

Recursive Methods in Urn Models
and First-Passage Percolation

Henrik Renlund

Department of Mathematics
Uppsala University
UPPSALA 2011

Dissertation presented at Uppsala University to be publicly examined in Högssalen, Ångströmlaboratoriet, Lägerhyddsvägen 1, Uppsala, Friday, March 25, 2011 at 13:15 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract

Renlund, H. 2011. Recursive Methods in Urn Models and First-Passage Percolation. Department of Mathematics. *Uppsala Dissertations in Mathematics* 69. 30 pp. Uppsala. ISBN 978-91-506-2190-7.

This PhD thesis consists of a summary and four papers which deal with stochastic approximation algorithms and first-passage percolation.

Paper I deals with the a.s. limiting properties of bounded stochastic approximation algorithms in relation to the equilibrium points of the drift function. Applications are given to some generalized Pólya urn processes.

Paper II continues the work of Paper I and investigates under what circumstances one gets asymptotic normality from a properly scaled algorithm. The algorithms are shown to converge in some other circumstances, although the limiting distribution is not identified.

Paper III deals with the asymptotic speed of first-passage percolation on a graph called the ladder when the times associated to the edges are independent, exponentially distributed with the same intensity.

Paper IV generalizes the work of Paper III in allowing more edges in the graph as well as not having all intensities equal.

Keywords: stochastic approximation algorithm, generalized Pólya urn, limit theorem, first-passage percolation, rate of percolation, time constant

Henrik Renlund, Department of Mathematics, Mathematical Statistics, Box 480, Uppsala University, SE-75106 Uppsala, Sweden.

© Henrik Renlund 2011

ISSN 1401-2049

ISBN 978-91-506-2190-7

urn:nbn:se:uu:diva-145430 (<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-145430>)

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

- I Renlund, H. (2009) Generalized Pólya urns via stochastic approximation.
- II Renlund, H. (2011) Limit theorems for stochastic approximation algorithms.
- III Renlund, H. (2010) First-passage percolation with exponential times on a ladder. *Combin. Probab. Comput.*, 19(4): 593–601.
- IV Renlund, H. (2011) First-passage percolation on ladder-like graphs with heterogeneous exponential times.

Paper III is reprinted with permission from the publisher.

Cover by Ann McLeod Photography.

Contents

1	Introduction	7
2	Stochastic approximation algorithms	9
2.1	The Robbins-Monro algorithm	9
2.2	A definition	10
2.3	An urn model	11
2.4	A summary of Paper I	13
2.5	A summary of Paper II	16
3	Models of percolation	17
3.1	Percolation	17
3.2	First-passage percolation	18
3.3	A summary of Paper III	19
3.4	A summary of Paper IV	21
4	Summary in Swedish	23
5	Acknowledgements	27
	Bibliography	29

1. Introduction

This thesis deals with two different subjects in the subfield of probability theory called stochastic processes, namely *stochastic approximation algorithms* and *first-passage percolation*. An explanation for each follows in Section 2 and Section 3, respectively. Albeit quite disparate, a theme can be recognized in that our analysis of each process applies *recursive methods*.

In mathematics, a *recursion* may be contrasted with a *formula*. Consider the following sequence of real numbers $0, 1, 4, 9, 25, \dots$. If we denote the n 'th number in this sequence by a_n we can give a formula for each of these numbers, namely $a_n = n^2$. This way, we can immediately express an entry with any index without bothering about any other entry.

Other sequences are naturally expressed recursively. The famous sequence of Fibonacci numbers starts with $b_0 = 0$, $b_1 = 1$, and the rule that any entry of the sequence – with index greater than, or equal to, two – is the sum of the last two entries, i.e. $b_n = b_{n-1} + b_{n-2}$, for $n \geq 2$. To calculate the, say, 100th entry of the Fibonacci sequence, would require of us to also calculate the preceding 99 values, unless we can find a formula for b_n .

Some recursions do admit explicit formulas, e.g. the Fibonacci sequence, which can be expressed as $b_n = (\varphi^n - (1 - \varphi)^n)/\sqrt{5}$, for any $n \geq 0$, where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. The work related to first-passage percolation, Papers III and IV, are in part concerned with this very problem, namely; given a recursively defined sequence, determine the formula.

The first part of the thesis, Papers I and II, is concerned with random sequences whose behavior is determined recursively. Essentially, a sequence X_0, X_1, X_2, \dots is a *stochastic approximation algorithm* if there is a function f that determines the *average* behavior of the sequence, in the following sense: given X_n , the next value X_{n+1} is given, on average, by

$$X_{n+1} = X_n + f(X_n)/n. \tag{1.1}$$

Hence the process takes a “step” of length $1/n$ in the direction determined by $f(X_n)$ (again, on average). In the one-dimensional case, which is all that we shall be concerned with, we may simply think of f as being positive, negative or zero at any given point x . If the process is in a region where f is positive it will tend to increase, and similarly decrease when f is negative. In this way it will tend to locate the zeros of the function f , if such exist (and they do in all sensible applications).

