Reliability Engineering and System Safety 123 (2014) 158–170

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

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

A Stochastic Hybrid Systems framework for analysis of Markov reward models

S.V. Dhople [a,*,1], L. DeVille [c,2], A.D. Domínguez-García [b,1]

[a] Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, United States
[b] Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, United States
[c] Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, United States

ARTICLE INFO

Article history:
Received 17 June 2013
Received in revised form
21 October 2013
Accepted 31 October 2013
Available online 9 November 2013

Keywords:
Markov availability models
Markov reliability models
Reward models
Performability analysis
Stochastic hybrid systems

ABSTRACT

In this paper, we propose a framework to analyze Markov reward models, which are commonly used in system performability analysis. The framework builds on a set of analytical tools developed for a class of stochastic processes referred to as Stochastic Hybrid Systems (SHS). The state space of an SHS is comprised of: (i) a discrete state that describes the possible configurations/modes that a system can adopt, which includes the nominal (non-faulty) operational mode, but also those operational modes that arise due to component faults, and (ii) a continuous state that describes the reward. Discrete state transitions are stochastic, and governed by transition rates that are (in general) a function of time and the value of the continuous state. The evolution of the continuous state is described by a stochastic differential equation and reward measures are defined as functions of the continuous state. Additionally, each transition is associated with a reset map that defines the mapping between the pre- and post-transition values of the discrete and continuous states; these mappings enable the definition of impulses and losses in the reward. The proposed SHS-based framework unifies the analysis of a variety of previously studied reward models. We illustrate the application of the framework to performability analysis via analytical and numerical examples.

1. Introduction

Continuous-time Markov chains (CTMCs) are commonly used for system reliability/availability modeling in many application domains, including: computer systems [12,36,51], communication networks [33,37], electronic circuits [6,52], power and energy systems [1,2,10,34], and phased-mission systems [30,58]. A Markov reward model is defined by a CTMC, and a reward function that maps each element of the Markov chain state space into a real-valued quantity [41,50,55]. The appeal of Markov reward models is that they provide a unified framework to define and evaluate reliability/availability measures that capture system performance measures of interest; in the literature, this is typically termed *performability analysis* [4,38–40,48,50,57,23,47,31,22,35]. In this paper, we propose a framework that enables the formulation of very general reward models, and unifies the analysis of a variety of previously studied Markov reward models. The framework foundations are a set of theoretical tools developed to analyze a class of stochastic processes referred to as Stochastic Hybrid Systems (SHSs) [26], which are a subset of the more general class of stochastic processes known as Piecewise-Deterministic Markov processes [11].

The state space of an SHS is comprised of a *discrete state* and a *continuous state*; the pair formed by these is what we refer to as the *combined state* of the SHS. The transitions of the discrete state are stochastic, and the rates at which these transitions occur are (in general) a function of time, and the value of the continuous state. For each value that the discrete state takes, the evolution of the continuous state is described by a stochastic differential equation (SDE). The SDEs associated with each value that the discrete state takes need not be the same; indeed, in most applications they differ significantly. Additionally, each discrete-state transition is associated with a reset map that defines how the pre-transition discrete and continuous states map into the post-transition discrete and continuous states. Within the context of performability modeling, the set in which the discrete state takes values describes the possible configurations/modes that a system can adopt, which includes the nominal (non-faulty) operational mode, but also those operational modes that arise due to faults (and repairs) in the components that comprise the system.

* Corresponding author. Tel.: +1 612 624 8837.
E-mail addresses: sdhople@UMN.EDU (S.V. Dhople), rdeville@ILLINOIS.EDU (L. DeVille), aledan@ILLINOIS.EDU (A.D. Domínguez-García).

[1] The work of these authors was supported in part by the National Science Foundation (NSF) under Career Award ECCS-CAR-0954420.
[2] The work of this author was supported in part by NSF under Award CMG-0934491.

0951-8320/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ress.2013.10.011

The continuous state captures the evolution of some variables associated with system performance, and as such, can be used to define reward measures that capture a particular performance measure of interest. Finally, the reset maps can define instantaneous gains and losses in reward measures that result from discrete-state transitions associated with failures/repairs.

In order to fully characterize an SHS-based reward model, we need to obtain the distribution of the combined state. However, this is an intractable problem in general, due to the coupling between the evolution of the discrete and continuous states and the presence of reset maps. In fact, this problem can only be solved in a few special cases. For instance, if we assume that the discrete state does not depend on the continuous state, the evolution of the former can be written as a CTMC; and as such, its probability distribution is fully characterized by the solution of the Chapman–Kolmogorov equations. However, unless we also assume that the resets do not change the value of the continuous state, it is not straightforward to obtain the continuous-state probability distribution. Given the difficulty in obtaining the distribution of the combined state, we settle for a method that allows the computation of any arbitrary number of their moments. To this end, we rely on the extended generator of the SHS, which together with Dynkin's formula can be used to obtain a differential equation that describes the evolution of the expectation of any function of the combined state, as long as such a function is in the domain of the extended generator. Following the approach outlined in [25,26], we show that under certain general assumptions, monomial functions are always in the domain of the extended generator, and thus, Dynkin's formula holds. Additionally, for SHS where the reset maps, transition rates, and the vector fields defining the SDEs are polynomial, the generator maps the set of monomial functions to itself. Therefore, Dynkin's formula gives a closed set of ordinary differential equations (ODEs) that describes the evolution of each moment in terms of the values of the other moments. Since there are infinitely many monomial functions, this formally produces an infinite-dimensional system of ODEs in what is referred to in the stochastic process literature as a *closure problem*.

The examples and case studies in this work demonstrate how the proposed SHS-based framework applies to reward models where the rate at which the reward grows is: (i) constant—this case is referred as the *rate reward model* [49], (ii) governed by a first-order linear differential equation—we refer to this case as a *first-order reward model*, and (iii) governed by a linear SDE—this case is referred as the *second-order reward model* [3,28]. As demonstrated in Section 3.1, the SHS-based framework can specify even more general reward models, but we restrict our attention to the above cases as they have been previously studied in the literature; this allows us to validate and verify our results. We will show that the structure of the standard reward models described above is such that there are finite-dimensional truncations of the ODEs governing the moment evolution that are closed, i.e., there are finite subsets of moments such that the evolution of any member of this subset is a function only of the other members of this subset. In other words, these conventional reward models do not lead to a closure problem, and we only have to solve a finite-dimensional ODE to determine the evolution of the reward moments.

Several numerical methods have been proposed to compute the reward distributions for rate reward models (see, e.g., [20,45,48,42,53,59,8] and the references therein). However, for more general reward models, e.g., second-order reward models with impulses and/or losses in the accumulated reward, it is very difficult to obtain explicit, closed-form, analytical solutions for the partial differential equations (PDEs) that describe the evolution of the reward distributions [27]. In practice, in order to analyze such reward models, numerical methods are utilized to integrate the PDEs governing the evolution of the accumulated reward probability density function [13,27] (see also [14,54] for discussions on specific reward modeling and analysis software packages). It is worth noting that systems with deterministic flows and random jumps in the state have been widely studied in the nuclear engineering community (in light of the description above, these are a type of SHS). For instance, Chapman–Kolmogorov equations with appropriate Markovian assumptions are utilized to derive the PDEs that govern the continuous states in [16,15,17]. However, even in this body of work, it has been acknowledged that closed-form analytical solutions to the PDEs can be derived only for simple models [15].

An alternative to numerical integration for characterizing the distribution of the reward is to compute its moments, which then can be used, e.g., to compute bounds on the probabilities of different events of interest using probability inequalities. In this regard, a number of methods have been proposed in the literature for computing moments in reward models. For example, techniques based on the Laplace transform of the accumulated-reward distribution are proposed in [21,28,29,49]. In [41], the first moment of the accumulated reward in these models is computed following a method based on the frequency of transitions in the underlying Markov chain. A numerical procedure based on the uniformization method is proposed to compute the moments of the accumulated reward in [9]. Methods from calculus of variations are used to derive differential equations that provide moments of rewards for rate-reward models in [46]. In the same vein of these earlier works, the SHS-based framework proposed in this paper provides a method to compute any desired number of reward moments. The advantages of the SHS approach are two-fold: (i) it provides a unified framework to describe and analyze a wide variety of reward models (even beyond the rate-, first-, and second-order reward models that our case studies focus on), and (ii) the method is computationally efficient as it involves solving a linear ODE, for which there are very efficient numerical integration methods.

The remainder of this paper is organized as follows. In Section 2, we provide a brief overview of Markov availability/reliability and reward models. In Section 3, we describe fundamental notions of SHS, and demonstrate how the Markov reward models studied in this work are a type of SHS. Case studies are discussed in Section 4, while Section 5 illustrates the moment closure problem in SHS. Concluding remarks and directions for future work are described in Section 6.

2. Preliminaries

In this section, we provide a brief overview of Markov availability and reliability models, as well as Markov reward models; while in the process, we introduce some relevant notation and terminology used throughout the paper. For a detailed account on these topics, interested readers are referred to [56,50].

2.1. Markov availability and reliability models

Let $Q(t)$ denote a stochastic process taking values in a finite set $\mathcal{M}$; the elements in this set index the system *operational modes*, including the nominal (non-faulty) mode and the modes that arise due to faults (and repairs) in the components comprising the system. The stochastic process $Q(t)$ is called a Continuous-Time Markov Chain (CTMC) if it satisfies the Markov property, which is to say that

$$\Pr\{Q(t_r)=i|Q(t_{r-1})=j_{r-1},\ldots,Q(t_1)=j_1\}=\Pr\{Q(t_r)=i|Q(t_{r-1})=j_{r-1}\}, \quad (1)$$