



Analysis of stochastic Petri nets with signals

Andrea Marin^{a,*}, Simonetta Balsamo^a, Peter G. Harrison^b

^a Università Ca' Foscari Venezia DAIS, Italy

^b Imperial College, Department of Computing, London, United Kingdom

ARTICLE INFO

Article history:

Received 15 December 2010
Received in revised form 28 June 2012
Accepted 29 June 2012
Available online 13 July 2012

Keywords:

Stochastic Petri nets
Product-form
G-networks

ABSTRACT

Product-form models facilitate the efficient analysis of large stochastic models and have been sought after for some three decades. Apart from the dominating work on queueing networks, some product-forms were found for stochastic Petri nets (SPNs) that allow fork-join constructs and for queueing networks extended to include special customers called *signals*, viz. *G*-networks. We appeal to the Reversed Compound Agent Theorem (RCAT) to prove new product-form solutions for SPNs in which there are special transitions, the firings of which act in a similar way to signals in *G*-networks, but which may be generated by synchronised firings (or service completions) and may affect several places simultaneously. We show that SPNs with signals are strict generalisations of *G*-networks with negative customers, triggers and catastrophes, and illustrate with copious examples.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Stochastic modelling has provided powerful methods for the performance evaluation of both hardware and software architectures. More specifically, models based on Continuous Time Markov Chains (CTMCs) play a pivotal role in this context and have been applied successfully to the performance evaluation of various real-world systems (see e.g. [1]). However, it is well-known that even structurally simple models may have a state space with cardinality that makes it infeasible to apply standard techniques in the computation of the stationary (equilibrium) probability distribution of the state. During the past three decades, considerable attention has been devoted to the identification of classes of stochastic models for which the stationary distributions are in *product-form*. This property essentially decomposes a model \mathcal{M} into a set of N sub-models $\{\mathcal{M}_1, \dots, \mathcal{M}_N\}$, the stationary distributions of each of which can be computed efficiently. At equilibrium, let $\mathbf{m} = (m_1, \dots, m_N)$ denote a state of \mathcal{M} , where the component m_i is the state of sub-model \mathcal{M}_i , $1 \leq i \leq N$. Then, \mathcal{M} is in product-form if the stationary probability $\pi(\mathbf{m})$ of each state \mathbf{m} is proportional to the product of the stationary probabilities $\pi_i(\mathbf{m}_i)$ of each state m_i in sub-model \mathcal{M}_i , i.e.:

$$\pi(\mathbf{m}) \propto \prod_{i=1}^N \pi_i(m_i)$$

Product-form theory has greatly enhanced the tractability of large stochastic models, largely due to the seminal results proved in queueing theory, such as Jackson's theorem and its multi-class extension, the BCMP theorem [2,3].

Another important class of models for which the underlying processes are CTMCs is that of stochastic Petri nets (SPNs). Roughly speaking, the structure of a SPN consists of places, transitions and arcs, the state being given by the numbers of tokens in each place. Arcs connect places to transitions and transitions to places. A state-transition event occurs when a transition fires, after a negative exponentially distributed random time, causing tokens to move from its input places to its

* Corresponding author.

E-mail addresses: marin@dsi.unive.it (A. Marin), balsamo@dsi.unive.it (S. Balsamo), pgh@doc.ic.ac.uk (P.G. Harrison).

output places. SPNs are an important formalism for the performance evaluation of complex systems mainly because they are able to describe, in a natural way, concurrency and synchronization, such as the fork-join construct and parallel computation. In the late 1980s, product-form solutions were obtained for SPNs such that synchronisation, batch token movements and resource competition were allowed [4,5]. In common with Jackson's theorem and the BCMP theorem, these product-forms follow from a set of linear equations that yield the correct parameterisation of associated sub-models, namely, the routing process invariant measures, analogous to the traffic equations of queueing networks.

In the early 1990s, the discovery of a novel product-form model, the *G*-network, showed that the traffic equations become non-linear when negative customers are introduced [6,7]. Differently from the ordinary customers of standard queueing networks, at a negative customer arrival epoch, the queue loses a positive customer if any is present (if the queue is empty, the negative customer simply vanishes). Since this seminal research, a number of new product-form queueing networks with quite sophisticated behaviours have been found, in which the concept of negative customer is generalised to that of a *signal*, which may act as a trigger that moves a customer from the queue of the arrival (when non-empty) to another queue, chosen probabilistically, or as a catastrophe, which empties the queue. Accessible surveys are given in the books [8,9], the latter featuring some telecommunications applications; see also [10].

