Computers and Chemical Engineering 106 (2017) 596–608

Contents lists available at ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Rapid and accurate reachability analysis for nonlinear dynamic systems by exploiting model redundancy

Kai Shen, Joseph K. Scott*

Department of Chemical and Biomolecular Engineering, Clemson University, Clemson, SC 29634, USA

ARTICLE INFO

Article history:
Received 10 December 2016
Received in revised form 27 July 2017
Accepted 3 August 2017
Available online 12 August 2017

Keywords:
Reachability
Interval analysis
Differential inequalities
Global optimization
Uncertainty propagation

ABSTRACT

A new method is presented for enclosing the reachable sets of nonlinear ordinary differential equations subject to a range of inputs. Reachable set enclosures are used for uncertainty propagation, robust control, and global optimization of dynamic systems arising in a variety of applications. However, existing methods often provide an unworkable compromise between cost and accuracy. For example, fast interval methods often produce divergent bounds, while methods based on more complex sets scale poorly with problem size. To overcome this, a novel method is introduced for reducing the conservatism of fast interval methods through the select addition of redundant model equations that can be exploited in the bounding procedure. Several case studies demonstrate that such redundancy can dramatically reduce conservatism. The additional cost is modest in most cases, but does become significant when many redundant equations are used.

1. Introduction

This article presents a new method for computing a rigorous enclosure of the set of solutions reachable by a given system of nonlinear ordinary differential equations (ODEs) subject to uncertain inputs (i.e., initial conditions and model parameters). Such sets are referred to as reachable sets, and methods for enclosing them are useful for quantifying the effects of uncertainty in dynamic models arising in various applications, including (bio)chemical reaction networks (Scott and Barton, 2010a,b; Moisan et al., 2009), autonomous vehicles (Wang et al., 2015; Althoff and Dolan, 2014), and power systems (Pico and Aliprantis, 2014; Althoff et al., 2012). Such methods are also useful for process control, where the reachable sets of interest describe the uncertainty in a system's state arising from disturbances, imprecisely known model parameters, and measurement errors. Enclosing these sets is a central step in set-based state estimation (Raissi et al., 2004; Moisan et al., 2009), which is used for robust model predictive control (Limon et al., 2005; Hariprasad and Bhartiya, 2014) and set-based fault detection (Scott et al., 2016; Tulsyan and Barton, 2016; Raimondo et al., 2016). Reachability calculations are also used within algorithms for determining the set of inputs that lead to a desired set of 'safe' states, as in the construction design spaces for pharmaceutical processes (Kishida and Braatz, 2014, 2015). Finally, reachable set enclosures are also used in algorithms for solving dynamic optimization problems to guaranteed global optimality, which have been used to solve parameter estimation and open-loop optimal control problems (Singer et al., 2006; Scott and Barton, 2015; Lin and Stadtherr, 2007; Houska and Chachuat, 2014). In this context, the reachable sets describe the range of solutions that can be achieved by decision variables in a given region of the search space, and enclosures are used to eliminate regions by proving infeasibility or suboptimality.

Unfortunately, existing methods do not provide enclosures with sufficient speed and accuracy for many important applications. For example, in set-based state estimation and robust control, the desired enclosures depend on process measurements. This requires methods that are both fast enough for real-time implementation and accurate enough to be useful for decision-making. Similarly, global dynamic optimization requires accurate enclosures to avoid excessive subdivision of the search space, and high speed because even accurate methods may still consider thousands of regions (Wechsung et al., 2014). This combination of speed and accuracy remains a challenge.

Methods for enclosing reachable sets of nonlinear ODEs fall into three broad categories: Taylor series methods, conservative linearization, and differential inequalities. Although we focus on continuous-time dynamics here, other approaches are available for discrete-time systems, particularly with rational right-hand sides, based on e.g. linear fractional transformations and skewed struc-

* Corresponding author.
E-mail addresses: shen5@clemson.edu (K. Shen), jks9@clemson.edu (J.K. Scott).

http://dx.doi.org/10.1016/j.compchemeng.2017.08.001
0098-1354/

tured singular values (Kishida et al., 2014; Kishida and Braatz, 2011).

Taylor series methods propagate enclosures over discrete time steps by constructing a Taylor expansion of the states with respect to time and bounding the coefficients with, e.g., interval arithmetic (Nedialkov et al., 1999). The resulting enclosure is then inflated by a rigorous bound on the truncation error. Classical methods propagate interval enclosures, which makes them relatively efficient but often conservative. Modern methods achieve high accuracy by using Taylor models, which are multivariate Taylor expansions in the model inputs with rigorous interval remainder bounds (Berz and Makino, 2006; Lin and Stadtherr, 2007). Further improvements have been achieved using ellipsoidal and other non-interval remainder bounds (Houska et al., 2013, 2015). However, high accuracy often requires high-order Taylor models, which can become intractable because the number of coefficients scales exponentially in the number of states and inputs.

Conservative linearization methods propagate enclosures over discrete time steps by first using a locally linearized model and then adding a rigorous bound on the linearization error (Althoff et al., 2008; Althoff and Krogh, 2014). The reachable sets of the linear system can be enclosed using efficient set representations such as ellipsoids or zonotopes, and modern methods of this type have been shown to produce highly accurate enclosures in many applications (Althoff et al., 2012; Althoff and Dolan, 2014). However, this often requires complex set representations and hence high cost (see e.g. the use of 400th order zonotopes, each described by 2406 real numbers, to enclose 6-dimensional reachable sets in (Scott and Barton, 2013a,b)).

Finally, differential inequalities (DI) methods construct and solve an auxiliary system of ODEs that describes componentwise upper and lower bounds on the reachable set as its solutions. Harrison (1977) originally observed that such a system can be constructed automatically using interval arithmetic. Moreover, this system can be solved with any state-of-the-art numerical integrator, whereas both Taylor series and conservative linearization methods require custom integration algorithms with significant step-size restrictions. Thus, DI methods are capable of producing bounds very rapidly (i.e., at a small multiple of the cost of integrating a single trajectory (Scott and Barton, 2013a,b)), making DI a potentially powerful tool for real-time control and global dynamic optimization. However, the resulting enclosures are often extremely conservative. Several methods have been proposed to address this by enabling the use of more complex set representations in place of intervals. Chachuat and Villanueva (2012) proposed an interesting use of DI to compute Taylor model enclosures. However, auxiliary ODEs are required for each Taylor coefficient, which is prohibitive for high-order expansions. Villanueva et al. (2015) introduced a general framework for using DI to compute general convex enclosures. Specific implementations have been developed by Scott and Barton (2010, 2013) and Harwood et al. (2016). In particular, Harwood et al. (2016) introduced an effective method for computing polytopic enclosures using DI. However, this method constructs auxiliary ODEs whose right-hand sides are evaluated by solving embedded linear programs rather than using interval arithmetic, which leads to significantly longer computation times.

This article presents a new method for reducing the conservatism of the DI approach while largely maintaining its efficiency. Rather than using complex non-interval enclosures, the central idea is to exploit model redundancy. This strategy is motivated by effective DI methods that have recently been developed for systems whose solutions satisfy natural bounds (e.g., nonnegativity) and linear relationships (e.g., conservation laws) that are implied by, and hence redundant with, the governing ODEs (Singer and Barton, 2006; Scott and Barton, 2010, 2013). In brief, it has been shown that such redundant relationships can be exploited within the DI bounding procedure to achieve much sharper bounds. Moreover, this is done using only fast interval operations, so the speed of the standard DI method is largely retained (Scott and Barton, 2013a,b). Unfortunately, these methods do not apply to the majority of systems of practical interest, which do not naturally obey any such redundant relations.

To address this limitation, this article presents a new technique for general nonlinear systems based on the deliberate introduction of carefully selected redundant equations, which are then exploited in a DI bounding procedure similar to that in (Scott and Barton, 2013a,b). This can be viewed as a dynamic analogue of methods used to generate redundant constraints in global optimization algorithms, such as the reformulation linearization technique (Sherali and Adams, 2009). Although an automated method for selecting redundant equations is not yet available, we demonstrate this technique through several detailed case studies, which clearly show that redundancy can dramatically reduce conservatism. The additional cost is modest in most cases, but does become significant when many redundant equations are used, highlighting the need for future work on selection heuristics. The mechanisms by which redundancy reduces conservatism are discussed in detail, and we provide preconditioning heuristics that significantly improve the efficacy of the added equations. Although we only consider DI, our results suggest that redundant equations could be used to reduce conservatism in other approaches as well, potentially enabling the use of lower complexity sets. Indeed, Villanueva et al. (2014) have shown that pre-existing affine solution invariants can stabilize the enclosures computed by Taylor methods.

2. Problem Statement

Let $I = [t_0, t_f]$, $P \subset \mathbb{R}^{n_p}$ compact, $D \subset \mathbb{R}^{n_x}$ open, and $\mathbf{f} : I \times P \times D \to \mathbb{R}^{n_x}$ and $\mathbf{x}_0 : P \to D$ locally Lipschitz continuous. Consider the nonlinear ODEs

$$\dot{\mathbf{x}}(t, \mathbf{p}) = \mathbf{f}(t, \mathbf{p}, \mathbf{x}(t, \mathbf{p})), \ \ \mathbf{x}(t_0, \mathbf{p}) = \mathbf{x}_0(\mathbf{p}), \tag{1}$$

where a solution is any continuously differentiable $\mathbf{x} : I \times P \to D$ satisfying (1) for all $(t, \mathbf{p}) \in I \times P$. It is assumed that a unique solution exists on I for every $\mathbf{p} \in P$.

Definition 1. The *reachable set* of (1) at $t \in I$ is

$$Re(t) \equiv \{\mathbf{x}(t, \mathbf{p}) : \mathbf{p} \in P\}. \tag{2}$$

Moreover, functions $\mathbf{x}^L, \mathbf{x}^U : I \to \mathbb{R}^{n_x}$ are *state bounds* for (1) if $\mathbf{x}^L(t) \le \mathbf{x}(t, \mathbf{p}) \le \mathbf{x}^U(t)$, $\forall (t, \mathbf{p}) \in I \times P$.

The best possible state bounds describe the interval hull of $Re(t)$, while all others are conservative. Our aim is to develop a method that exploits model redundancy to efficiently compute state bounds with minimal conservatism.

3. Background

3.1. Interval Arithmetic

Let $X = [\mathbf{x}^L, \ \mathbf{x}^U]$ denote the n-dimensional interval $\left\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^L \le \mathbf{x} \le \mathbf{x}^U\right\}$, and for any $D \subset \mathbb{R}^n$, let $\mathbb{I}D$ denote the set of interval subsets of D. Given $\mathbf{f} : D \to \mathbb{R}^m$, an interval function $F : \mathbb{I}D \to \mathbb{IR}^m$ is an *inclusion function* for $\mathbf{f}$ on D if $\mathbf{f}(X) \equiv \{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in X\} \subset F(X)$, $\forall X \in \mathbb{I}D$.

A function $\mathbf{f}$ is *factorable* if it is a finite recursive composition of basic operations, including $\{+, -, \times, \div\}$ and intrinsic univariate functions such as e^x, x^n, etc. In this case, a specific inclusion function called the *natural interval extension* of $\mathbf{f}$ can be computed by *interval arithmetic* (IA), which replaces each basic operation with an interval-valued counterpart (Moore, 1966). This is very efficient, but can also be very conservative due to the *dependency problem*;