Theoretical Computer Science 736 (2018) 25–40

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Linear multi-objective drift analysis

Jonathan E. Rowe

School of Computer Science, University of Birmingham, Birmingham, B15 2TT, United Kingdom

A R T I C L E I N F O

Article history:
Received 16 August 2017
Received in revised form 1 February 2018
Accepted 8 February 2018
Available online 2 April 2018
Communicated by C. Witt

Keywords:
Run-time analysis
Multiplicative drift analysis
Multi-objective optimisation
Evolutionary algorithms
Voter model
Network epidemiology

A B S T R A C T

The tools of drift analysis enable bounds on run-times (or first hitting times) of stochastic processes, such as randomised algorithms, based on bounds on the expected progress at each time step in terms of a distance measure. In this paper, we generalise the *multiplicative drift* theorem to apply to processes which are best described by more than one distance function. We provide four examples to illustrate the application of this method: the run-time analysis of an evolutionary algorithm on a multi-objective optimisation problem; the analysis of a variant of the *voter* model on a network; a parallel evolutionary algorithm taking place on islands with limited migration; a synchronous network epidemiology model. In the latter example, we show that populations with limited neighbourhoods (such as the ring topology) are able to resist epidemics much more effectively than well-mixed populations.

1. Introduction

Drift analysis has now become a standard tool in the analysis of randomised search heuristics (see, for example, [9,10]). The fundamental idea is that, given an estimate of the expected progress towards some target at each time step, we can use this to bound the expected time for the process to reach that target. Upper and lower bound versions exist. We will consider the case of *multiplicative drift* [2], and will generalise this to the situation where the process concerned is described by more than one distance measure.

We begin, then, by considering a finite state space $\mathcal{X}$ and a random process $X_0, X_1, X_2, \ldots$ on this set. Suppose we are interested in the first hitting time, T, of a target set $S \subseteq \mathcal{X}$. We describe the progress towards this target set by means of a *distance function*.

Definition 1. Given a set $\mathcal{X}$ and a target set $S \subseteq \mathcal{X}$ a *distance function*, with respect to S, is a function $d : \mathcal{X} \to \mathbb{R}$ such that:

$$x \in S \Longrightarrow d(x) = 0 \text{ and}$$
$$x \notin S \Longrightarrow d(x) > 0.$$

It should be noted that we do not necessarily expect a distance function to be a metric.

We estimate the progress of the random sequence towards S by bounding the expected change in the distance. This then allows us to bound the expected first hitting time of the target.

E-mail address: J.E.Rowe@cs.bham.ac.uk.

https://doi.org/10.1016/j.tcs.2018.02.011
0304-3975/

Theorem 1 *(Multiplicative drift [1,2]). Let $X_0, X_1, X_2, \ldots$ be a random sequence from the finite set $\mathcal{X}$. Let $S \subseteq \mathcal{X}$ be a target set, and let $d : \mathcal{X} \to \mathbb{R}$ be a distance function with respect to S. Let T be the first hitting time of the target set S.*

Suppose, for all states $x \notin S$, we have

$$E[d(X_{t+1}) \mid X_t = x] \leq \delta d(x)$$

for some constant $0 < \delta < 1$. Then:

$$E[T \mid X_0 = x_0] \leq \frac{1 + \log(d(x_0)/d^*)}{1 - \delta}$$

where

$$d^* = \min\{d(x) \mid x \notin S\}.$$

The probability that the first hitting time exceeds

$$\frac{\log(d(x_0)/d^*) + c}{1 - \delta}$$

for any $c > 0$ is no more than

$$e^{-c}.$$

It should be noted that, while there is a sharp tail bound on the probability that the expected hitting time exceeds the upper bound, the bound can, in some situations, be rather conservative. A simple example that illustrates this is the *unexploded bomb* problem, in which there are n bombs that, at a given time step, each have probability p of exploding. If there are currently k unexploded bombs, then the expected number of unexploded bombs at the next time step is $(1 - p)k$. The multiplicative drift theorem them tells us that the expected time until they have all exploded is bounded above by $(1 + \log n)/p$. This is quite a good estimate when p is small, but as p approaches 1, the expected time actually approaches 1 and not $1 + \log n$. The upper bounds are better when the process described typically only makes small jumps.

Our goal in this paper is to extend the multiplicative drift theorem to the case where the underlying random sequence is best described by more than one distance function, and where the progress in each distance function considered separately is not necessarily monotonic. We will, in particular, look at the situation where the expected change in each distance function is a linear function of all the distances, and provide conditions under which first hitting time bounds can be proved. This turns out to be relatively straightforward for the case of two distance functions. Dealing with a larger number of functions requires sufficient structure in the problem to make progress. However, it is possible, in some situations, to deal even with an arbitrary number of dimensions. Our result will inherit the strengths (in terms of the sharp tail bound) and the weaknesses (in terms of the conservatism) of the original multiplicative drift theorem.

We will illustrate our result with four examples. Firstly, we will look at proving the run-time of a simple evolutionary algorithm on a multi-objective optimisation problem. Then we will look at the so-called *voter* model on a graph, examining a variant in which each voter has an inherent preference. A generalisation of that model allows us to analyse an evolutionary process taking place in parallel on multiple islands, with migration between neighbouring islands. Finally, we look at a synchronous model of epidemiology on a network and derive conditions for which epidemics will fail to take hold in the population. In particular, we will show that a population with very limited neighbourhood structure (we will consider a ring topology) is far better able to resist an epidemic than a well-mixed population.

2. Multi-objective drift

Now suppose our process is described by several distance functions $d_1, \ldots, d_m$ such that the intersection of the corresponding target sets is not empty. This intersection is now our new target set S. The expected change in each of these distances can depend on each other as follows. Let A be a non-negative $m \times m$ matrix. Then we suppose that, for each d_i and $x \notin S$, we have

$$E[d_i(X_{t+1}) \mid X_t = x] \leq \sum_j A_{i,j} d_j(x).$$

In order to find a measure of our overall progress, we will define a new distance function, which will be a convex combination of $d_1, \ldots, d_m$.

Suppose that A has a left eigenvector $\mathbf{v}$ which contains only real, positive entries. We can assume that this eigenvector is normalised so that its components sum to 1. Let λ be the corresponding eigenvalue.

$$\mathbf{v}A = \lambda \mathbf{v}.$$

We will use the normalised $\mathbf{v}$ to define our new distance function. Letting $\mathbf{d}(x)$ be the (column) vector $(d_1(x), \ldots, d_m(x))$, our new distance function is simply $\mathbf{v} \cdot \mathbf{d}(x)$.