{H} , the solution of (7) exists and can be computed in polynomial time.

The solution of the optimization problem (7) stated in Theorem 5 is informally depicted in Fig. 1a.

The reference PIVP is the extended, exactly reducible PIVP with the optimum initial condition $\hat{\sigma}_*$, i.e., $\hat{\mathcal{P}}(\hat{\sigma}_*)$. Its ODE solution is called the *reference trajectory*.

4 Error Bounds

The objective of this section is to provide a tight bound on the difference between the solution of the original PIVP and the reference. More specifically, we will show how to compute two values $\delta > 0$ and $\lambda > 0$ such that for all initial conditions $\hat{\sigma}_1 \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ with $\|x^{\hat{\sigma}_1}(0) - x^{\hat{\sigma}_*}(0)\| \leq \delta$, it holds that $\max_{0 \leq t \leq \hat{\tau}} \|x^{\hat{\sigma}_1}(t) - x^{\hat{\sigma}_*}(t)\| \leq \lambda \|x^{\hat{\sigma}_1}(0) - x^{\hat{\sigma}_*}(0)\|$, where $x^{\hat{\sigma}}$ denotes the solution underlying $\hat{\mathcal{P}}(\hat{\sigma})$ and $\hat{\tau} > 0$ is a previously fixed finite time horizon.

The quantity δ gives the size of the ball around the initial condition $\hat{\sigma}_*$ of the reference PIVP, whereas λ is the *amplifier* that relates the maximum distance between trajectories




Fig. 1: Given a PIVP \mathcal{P} , a partition \mathcal{G} of \mathcal{S} , and an $\varepsilon > 0$, the coarsest ε -FDE/BDE partition \mathcal{H} that refines \mathcal{G} is constructed. Afterwards, the solution $\hat{\sigma}_*$ of the optimization problem (7) is computed in Fig. 1(a). This allows to compute the ε -FDE/BDE quotient $\hat{\mathcal{P}}(\hat{\sigma}_*)$ of \mathcal{H} . With this, λ and δ from Theorem 7 are calculated. In the case the distance between $\hat{\sigma}_0$ and $\hat{\sigma}_*$ does not exceed δ , the tight bounds of Theorem 7 can be applied and relate the trajectories of $\hat{\mathcal{P}}(\hat{\sigma}_*)$ and $\hat{\mathcal{P}}(\hat{\sigma}_0) = \mathcal{P}(\sigma)$, as depicted in Fig. 1(b).

to the distance between the initial conditions. Therefore, if the initial condition of the original PIVP $\hat{\mathcal{P}}(\hat{\sigma}_0)$ falls within the prescribed δ ball, then the above statement will provide a formal bound of the error made in approximating the original PIVP $\hat{\mathcal{P}}(\hat{\sigma}_0)$ with the reference PIVP. This idea is visualized in Fig. 1(b).

Let us recall the notion of Jacobian matrix.

Definition 9. Given an extended PIVP with variables $\mathcal{S} \cup \Theta$, the entries of the Jacobian matrix $A = (A_{i,j})_{x_i, x_j \in \mathcal{S} \cup \Theta}$ are given by $A_{i,j} = \partial_{x_j} \hat{q}_i$, where ∂_x denotes the partial derivative with respect to x .

Let $A(t) \in \mathbb{R}^{\mathcal{S} \cup \Theta \times \mathcal{S} \cup \Theta}$ denote the Jacobian obtained by plugging in the reference trajectory $x^{\hat{\sigma}_*}(t)$. We will need the following result from the theory of ODEs.

Theorem 6. There exists a family of matrices $\Lambda(t_0, t_1)$, with $0 \leq t_0 \leq t_1 \leq \hat{\tau}$, such that the solution of $\dot{y}(t) = A(t)y(t)$, where $y(t_0) = y_0$ and $t_0 \leq t \leq \hat{\tau}$, is given by $y(t) = \Lambda(t_0, t)y_0$ for all $t_0 \leq t \leq \hat{\tau}$.

This is needed in the following.

