Branching Markov Decision Processes with Real-Time Constraints*

Ashutosh Trivedi
IIT Bombay, India

Dominik Wojtczak
University of Liverpool, UK

Abstract— **In this paper we introduce Timed Branching Processes, a natural extension of (multi-type) Branching Processes where each entity is equipped with a finite set of private continuous variables, called clocks. Clocks grow uniformly with the same rate and using them various timing constraints can be imposed on the branching rules of the system, e.g. the way an entity reproduces (branches) can depend on its age. The branching can take place at any time point which does not invalidate any of the time constraints. By choosing these time points we can try to maximise or minimise some objective, e.g. the probability of extinction, which allows us to study the best and the worst behaviour of the system.**

Keywords– **Branching Processes, Markov Decision Processes, Timed automata, Population Dynamics**

I. Introduction

We study Timed Branching Processes (TBPs), a natural extension of (Multi-typed) Branching Processes (BPs). BPs are a natural model used for studying behaviour of population dynamics. A population consists of entities of various types (possibly many entities of the same type can coexist at the same time) and each of them branches after some time into a set (possibly empty) of entities of various types while disappearing itself. This assumption is natural, for instance, for annual plants that reproduce only at a specific time of the year. The set of offspring of an entity is chosen at random among many possibilities with some fixed distribution that depends only on the type of the entity that has branched. The type can describe fundamental differences between entities, e.g. stem cells are very different from regular cells, or it can correspond to some characteristics of the entities, such as their age or size. Although the entities coexist with each other, the BP model assumes that there is no interaction between them, so how they reproduce and for how long they live is the same as if they were the only entities in the system. This assumption greatly improves the computational complexity of the analysis of such models and is natural in situations where the population exists in an environment that has virtually unlimited resources to sustain the growth of the population, e.g. common situation for bacteria or insects.

II. Summary of the results

In this paper we show algorithms for the following problems. The problem of approximating the optimal probability of extinction can be done in exponential time. Moreover, the problem of approximating the optimal time-bounded probability of extinction, with time bound T, can also be done in exponential time. Finally, for the total reward objective, which e.g. can be used to compute the supremum of the expected number of entities of a given type created before the whole population becomes extinct, the optimal value can be computed in exponential time exactly. The exponential blow-up in the computational complexity of all these problems comes from the (boundary) region abstraction of the underlying Timed Automata. On the other hand, if all entities have just one private clock, the computational complexity remains the same as for BDPs without any clocks.

III. Related work

Branching Decision Processes (BDPs), a natural generalisation of BPs to controlled setting was studied before for discrete-time BPs in the OR literature (e.g. [5], [6]) and found applications in manpower planning, controlled queuing networks, management of livestock and epidemic control, among others. The focus of these works was on optimising the expected average or discounted reward over a run of the process, or optimising the population growth rate.

In the branching processes arising in biology, the probability that an entity will branch is usually dependent on its age. Age-dependent models of BPs were considered before in the literature (see, e.g. [3], [2]), but not in the presence of time constraints on transitions and invariants on states. TBPs can be used to perform the best/worst behaviour analysis of age-dependent continuous-time branching processes.

Timed automata, defined in [1], and its extension to probabilistic setting ([4]) are the most commonly used formalisms for describing systems that evolve in real-time. The use of timing constraints to describe the behaviour of the system is both intuitive and expressive enough to capture a large class of real-time systems occurring in practice, and yet many of the decision problems for these models have moderate computational complexity.

References

[1] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
[2] P. Haccou, P. Jagers, and V. A. Vatutin. *Branching Processes: Variation, Growth, and Extinction of Populations*. Cambrdige University Press, 2005.
[3] T. E. Harris. *The Theory of branching processes*. Springer, Berlin, 1963.
[4] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282, 2002.
[5] S. R. Pliska. Optimization of multitype branching processes. *Management Science*, 23(2):117–124, 1976.
[6] U. G. Rothblum and P. Whittle. Growth optimality for branching markov decision chains. *Mathematics of Operations Research*, 7(4):582–601, 1982.

* This work is based on our "Timed Branching Processes" paper which appeared in a computer science conference 7th Quantitative Evaluation of Systems in 2010.