RANDOM PROCESSES WITH REINFORCEMENT

Robin A. Pemantle
B.A., University of California at Berkeley
(1984)

Submitted to the
Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

Massachusetts Institute of Technology
August 5, 1988

Signature of author____

Department of Mathematics
August 5, 1988

Certified by____

Professor Persi Diaconis
Thesis Supervisor

Accepted by____

Professor Sigurdur Helgason, Chairman
Departmental Graduate Committee
Department of Mathematics

ABSTRACT

We consider a class of random processes that have a kind of reinforcement. The first chapter is devoted to an expository survey that delimits this class of processes, of which the Pólya urn process is prototypical. We then consider some generalizations. The generalizations in chapter three are urn models. The ones in chapters four and five are non-Markovian finite state processes, or alternatively Markov processes with very large state spaces. We derive two types of limit theorems for these processes. The first kind is essentially a strong law of large numbers, stating that the fractional occupation of each state converges. The second type is a characterization of the law of the limiting fractional occupation. We consider one version that has an infinite state space and derive conditions for recurrence and for transience of the process. Our results have the following consequence: any proposition logically implies a true proposition.

ACKNOWLEDGEMENTS

I would like to thank four of the best math teachers: Oscar Pemantle, Dan Bach, John Addison, Jr. and Dan Kan.

Special thanks are due to Persi Diaconis: not only has he been very generous with his time, but also his presence somehow causes me to do math twice as well. Thanks are also due to Rick Durrett for his extensive and very helpful criticism of chapter four, to Dick Dudley for reading and criticizing most of the other chapters and to John Palmieri for keeping me honest in the introduction. Thanks to the National Science Foundation for financial support, and to the Institute for Elementary Studies at Lake Tahoe as well as several universities in New Zealand for hospitality during summers and vacations. Lastly, I thank a source who wishes to remain anonymous for technical assistance.

Contents

Chapter 1

INTRODUCTION

In this paper we study a variety of random processes that can be termed *random processes with reinforcement.* One example, an edge-reinforced random walk, is described as follows: Imagine a person getting acquainted with a new city. She walks about the area near the hotel somewhat randomly, but tends to traverse the same blocks over and over as they become familiar. To model this, a random walk is defined on the vertices of an undirected graph in such a way that the probability of a transition from one vertex to another depends on the number of previous transitions along the connecting edge. Edge-reinforced random walks were introduced by Coppersmith and Diaconis in 1986 [CD]; they are discussed in chapter four.

It turns out that such problems bear similarities to the so-called *Pólya urn process*: an urn contains R red balls and B black balls. A ball is drawn from the urn and then put back along with Δ balls of the same color, where Δ is some fixed constant. This process of drawing and replacing is repeated ad infinitum. In chapter two we give a summary of known results on Pólya's urn and other similarly behaved processes. This chapter is a philosophical prerequisite for the rest of the paper, although mathematically the last three chapters are almost completely self-contained.

In chapter three we discuss generalizations of the basic Pólya urn process; for example, we allow Δ to change with time, or we allow more than two different colors for the balls in the urn. In the course of doing this, we prove an important lemma: a sequence of random variables cannot converge to a point with non-zero probability if certain hypotheses are satisfied which make this point a *nonattracting point* for the given sequence of random variables. This lemma is used later in chapter five.

As mentioned above, chapter four is devoted to edge-reinforced random walks. If the graph is finite, then this process exhibits Pólya urn-like behavior; results are quoted from [CD]. Otherwise, we deal with the case where the graph is infinite and acyclic (i.e., it contains no circuits). Then we can describe the random walk as embedded Pólya urn processes. Essentially, one can imagine a policeman with an urn standing at each vertex of the tree directing where you should go next by pulling balls out of the urn – whose labels are the various choices of where to go next, instead of the colors red and blue – and replacing them according to the Pólya urn scheme.

The useful (and surprising) result on Pólya urns is that they are equivalent to a mixture of independent, identically distributed draws. To be graphic: we can replace the policemen by clay tablets, each bearing an inscription giving the probabilities of the various choices from that vertex. There is a different tablet at each vertex, each being randomly chosen from a collection of all possible tablets, but once chosen they do not change, and they obviate the need to keep track of how many times the walk has traversed each edge. In other words, the walk is shown to be equivalent to a random walk in a random environment. The random environment can then be analyzed using large deviation estimates and we can get recurrence and transience conditions that are quite sharp.

Chapter five deals with the dual process to the one in chapter four, a vertex-reinforced random walk. Consider a complete graph with loops on a finite number of vertices and give each edge a non-negative weight. Define a Markov chain on the vertices of the graph by letting the probabilities of transitions along the edges leading away from a vertex be proportional to the weights of the edges. The vertex-reinforcement enters the picture by updating the transition probabilities at each step to favor those vertices that have previously been visited the most. This models a person getting acquainted with a new city – say Los Angeles – by car rather than on foot; in this case familiarity will reinforce visits to the same destination rather than journeys along the same route.

Just as edge-reinforcement induces Pólya urn-like behavior, so vertex-reinforcement induces Friedman urn-like behavior (for information on Friedman urn processes, see chapter 2). In particular, the fractional occupation vector converges under certain conditions to a limit which is random, but whose distribution is supported on a set of small dimension. To show that the vector converges, we find a scalar function of this vector that measures, in some sense, the correlation between occupations of vertices connected by edges with the largest weights. The key property of this function is that it always increases as the process evolves. Thus it serves as a Lyapunov function, whose existence prevents the fractional occupation vector from wandering around in circles without converging. This method proves that the vector must converge to a place where the Lyapunov function has no gradient, but in fact in most cases this point must be a maximum, not a minimum or a saddle point. The latter part of chapter five uses the lemma on nonattracting points from chapter three to obtain this further result, under some extra hypotheses on the eigenvalues of a certain matrix. We end with a few fun examples of vertex-reinforced random walk that are calculated using some elementary character theory.

The material in chapter four appears in [Pe1].