Theorem 7. Consider an extended PIVP $\hat{\mathcal{P}}$ with variables $\mathcal{S} \cup \Theta$ and define $\lambda_0 = \max_{0 \leq t \leq \hat{\tau}} \|A(0, t)\|$ and $\lambda_1 = \max_{0 \leq t_0 \leq t_1 \leq \hat{\tau}} \|A(t_0, t_1)\|$. Further, define the remainder function $r : [0; \hat{\tau}] \times \mathbb{R}^{\mathcal{S} \cup \Theta} \rightarrow \mathbb{R}^{\mathcal{S} \cup \Theta}$ via

$$r(t, x - x^{\hat{\sigma}_*}(t)) = \hat{q}(x) - \hat{q}(x^{\hat{\sigma}_*}(t)) - A(t)(x - x^{\hat{\sigma}_*}(t))$$

and let $0 \leq d_2, d_3, \dots$ be such that $\|r(t, y)\| \leq \sum_{k=2}^{\deg(\hat{\mathcal{P}})} d_k \|y\|^k$ for all $y \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ and $0 \leq t \leq \hat{\tau}$. Then, with $\lambda = 2\lambda_0$, for any $x^{\hat{\sigma}_1}(0) \in \mathbb{R}^{\mathcal{S} \cup \Theta}$, it holds that

$$\|x^{\hat{\sigma}_1}(0) - x^{\hat{\sigma}_*}(0)\| \leq \delta \Rightarrow \max_{0 \leq t \leq \hat{\tau}} \|x^{\hat{\sigma}_1}(t) - x^{\hat{\sigma}_*}(t)\| \leq \lambda \|x^{\hat{\sigma}_1}(0) - x^{\hat{\sigma}_*}(0)\|$$

whenever $\delta > 0$ satisfies $\sum_{k=2}^{\deg(\hat{\mathcal{P}})} d_k (2\lambda_0 \delta)^{k-1} \leq (2\lambda_1 \hat{\tau})^{-1}$.

Theorem 7 provides a bound on the difference $x^{\hat{\sigma}^1}(t) - x^{\hat{\sigma}^*}(t)$ in terms of the initial perturbation $x^{\hat{\sigma}^1}(0) - x^{\hat{\sigma}^*}(0)$ if the latter is sufficiently small, i.e., does not exceed δ . We wish to point out that the maximal δ satisfying $\sum_{k=2}^{\deg(\hat{\mathcal{P}})} d_k (2\lambda_0 \delta)^{k-1} \leq (2\lambda_1 \hat{\tau})^{-1}$ is a root of a polynomial in one variable and thus can be efficiently approximated from below via Newton's method. Instead, the assumption $\|r(s, y)\| \leq \sum_{k=2}^{\deg(\hat{\mathcal{P}})} d_k \|y\|^k$ on the remainder function r states essentially that, for any $k \geq 2$, the sum of all k -th order derivatives of r are bounded by d_k along the reference trajectory $x^{\hat{\sigma}^*}$.

The next result shows that the bound of Theorem 7 is tight and relies on the sharp bound available in the special case of linear systems (i.e., $\deg(\hat{\mathcal{P}}) = 1$) as discussed in [14].

Theorem 8. *If an extended PIVP $\hat{\mathcal{P}}$ satisfies $\deg(\hat{\mathcal{P}}) = 1$ and $\lambda = 2\lambda_0$, it holds that*

$$\max_{0 \leq t \leq \hat{\tau}} \|x^{\hat{\sigma}^1}(t) - x^{\hat{\sigma}^*}(t)\| \leq \frac{\lambda}{2} \|x^{\hat{\sigma}^1}(0) - x^{\hat{\sigma}^*}(0)\|$$

for any $x^{\hat{\sigma}^1}(0) \in \mathbb{R}^{\mathcal{S} \cup \Theta}$. The bound is tight in the sense that there exist $0 \leq t \leq \hat{\tau}$ and $x^{\hat{\sigma}^1}(0) \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ such that $\|x^{\hat{\sigma}^1}(t) - x^{\hat{\sigma}^*}(t)\| = \frac{\lambda}{2} \|x^{\hat{\sigma}^1}(0) - x^{\hat{\sigma}^*}(0)\|$.

Note that the amplifier in Theorem 7 is twice as large as the amplifier in Theorem 8. This is because the proof of Theorem 7 has to estimate nonlinear terms present in the remainder function r . More importantly, Theorem 8 shows that the amplifier of Theorem 7 cannot be substantially improved.

Remark 2. We note that λ_0, λ_1 can be estimated efficiently. Indeed, let $e_{x_i} \in \mathbb{R}^{\mathcal{S} \cup \Theta}$ be the x_i -th unit vector in $\mathbb{R}^{\mathcal{S} \cup \Theta}$, i.e., $e_{x_i}(x_j) = \delta_{i,j}$ where $\delta_{i,j}$ is the Kronecker delta. Then, if $y(t_0) = e_{x_i}$, Theorem 6 implies $y(t_1) = A(t_0, t_1)e_{x_i}$. Since $A(0, t_1)e_{x_i}$ is the x_i -column of $A(0, t_1)$ and $A(t_0, t_1) = A(0, t_1)A(0, t_0)^{-1}$, this shows that the matrices $A(t_0, t_1)$ can be computed by solving $|\mathcal{S} \cup \Theta|$ instances of the linear ODE system from Theorem 6 up to time $\hat{\tau}$.

By calculating a bound $L > 0$ on $\max_{0 \leq t \leq \hat{\tau}} \|A(t)\|$ and by computing the matrices $A(t_i, t_j)$ for all time points t_k underlying a fixed discretization step $\Delta t > 0$ of $[0; \hat{\tau}]$, the following can be shown.

Lemma 1. *Together with $\lambda_0^+ = \max_i \|A(0, t_i)\|$ and $\lambda_1^+ = \max_{i \leq j} \|A(t_i, t_j)\|$, it holds that $\lambda_0 \leq \lambda_0^+ e^{L\Delta t}$ and $\lambda_1 \leq \lambda_1^+ [1 + L\Delta t(e^{L\Delta t} + 1)]$.*

The next result simplifies the constraints on δ from Theorem 7 if $\deg(\hat{\mathcal{P}}) \leq 3$.

Lemma 2. *In the case where $\deg(\hat{\mathcal{P}}) \leq 3$, the constraint on δ of Theorem 7 simplifies to $\delta \leq \left[2\hat{\tau}\lambda_0\lambda_1 \left(d_2 + \sqrt{d_2^2 + \frac{2d_3}{\lambda_1\hat{\tau}}} \right) \right]^{-1}$.*

The above lemma applies, for instance, to most biochemical systems, as discussed in Section 3. The next result, instead, allows for an efficient estimation of d_k , with $2 \leq k \leq \deg(\hat{\mathcal{P}})$.

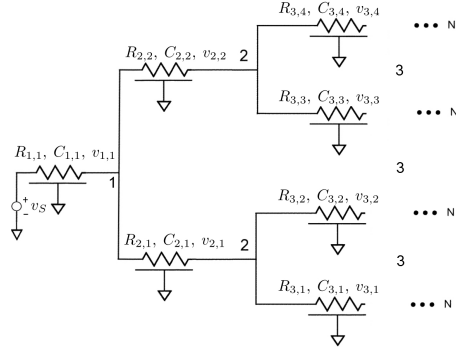


Fig. 2: H-tree network adapted from [31].

i	R_i^* (m Ω)	C_i^* (fF)
1	3.19	0.280
2	6.37	0.300
3	12.75	0.130
4	25.50	0.140
5	50.00	0.070
6	100.00	0.070
7	200.00	0.035
8	400.00	0.035

Table 1: Nominal parameters of electronic components at different depths i .

Lemma 3. Given an extended PIVP $\hat{\mathcal{P}}$ with variables $\mathcal{S} \cup \Theta$, let $\#_k(\hat{q}_i)$ be the number of degree k monomials in $\mathcal{N}(\hat{q}_i)$ and $D(\hat{q}_i, \hat{\sigma})$ the largest coefficient of $\mathcal{N}(\hat{q}_i)$ for configuration $\hat{\sigma} \in \mathbb{R}^{\mathcal{S} \cup \Theta}$. With $C = \max_{0 \leq t \leq \hat{\tau}} \|x^{\hat{\sigma}^*}(t)\|$, $M = \max_{x_i \in \mathcal{S}} \max_{k \geq 2} \#_k(\hat{q}_i)$ and $D = \max_{x_i \in \mathcal{S}} D(\hat{q}_i, \hat{\sigma}^*)$, it suffices to set d_k in Theorem 7 to $C^{\deg(\hat{\mathcal{P}}) - k} M D$.

For the case of linear systems whose parameters are subject to perturbation, instead, the following lemma can be applied. It provides a sharper estimate on d_2 but comes at the price of more involved computation.

Lemma 4. Given an extended PIVP $\hat{\mathcal{P}}$ with variables $\mathcal{S} \cup \Theta$, the Hessian matrix $H^k = (H_{i,j}^k)_{x_i, x_j}$ of \hat{q}_k is given by $H_{i,j}^k = \partial_{x_i} \partial_{x_j} \hat{q}_k$. With this, d_2 can be chosen as $d_2 = \frac{1}{2} \cdot \max_{x_i \in \mathcal{S} \cup \Theta} \max_{0 \leq t \leq \hat{\tau}} \|H^i(x^{\hat{\sigma}^*}(t))\|$.

Example 9. Since $\deg(\hat{\mathcal{P}}) = 2$ in Example 5, coefficients d_3, d_4, \dots are zero and we only need to bound d_2 . Moreover, the constraint in Theorem 7 simplifies to $\delta \leq (4\hat{\tau}\lambda_0\lambda_1 d_2)^{-1}$ thanks to Lemma 2. By applying Lemma 3, instead, we see that it suffices to choose $d_2 = 2.00$ because $M = 2.00$ and $D = 1.00$. In the case of $\hat{\tau} = 3.00$, we thus get $\lambda_0 = \lambda_1 = 1.40$ which yields $\delta \leq 0.02$.

5 Evaluation

We consider a simplified (inductance free) version of a power distribution electrical network from [31], arranged as a tree called *H-tree* (Figure 2). We let N be the depth of the tree and denote the resistances and the capacitances at depth i by $R_{i,k}$ and $C_{i,k}$, respectively. The source voltage is v_s , here assumed to be constant, $v_s = 2.0\text{V}$. Then, the voltage across $C_{i,k}$, denoted by $v_{i,k}$, obeys the affine ODE

$$\dot{v}_{1,1} = \frac{v_s - v_{1,1}}{R_{1,1}C_{1,1}} - \frac{v_{1,1} - v_{2,1}}{R_{2,1}C_{1,1}} - \frac{v_{1,1} - v_{2,2}}{R_{2,2}C_{1,1}}, \quad \dot{v}_{i,k} = \frac{v_{i-1,l} - v_{i,k}}{R_{i,k}C_{i,k}}, \quad (8)$$

where $1 \leq i \leq N$, $k = 1, \dots, 2^{i-1}$, and $l = \lceil k/2 \rceil$, where $\lceil \cdot \rceil$ denotes the ceil function. Here we considered networks with depth up to $N = 8$. For depths $i \leq 4$, the nominal

N	Reference model					Time (s) & Maximal Over-approximation					
	Time (s)	λ	δ	$\ \cdot\ $	$\lambda \cdot \ \cdot\ $	C2E2	CORA		FLOW*		
2	1.96E+0	5.41	7.95E-4	4.43E-4	2.40E-3	1.00E+1 safe	4.82E+1	8.50E-3	2.51E+1	8.21E-3	
3	3.51E+0	6.27	6.33E-4	7.42E-5	4.65E-4	4.20E+2 safe	1.42E+2	7.80E-3	1.10E+2	9.74E-3	
4	7.75E+0	7.78	4.71E-4	2.17E-5	1.68E-4	—	6.23E+2	5.50E-3	9.40E+2	1.16E-2	
5	2.46E+1	7.78	4.71E-4	3.31E-4	2.58E-3	—	5.39E+3	5.00E-3	—	—	
6	8.42E+1	7.78	4.71E-4	8.97E-5	6.98E-4	—	—	—	—	—	
7	4.74E+2	7.78	4.71E-4	4.22E-4	3.28E-3	—	—	—	—	—	
8	—	—	—	—	—	—	—	—	—	—	

Table 2: H-tree model results.

parameter values R_i^* and C_i^* were taken from [31]; for $i \geq 5$, instead, we extrapolated them. The parameters are summarized in Table 1.

In [31] the authors show that when the values of resistors and capacitors of any depth are equal, i.e., $R_{i,\cdot} \equiv R_i^*$ and $C_{i,\cdot} \equiv C_i^*$ then the network is *symmetric*. That is, the voltages at the capacitors in any level are equal at all time points. Indeed, $\{\{v_{i,k} \mid 1 \leq k \leq 2^{i-1}\} \mid 1 \leq i \leq N\}$ is an exact BDE partition (with N equivalence classes).

We now study the robustness of the symmetry under the realistic assumption that resistances and capacitances are only approximately equal. In particular, we test whether it is possible to explain quasi-symmetries when the parameters have tolerance $\eta = 0.01\%$. This corresponds to a practical situation when components or measurements parameters enjoy high accuracy. We considered networks of different size by varying the maximum depth N from 2 to 8. For each size, in order to simulate a quasi-symmetric scenario we built 30 distinct ODE models by sampling values for $R_{i,k}$ and $C_{i,k}$ uniformly at random within η percent from their nominal values. These repetitions were made in order to avoid fixing a single instance that might unfairly favor our algorithm. To each model we applied the ε -BDE reduction algorithm; choosing $\varepsilon = 6.00E-4$, it returned a quotient corresponding to a perfectly symmetrical case. The reduction times were below 0.5 s in all cases. (Throughout this section, the runtimes reported were measured on a VirtualBox virtual machine running Ubuntu 64 bits over an actual 2.6 GHz Intel Core i5 machine with 4 GB of RAM.) Then, we computed the values of δ and λ over a time horizon of 7 times units. This was chosen as a representative time point, for any N , of the transient state of the network (to account for the fact that, typically, circuits are analyzed in the time domain for transient analysis only).

The presence of uncertain parameters required us to transform the originally affine system (8) into a polynomial system of degree two (by substituting each $1/(R_{i,k}C_{j,l})$ with a corresponding new state variable) with 2^{N+1} states. This nonlinearity ruled out the application of standard over-approximation techniques for linear systems.

We present the results for the random model with the smallest value of δ in Table 2. The runtimes (second column) refer to the computation cost of the λ - δ pair. In all cases, δ turned out to be larger than the distance between the original model and its quotient $\|\hat{\sigma}_0 - \hat{\sigma}_*\| = \|x^{\hat{\sigma}_0}(0) - x^{\hat{\sigma}_*}(0)\|$ (shown in column $\|\cdot\|$). This demonstrates that the 0.01% tolerance can be formally explained by approximate differential equivalence, as confirmed by the small values of the amplifiers λ .

We compared against C2E2, CORA, and Flow*. (We did not compare our technique with other approaches applicable to nonlinear dynamics such as VNODE-LP [32], Ariadne [6] or HySAT/iSAT [28] because those have been compared to Flow* in the past [12].) While for CORA and Flow* the comparison can be made directly, C2E2 seeks to decide whether a given set is reachable. In order to perform as fair a comparison as possible we chose unreachable sets generously far away from the over-approximations computed with our bound, in order to ensure that C2E2 computes an over-approximation proving this. In all three cases, the initial uncertain set was fixed so as to correspond to the ball around the initial condition of the reference model $\hat{\sigma}_*$ that included the original model $\hat{\sigma}_0$. This is the most favorable condition for the tools since it corresponds to the smallest uncertainty set which can provide a guaranteed error bound. Other settings of C2E2, CORA, and Flow* were chosen such to ensure successful termination. In the case of Flow* we used estimation remainder 0.10 and allowed for adaptive Taylor model degree between 2 and 6. In C2E2 we set the K value to 2000. We used the time step 0.10 for C2E2 since this ensured safety for the models that could be analyzed. For CORA and Flow* we set time step equal to 0.01 as this led to tight enough bounds. For our approach, instead, we used time step 0.023 because this ensured tight approximations of λ_0 and λ_1 via Lemma 1. In all cases the time-out was set to 3 hours.

The comparison results are also reported in Table 2. The over-approximations computed by C2E2 are not shown because they need not to be tight in order to verify the reachability problem. This is because C2E2 is refining an over-approximation only if this is necessary to decide safety. Thus, we report that C2E2 was able to verify a set to be safe, i.e., unreachable. For a network of depth N , the over-approximations for CORA and Flow* are reported as the maximal diameter of the flowpipe underlying $v_{N,1}$ across all time points. As such, it can be compared to the product $\lambda \cdot \|\hat{\sigma}_0 - \hat{\sigma}_*\|$ given by our bound, which is also explicitly reported in the table for the sake of easy comparability (column $\lambda \cdot \|\cdot\|$). Both CORA and Flow* reported tight bounds, of the same orders as ours. These correspond to at most ca. 1% error on the observed variable, which cannot exceed the value of 2.0 (the source voltage applied to the network). C2E2 did not terminate within the time-out with models with $N > 3$, while Flow* ran out of memory for $N > 4$; CORA, instead, failed to compute the symbolic Jacobian matrices for $N > 5$. Our approach, instead, timed out for $N = 8$. However, we wish to point out that our algorithm naturally applies to parallelization. Indeed, its bottleneck is in the computation of the set of linear ODE systems discussed in Remark 2, which can be trivially solved independently from each other.

Discussion. In summary, the experimental results suggest that our bounding technique may complement the current state of the art in reachability analysis. Indeed, it has been shown to handle ODE systems of larger size, but it provides a δ neighborhood that can explain perturbations up to ca 0.1% at best. This makes our approach beneficial for the automatic detection and abstraction from quasi-symmetries due to small uncertainties, e.g., measurement errors. On the other hand, algorithms such as those implemented in C2E2, CORA, and Flow* can theoretically cover arbitrarily larger initial uncertainties, but at a computational cost that blocked their applicability to our larger benchmarks. In future work it is worth investigating a possible combination of these techniques.

6 Conclusion

Reasoning about quantitative properties approximately can represent an effective way of taming the complexity of real systems. Here we have considered ordinary differential equations (ODEs) with polynomial derivatives. We developed notions of equivalence as a relaxation of their exact counterparts, allowing the derivatives of related ODE variables to vary up to a desired tolerance. Our algorithmic approach can be useful to systematically discover quasi-symmetries in situations such as those presented in our case study. In future work, it would be also possible to integrate other bounding techniques, such as [35] which lacks an automatic synthesis of a reference model but can offer a tradeoff between tightness of the bound and computation cost in its derivation.

Acknowledgement

Luca Cardelli is partially funded by a Royal Society Research Professorship. Mirco Tribastone is supported by a DFG Mercator Fellowship (SPP 1593, DAPS2 Project). Max Tschaikowski is supported by a Lise Meitner Fellowship funded by the Austrian Science Fund (FWF) under grant number M 2393-N32 (COCO).

References

1. Abate, A., Brim, L., Ceska, M., Kwiatkowska, M.Z.: Adaptive Aggregation of Markov Chains: Quantitative Analysis of Chemical Reaction Networks. In: CAV. pp. 195–213 (2015)
2. Althoff, M.: Reachability Analysis of Nonlinear Systems using Conservative Polynomialization and Non-Convex Sets. In: HSCC. pp. 173–182 (2013)
3. Althoff, M.: An introduction to CORA 2015. In: Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems (2015)
4. Arand, J., Spieler, D., Karius, T., Branco, M.R., Meilinger, D., Meissner, A., Jenuwein, T., Xu, G., Leonhardt, H., Wolf, V., et al.: In vivo control of CpG and non-CpG DNA methylation by DNA methyltransferases. *PLoS Genet* 8(6) (2012)
5. Asarin, E., Dang, T., Girard, A.: Reachability Analysis of Nonlinear Systems Using Conservative Approximation. In: HSCC (2003)
6. Benvenuti, L., Bresolin, D., Casagrande, A., Collins, P., Ferrari, A., Mazzi, E., Sangiovanni-Vincentelli, A., Villa, T.: Reachability computation for hybrid systems with Ariadne. {IFAC} Proceedings Volumes 41(2), 8960 – 8965 (2008), 17th {IFAC} World Congress
7. Bogomolov, S., Frehse, G., Grosu, R., Ladan, H., Podelski, A., Wehrle, M.: A Box-Based Distance between Regions for Guiding the Reachability Analysis of SpaceX. In: CAV. pp. 479–494 (2012)
8. Boreale, M.: Algebra, coalgebra, and minimization in polynomial differential equations. In: FOSSACS. pp. 71–87 (2017)
9. van Breugel, F., Worrell, J.: Towards quantitative verification of probabilistic transition systems. In: ICALP (2001)
10. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Maximal aggregation of polynomial dynamical systems. *Proceedings of the National Academy of Sciences* 114(38), 10029 – 10034 (2017)
11. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Symbolic computation of differential equivalences. In: POPL (2016)

12. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Flow*: An Analyzer for Non-linear Hybrid Systems. In: CAV. pp. 258–263 (2013)
13. Danos, V., Laneve, C.: Formal molecular biology. *Theoretical Computer Science* 325(1), 69–110 (2004)
14. Donzé, A., Maler, O.: Systematic Simulation Using Sensitivity Analysis. In: HSCC. pp. 174–189. Springer (2007)
15. Duggirala, P.S., Mitra, S., Viswanathan, M.: Verification of Annotated Models from Executions. In: EMSOFT. pp. 26:1–26:10. IEEE Press (2013)
16. E, W., Li, T., Vanden-Eijnden, E.: Optimal partition and effective dynamics of complex networks. *PNAS* 105(23), 7907–12 (Jun 2008)
17. Fan, C., Qi, B., Mitra, S., Viswanathan, M., Duggirala, P.S.: Automatic reachability analysis for nonlinear hybrid models with C2E2. In: CAV. pp. 531–538 (2016)
18. Girard, A., Pappas, G.: Approximate bisimulations for nonlinear dynamical systems. In: IEEE Conference on Decision and Control and European Control Conference (2005)
19. Iacobelli, G., Tribastone, M.: Lumpability of fluid models with heterogeneous agent types. In: DSN (2013)
20. Islam, M.A., Murthy, A., Bartocci, E., Cherry, E., Fenton, F.H., Glimm, J., Smolka, S.A., Grosu, R.: Model-order reduction of ion channel dynamics using approximate bisimulation. *Theor. Comput. Sci.* 599, 34–46 (2015)
21. Iwasa, Y., Levin, S.A., Andreasen, V.: Aggregation in Model Ecosystems II. Approximate Aggregation. *Mathematical Medicine and Biology* 6(1), 1–23 (1989)
22. Kozlov, M., Tarasov, S., Khachiyan, L.: The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics* 20(5), 223 – 228 (1980)
23. Kuo, J.C.W., Wei, J.: Lumping analysis in monomolecular reaction systems. Analysis of approximately lumpable system. *Industrial & Engineering Chemistry Fundamentals* 8(1), 124–133 (1969)
24. Lal, R., Prabhakar, P.: Bounded error flowpipe computation of parameterized linear systems. In: EMSOFT. pp. 237–246 (2015)
25. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. *Inf. Comput.* 94(1), 1–28 (1991)
26. Li, G., Rabitz, H.: A general analysis of approximate lumping in chemical kinetics. *Chemical Engineering Science* 45(4), 977–1002 (1990)
27. Majumdar, R., Zamani, M.: Approximately Bisimilar Symbolic Models for Digital Control Systems. In: CAV. pp. 362–377 (2012)
28. Martin Fränzle and Christian Herde and Tino Teige and Stefan Ratschan and Tobias Schubert: Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure. *Journal on Satisfiability, Boolean Modeling and Computation* 1, 209–236 (2007)
29. Pardalos, P.M., Vavasis, S.A.: Quadratic programming with one negative eigenvalue is NP-hard. *Journal of Global Optimization* 1(1), 15–22 (1991)
30. Pedersen, M., Plotkin, G.D.: A language for biochemical systems: Design and formal specification. In: IEEE/ACM TCBB XII, LNCS, vol. 5945, pp. 77–145. Springer (2010)
31. Rosenfeld, J., Friedman, E.G.: Design Methodology for Global Resonant H-Tree Clock Distribution Networks. *IEEE Trans. VLSI Syst.* 15(2), 135–148 (2007)
32. S., N.: Implementing a Rigorous ODE Solver Through Literate Programming (2011)
33. Tschaikowski, M., Tribastone, M.: Tackling continuous state-space explosion in a Markovian process algebra. *Theor. Comput. Sci.* 517, 1–33 (2014)
34. Tschaikowski, M., Tribastone, M.: Approximate reduction of heterogeneous nonlinear models with differential hulls. *IEEE TAC* (2016)