In this paper, we appeal to the Reversed Compound Agent Theorem (RCAT) of [11] to derive new product-form solutions for SPN models in which signals are introduced. We first define a set of transitions with a simple semantics that describes signals in a compact way. We then prove that a SPN with such transitions may always be transformed into an equivalent SPN with only standard exponential transitions, but with inhibitor arcs as well. RCAT is a very general result for the derivation of product-forms in compositional models using a modular approach. This is because it is never necessary to consider the CTMC underlying the *whole* model in order to establish the state distribution at equilibrium. Moreover, RCAT ensures *a priori* the product-form solution of models in which some components with complicated behaviours (not necessarily specified as SPNs) interact with SPNs having signals—see, for example, [12].

The parameterisation of the sub-models is given by the solution of RCAT's system of *rate equations*, which are equivalent to the traffic equations in the case of Jackson networks [11]. Differently from *G*-network signals, SPN signals may be generated by synchronised job completion events, and may affect a set of places chosen probabilistically. Another significant difference from what is known in *G*-networks is that SPNs with signals may have rate-dependent product-form conditions. We prove that the product-form solution holds for models with both triggers (including negative customers) and catastrophes. As one might expect, in SPNs with signals, RCAT's rate equations form, in general, a non-linear system. In the case of *G*-networks, the existence of a solution to their non-linear traffic equations is considered in [13] by a generic method using Brouwer's fixed point theorem [14]. In a similar fashion, an analogous result on the existence and uniqueness of a solution to RCAT's rate equations is given in [15]. We show that the existence of a solution to the rate equations obtained in this paper follows from the latter. Our method is computationally efficient because it exploits a decomposition of a net into a set of *building blocks*, the analysis of which can be carried out independently of each other, once the set of rate equations is solved.

Related works. Since they were first introduced in [6,7,16], *G*-networks have proved to be a valuable tool for performance evaluation and optimisation purposes. For instance, in [17], they are used to model a multimedia server system in which customers require “documents” that are transmitted on the shared media. Several other applications of *G*-networks have been presented in the literature by various authors (see [18–21]). Following on from the seminal works of [6,7], a number of extensions have been proposed. The aforementioned triggers were introduced in [22] and generalised to the signals studied in [23], which can remove a batch of positive customers from one queue whilst adding a positive one to another. According to [24], we call a signal that completely flushes a node of all its customers at the arrival epoch a *catastrophe*. Triggers and signals as presented in [22,23] play an important role in the novel results that we introduce in this paper. In fact, as observed in [22], these mechanisms can indeed be useful for representing certain token movements in stochastic Petri nets. Other extensions are important in queueing network theory, particularly those with multiple classes of positive customers or signals [25,26] and the so-called *reset* customers of [27,20]. Resets move from one node to another and, on arriving at an empty queue, change the number of enqueued customers from zero to some random value that has a probability distribution that coincides with the steady-state distribution of the number of customers in that queue. Further properties of *G*-Networks with resets were considered in [28], where the notions of *stationary equivalence* and *flow equivalence* were considered and it was shown that, for this class of models, the latter implies the former.

Product-forms in stochastic Petri nets have also received considerable attention, see e.g. [4,5,29,30]. One of the interesting aspects of the models studied in these papers is that synchronised token arrivals and departures may lead to rate-dependent product-form conditions. In [4,5] the rate conditions arise from the application of a so-called “rank theorem”, in which the rank of a matrix of rates must be checked, whereas in [30], these conditions can be checked in a modular way as a consequence of a structural decomposition of the model. This helps in giving a physical interpretation to conditions that would otherwise be purely algebraic. In relation to other product-form models, we can say first that SPNs with signals generalise *G*-networks with negative customers, triggers and catastrophes [7,22,23,31]. Moreover, they subsume the product-form SPN of [5], without batch token movements. It is worthwhile noting that, differently from [4,5], we deal mainly with open models, since triggers may cause token destructions, leading to an empty network in the long-term.

The problem of synchronised arrivals in *G*-networks has also been addressed in [20] and, more recently, in [32]. The present paper differs from these in that it deals with exact analyses and does not need to introduce a network perturbation in order to derive a product-form, stationary state distribution. Moreover, interpreting the places of a SPN as queues, and

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات