Information and Computation 243 (2015) 50–74

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

A maximal entropy stochastic process for a timed automaton[☆,☆☆]

Nicolas Basset[a,b,*]

[a] LIGM, University Paris-Est Marne-la-Vallée and CNRS, France
[b] LIAFA, University Paris Diderot and CNRS, France

ARTICLE INFO

Article history:
Received 31 August 2013
Available online 12 December 2014

Keywords:
Timed automata
Maximal entropy
Stochastic process

ABSTRACT

Several ways of assigning probabilities to runs of timed automata (TA) have been proposed recently. When only the TA is given, a relevant question is to design a probability distribution which represents in the best possible way the runs of the TA. We give an answer to this question using a maximal entropy approach. We introduce our variant of a stochastic model, the stochastic process over runs, which permits to simulate random runs of any given length with a linear number of atomic operations. We adapt the notion of Shannon (continuous) entropy to such processes. Our main contribution is an explicit formula defining a process Y^* which maximizes the entropy. This ensures that, among the stochastic process over runs of a given TA, Y^* is the one that permits to sample runs of the TA in the most uniform way possible. Hence, our method could be used in a statistical model checking framework, providing a non-trivial yet natural way to generate runs in a quasi-uniform manner (described in the article). The formula defining Y^* is an adaptation of the so-called Shannon–Parry measure to the timed automata setting. We also show that Y^* enjoys well known properties in ergodic and information theory, namely, Y^* is ergodic and satisfies an asymptotic equipartition property.

1. Introduction

Timed automata (TA) were introduced in the early 1990's by Alur and Dill [3] and have been extensively studied since then to model and verify the behaviours of real-time systems. In this context of verification, several probability annotations have been added to TAs (see references below). There are several reasons to add probabilities: this permits (i) to reflect in a better way physical systems which behave randomly, (ii) to reduce the size of the model by pruning out the behaviours of null probability [9] and (iii) to resolve non-determinism when dealing with parallel composition [21,26].

In most of previous works on the subject (see e.g. [14,15,19,21]), probability distributions on continuous and discrete transitions are given at the same time as the timed annotations. In these works, the choice for the probability functions

---

[☆] The support of Agence Nationale de la Recherche under the project EQINOCS (ANR-11-BS02-004) is gratefully acknowledged.

[☆☆] This article is an extended version of [10] (best student paper at ICALP 2013, track B). This work corresponds also to Chapter 4 of the PhD thesis [11].

[*] Correspondence to: LIGM, University Paris-Est Marne-la-Vallée and CNRS, France.
E-mail address: nbasset@liafa.univ-paris-diderot.fr.
URL: http://www.liafa.univ-paris-diderot.fr/~nbasset/.

http://dx.doi.org/10.1016/j.ic.2014.12.006
0890-5401/

is left to the designer of the model. However, she or he may want to provide only the TA and ask the following question: what is the "best" choice of the probability functions according to the given TA? Such a "best" choice must transform the TA into a random generator of runs in the least biased manner, i.e. it should generate the runs as uniformly as possible to cover with high probability the maximum of behaviours of the modelled system. More precisely, the probability for a generated run to fall in a set should be proportional to the size (volume) of this set (see [26] for a similar requirement in the context of job-shop scheduling). We formalise this question and propose an answer based on the notion of entropy of TA introduced in [7].

The theory developed by Shannon [37] and his followers permits one to solve the analogous problem of quasi-uniform path generation in a finite graph. This problem can be formulated as follows: given a finite graph G, how can one find a stationary Markov chain on G which allows one to generate the paths in the most uniform way? The answer is in two steps (see Chapter 1.8 of [30] and also Section 13.3 of [29]): (i) there exists a stationary Markov chain on G with maximal entropy, the so called Shannon–Parry Markov chain; and (ii) this stationary Markov chain allows one to generate paths quasi-uniformly.

In this article we lift this theory to the timed automata setting. We work with timed region graphs which are to timed automata what finite directed graphs are to finite state automata, that is, automata without labelling on transitions and without initial and final states. We define stochastic processes over runs of timed region graphs (SPOR) and their (continuous) entropy. This generalisation of Markov chains for TA is of independent interest: as far as we are aware, it is the first one which provides a continuous probability distribution on the start states. As a main result we describe a maximal entropy SPOR which is stationary and ergodic, and which generalises the Shannon–Parry Markov chain to TA (Theorem 4). Concepts of maximal entropy, stationarity and ergodicity can be interesting by themselves; here, we use them as the key hypotheses to ensure quasi-uniform sampling (Corollary 6). More precisely, the result we prove is a variant of the so called Shannon–McMillan–Breiman theorem, also known as asymptotic equipartition property (AEP).

Potential applications.

One possible application of our method is in probabilistic model checking. Almost sure model checking aims to decide if a model satisfies a formula with probability one or zero (e.g. [2,17]). Quantitative probabilistic model checking (e.g. [15, 21]), on the other hand, aims to compute the probability of a formula being satisfied, and compare it to a given threshold. A related method is the statistical model checking, which aims to estimate the probability of a formula being satisfied by applying statistical inference, for example hypothesis testing. It is mainly based on Monte Carlo simulation and returns results with a given confidence (see [21] and references therein). Using our stochastic process, a random simulation with a linear number of operations wrt. the length of the run can be achieved. Further work would be needed to incorporate the simulation of our maximal entropy process in a statistical model checking algorithm.

Our result could then be used to obtain a statistical model checking algorithm for TA based on "proportionality". The inputs of the problem are a timed region graph $\mathcal{G}$, a formula φ and a threshold $\theta \in [0,1]$. The question is whether the proportion of runs of $\mathcal{G}$ which satisfy φ is greater than θ or not. A recipe to address this problem would be as follows: (i) take as a probabilistic model $\mathcal{M}$ the timed region graph $\mathcal{G}$ together with the maximum entropy SPOR Y^* defined in Theorem 4; (ii) run a quantitative probabilistic model checking algorithm on the inputs $\mathcal{M}$, φ, θ (the output of the algorithm is yes or no depending on whether $\mathcal{M}$ satisfies φ with probability greater than θ or not); and (iii) return this as output to our problem.

The concepts handled in this article, such as stationary stochastic processes and their entropy, AEP, etc., come from information and coding theory (see [20]). Our work can be viewed as providing a probabilistic counterpart of the timed channel coding theory we have proposed in [4]. Another application of our method would be a compression algorithm for timed words accepted by a given deterministic TA.

Related work.

As mentioned above, this work generalises the Shannon–Parry theory to the TA setting. As far as we are aware, this is the first time that a maximal entropy approach is used in the context of quantitative analysis of real-time systems.

Our models of stochastic real-time systems can be related to numerous previous works. Almost-sure model checking for probabilistic real-time systems based on generalised semi Markov processes (GSMPs) was presented in [2] at the same time as the timed automata theory and by the same authors. This work was followed by [14,19], which address the problem of quantitative model checking for GSMPs under restricted hypotheses. GSMPs have several differences with TA. Roughly speaking, they behave as follows: in each location, clocks decrease until a clock is null, and at this moment an action corresponding to this clock is fired, the other clocks are either reset, unchanged or purely cancelled. Our probability setting is more closely inspired by [9,17,21], where probability densities are defined directly on the TA. Here we add the new feature of an initial probability density function on states.

In [21], a probability distribution on the runs of a network of priced timed automata is implicitly defined by a race between the components, each of them having its own probability. This allows a random simulation of runs in a non-deterministic structure without state-space explosion. The obtained probability does not approximate uniformity, and thus this method cannot be used to satisfy our objective.

Our techniques are based on the pioneering articles [7,8] on entropy of regular timed languages. In the latter article and in [4], an interpretation of the entropy of a timed language as information measure of the language was given.