Automatica 48 (2012) 1372–1379

Contents lists available at SciVerse ScienceDirect

Automatica

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

Brief paper

Fault coverage modeling in nonlinear dynamical systems☆

Matthias A. Müller [a], Alejandro D. Domínguez-García [b,1]

[a] Institute for Systems Theory and Automatic Control, University of Stuttgart, 70550 Stuttgart, Germany

[b] Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

ARTICLE INFO

Article history:
Received 30 December 2010
Received in revised form
15 August 2011
Accepted 25 November 2011
Available online 21 May 2012

Keywords:
Fault coverage
Fault tolerance
Nonlinear systems
Invariant sets
Input-to-state stability (ISS)

ABSTRACT

In this paper, we propose an approach for modeling fault coverage in nonlinear dynamical systems. Fault coverage gives a measure of the likelihood that a system will be able to recover after a fault occurrence. In our setup, the system dynamics are described by a standard state-space model. The system input (disturbance) is considered to be unknown but bounded at all times. Before any fault occurrence, the vector field governing the system dynamics is such that, for any possible input signal, the corresponding system reach set is contained in some region of the state space defined by the system performance requirements. When a fault occurs, the vector field that governs the system dynamics might be altered. Fault coverage is defined as the probability that, given a fault has occurred, the system trajectories remain, at all times, within the region of the state-space defined by the performance requirements. Input-to-state stability (ISS) concepts are used to compute estimates of the proposed coverage model. Several examples are discussed in order to illustrate the proposed modeling approach.

1. Introduction

Fault tolerance can be defined as the ability of a system to adapt and compensate in a systematic way to random component faults, and keep delivering completely or partially the functionality for which it was designed (Laprie, 1991). Fault tolerance is paramount in control system design for safety- and mission-critical applications. For a specific fault, fault tolerance can be measured through the notion of *fault coverage*, which can be defined as the conditional probability that, given a fault that alters the system structure occurs, the system is able to recover and keeps functioning. The notion of fault coverage was first introduced in the field of fault-tolerant computing (see, e.g., Arnold, 1973, Bouricius, Carter, & Schneider, 1969), and there is an extensive literature on fault coverage modeling in this field. In this regard, most fault coverage models proposed in fault-tolerant computing are developed using a probabilistic characterization of the fault mechanisms and recovery process; in particular, Markovian models are commonly used (see, e.g., Dugan & Trivedi, 1989, Pradhan, 1995 and the references therein). For example, in Stiffler and Bryant (1982), a continuous-time Markov chain is utilized, where the states of the chain represent the possible outcomes – the system fails or recovers – after the fault occurrence. Then, fault coverage is obtained by computing the probability that the system is in a state of the chain that results in recovery from the fault. For a detailed discussion on fault coverage modeling in fault-tolerant computing, the reader is referred to Dugan and Trivedi (1989) and Pradhan (1995).

In contrast with the approaches used in fault-tolerant computing, this paper focuses on fault coverage modeling in nonlinear dynamical systems that can be described by the standard state-space representation used in control theory, $\dot{x} = f(x, u)$, with state x and input u. Throughout this paper, we assume that the system input u is unknown but bounded, which could correspond to an external disturbance or to some uncertainty in the system operation, e.g., uncertainties in the load of an electric power system (Sauer & Pai, 1998). The vector field f is defined by the components constituting the system and how these are interconnected. When a fault occurs, the vector field f is altered, resulting in a new vector field $\hat{f}$. The performance requirements the system has to fulfill are modeled as a certain region in the system state-space to which the system state x is confined. Fault coverage is then defined as the (conditional) probability that the system state x remains inside the region specified by the performance requirements for all times. The goal of the paper is to provide a method for obtaining fault cover-

☆ The work of M.A. Müller was supported by the Fulbright Commission and by the German Research Foundation (DFG) within the Priority Programme 1305 "Control Theory of Digitally Networked Dynamical Systems" and within the Cluster of Excellence in Simulation Technology (EXC 310/1) at the University of Stuttgart. The work of A.D. Domínguez-García was supported in part by the National Science Foundation (NSF) under Career Award ECCS-CAR-0954420. The material in this paper has not been presented at any conference. This paper was recommended for publication in revised form by Associate Editor Fen Wu under the direction of Editor Roberto Tempo.

E-mail addresses: matthias.mueller@ist.uni-stuttgart.de (M.A. Müller), aledan@ILLINOIS.EDU (A.D. Domínguez-García).

[1] Tel.: +1 217 333 3953; fax: +1 217 333 1162.

0005-1098/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2012.04.007

age estimates that can be analytically formulated using the system state-space representation.

The definition of fault coverage adopted in this paper is relevant in various contexts. In a fail-safe application, it might be the case that violating the given performance requirement bounds might be catastrophic, e.g., an aircraft going beyond a critical angle of attack might cause the aircraft to stall. In such a case, it is desirable to guarantee that the system never violates the performance requirements, i.e., to ensure that fault coverage is equal to one independently of the input distribution. However, there might be other applications where performance requirements are "soft" constraints and violating them may cause the system to operate in some degraded mode without catastrophic consequences. In this case, assuming some knowledge of the input distribution, if fault coverage is smaller than one, it does not mean that the system cannot survive to a particular fault, but that after the fault, it might deliver its functionality in a degraded fashion (without necessarily failing catastrophically). For example, in an AC power system, performance requirements impose constraints on frequency deviations, and it is always desirable to keep the frequency within some acceptable range, e.g., 59.6–60.4 Hz. However, temporarily violating this requirement does not mean the system fails catastrophically (Sauer & Pai, 1998).

Modeling faults as a change in the vector field f is common when designing fault-tolerant controllers (see, e.g., Gao & Ding, 2007, Mao, Jiang, & Shi, 2010), and fault detection filters (see, e.g., Willsky, 1976). The purpose of fault-tolerant controllers is to ensure stability of the considered system despite the possible occurrence of certain faults. Different methods in achieving this have been proposed in the literature, both for linear and nonlinear systems (see, e.g., Bonivento, Isidori, Marconi, & Paoli, 2004, Gao & Ding, 2007, Mao et al., 2010, Qu, Ihlefeld, Jin, & Saengdeejing, 2003, Yang, Jiang, & Staroswiecki, 2009 and the references therein). Considering this, the work in this paper can be seen as a method to quantify the effectiveness of such fault-tolerant controllers, i.e., to give a quantitative expression for the probability that the closed-loop system meets the performance requirements for all times, including the transient phase after the fault occurrence, even if additional uncertainties/disturbances are present.

Fault coverage in the context of fault-tolerant control has also been considered in Wu (2004), where a two-level architecture consisting of the considered dynamical system at the lower level and a discrete state Markov process at the upper level, representing the set of possible failures, was used. There, the probabilistic nature of fault coverage stems from a possible uncertainty in the parameters describing a certain fault, and the coverage is intended for decision making in order to prevent the system from entering a state corresponding to a permanent system failure. The framework considered in this paper is significantly different as we explicitly make use of the system dynamics, and utilize reachability analysis techniques to compute fault coverage estimates. In particular, in our framework the probabilistic nature of fault coverage is due to the distribution of the system state at the time of fault occurrence corresponding to the distribution of the input, and not due to parameter uncertainties corresponding to fault occurrences. Our framework provides an analytically tractable method for computing fault coverage estimates, which in turn can directly be incorporated when designing fault-tolerant controllers.

The idea of obtaining fault coverage estimates via the state space representation of the considered system, which is used in this paper, was already employed in Dominguez-Garcia, Kassakian, and Schindall (2009) for linear system dynamics. However, the techniques used for computing fault coverage estimates in the nonlinear case are substantially different from the linear case. While ellipsoidal-based reachability analysis techniques are used in Domínguez-García et al. (2009), in this paper, we use input-to-state stability (ISS) notions.

The structure of this paper is as follows. In Section 2, the system model is presented. This section also provides some background on input-to-state stability (ISS) notions, which are key in the development of the proposed fault coverage model. In Section 3, the formal definition of the proposed fault coverage model is given. Section 4 provides analytically tractable methods to compute estimates of the proposed fault coverage model. Section 5 presents several examples that illustrate the ideas developed in the previous sections. Concluding remarks are presented in Section 6.

2. Preliminaries

In this section, we introduce the dynamical system model used throughout the paper, and specify the performance requirements the system is supposed to fulfill. Namely, we assume that our system is described by a state-space representation with an unknown-but-bounded input, and the performance requirements constrain the system trajectories to a region of the state-space defined by a symmetric polytope.

2.1. Fault-free system dynamics

Let the dynamics of a system operating with no faults be represented by

$$\begin{aligned} &\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0, \\ &u(t) \in B_u = \{u : |u| \leq u_{\max}\}, \end{aligned} \tag{1}$$

where the state $x \in \mathbb{R}^n$, the input $u \in \mathbb{R}^m$, $x_0 \in \mathbb{R}^n$ and $u_{\max} \geq 0$. The third equation implies that the input signal u (measurable and locally bounded) is contained in a ball with radius $u_{\max}$ for all t. Assume that $f(\cdot, \cdot)$ is locally Lipschitz and the unforced system $\dot{x} = f(x, 0)$ has an asymptotically stable equilibrium point at the origin. We assume that the system (1) is forward complete, i.e., the solution $x(t)$ exists for all $t \geq 0$ and it will be contained in the reachable set $\mathcal{R}(t)$.

2.1.1. Performance requirements

If the system is properly designed, it must meet some performance requirements. These requirements constrain the state-vector x to some region of the state-space Φ. We assume that Φ is given by a symmetric polytope, defined by

$$\Phi = \{x : |\pi_i^T x| \leq 1, \ i = 1, 2, \ldots, p\}, \tag{2}$$

with $\pi_i \in \mathbb{R}^n$. Then, for every $u(\cdot)$ with $u(t) \in B_u$ for all $t \geq 0$, in order for the system to deliver its intended function, it has to hold that $x(t) \in \Phi$ for all $t \geq 0$, i.e., $\mathcal{R}(t) \subseteq \Phi$ for all $t \geq 0$.

Remark 1. The subsequent analysis also works if more general constraints than those in the form of symmetric polytopes are considered. However, in this case the computation of fault coverage might become more complex or even computationally intractable. Furthermore, performance requirements resulting in polytopical state constraints are fairly general and include many practical problems. □

2.2. System dynamics after a fault

Let T be a random variable representing the time to fault occurrence. This fault alters the vector field f in (1), resulting in a new vector field $\hat{f}$. Let τ be a realization of T. Then, after a fault, the system state space representation is given by

$$\begin{aligned} &\dot{x}(\hat{t}) = \hat{f}(x(\hat{t}), u(\hat{t})), \quad x(\hat{t} = 0) = x(t = \tau) \in \mathcal{R}(\tau), \\ &u(\hat{t}) \in B_u = \{u : |u| \leq u_{\max}\}, \end{aligned} \tag{3}$$