2. Stochastic approximation algorithms

2.1 The Robbins-Monro algorithm

Let us introduce the concept of stochastic approximation algorithms via the 1951 paper “A stochastic approximation method” by Robbins and Monro [RM51] that initiated the study: consider an experiment in which we use x as an input to some system and get $m(x)$ as a response. In a deterministic system, measuring, say, the pull of two opposite electrical charges when they are at distance x from each other, we expect to be able to determine the function m from observations (conveniently ignoring errors from measurement). In a system prone to variation, the function m is not so easily observed. Consider for instance measuring the blood pressure as an effect of some drug that has been administered to a patient. We may think that in the population as a whole there is an average effect function m responding in a precise manner to a certain level x of the drug, but any one patient will most likely deviate from this. (Even if m describes the average behavior of that particular individual there is bound to be natural variation, as blood pressure depends on numerous factors.)

It is this situation that is considered in [RM51]; whenever we give input x to a system, we get as response

$$M(x) = m(x) + \text{“noise”},$$

where the noise *on average* is zero, but, typically, never exactly zero in any given experiment.

It is not the purpose to determine m but to establish what input yields a given response θ , i.e. what x solves the equation $m(x) = \theta$? For simplicity we assume that there is a unique solution x^* to this problem, and – even more restrictive – that

$$m(x) \begin{cases} < \theta, & x < \theta, \\ > \theta, & x > \theta. \end{cases} \quad (2.1)$$

The proposed method for finding x^* in [RM51] is an iterative procedure. Start with some arbitrary value X_0 as a first approximation. Create a sequence of refinements by, given the current approximation X_n , letting the next one be

$$X_{n+1} = X_n + \frac{1}{n} (\theta - m(X_n) + U_{n+1}). \quad (2.2)$$

In (2.2), U_{n+1} is the noise associated with the n 'th measurement. As we expect U_{n+1} to be zero, we expect X_{n+1} to move towards x^* . This is due to our assumption (2.1); if we think of U_{n+1} as being zero in (2.2), then what is in the parenthesis on the right hand side is positive when $X_n < x^*$ and negative when $X_n > x^*$. Hence, on average, X_{n+1} is greater than X_n when $X_n < x^*$ and smaller when it is not, i.e. it moves (on average) towards x^* .

Notice that the step lengths “ $1/n$ ” serves the purpose of, hopefully, making the process converge to the point x^* (without decreasing steps, the sequence X_0, X_1, \dots might eventually just be jumping back and forth around x^*), yet are large enough to allow the process to reach x^* , regardless of the initial value X_0 . (One may walk infinitely far using step lengths that decrease to zero at the rate of $1/n$, since this series – known as the harmonic series – diverges, i.e. $\sum_1^\infty 1/n = \infty$.)

In [RM51] sufficient criteria are given for the convergence (in probability) of X_0, X_1, X_2, \dots to x^* .

The study of related algorithms was continued the following year by Kiefer and Wolfowitz in “Stochastic estimation of the maximum of a regression function” [KW52], which sought not to estimate the zero of an unknown function, but rather the maximum.

Since then, similar algorithms have found many applications in areas such as signaling processing, resource allocation, system identification, reinforcement learning, neural networks, adaptive control, etc. An introduction to the subject, with a view towards applications as such mentioned above, can be found in [Bor08].

In general, a stochastic approximation algorithm is an adapted process $\{X_n, n \geq 0\}$ that evolves according to

$$X_{n+1} = X_n + \gamma_{n+1}(f(X_n) + U_{n+1}),$$

where γ_n are the step lengths, U_{n+1} is the mean zero noise and f is referred to as the *drift function*. ($f(x) = \theta - m(x)$ in the Robbins-Monro algorithm in (2.2).) In general, X_n need not be one dimensional. In the multidimensional case, it is fruitful to consider the process to be a discrete time approximation to a solution curve of the corresponding ordinary differential equation

$$x'(t) = f[x(t)].$$

A standard reference for this approach is the lecture notes [Ben99].

2.2 A definition

Papers I and II are concerned with processes according to the following definition.

Definition 1 Given a probability space $(\Omega, \mathbb{P}, \mathcal{F})$ and a filtration $\{\mathcal{F}_n, n \geq 1\}$, an adapted process $\{X_n, n \geq 0\}$ is said to be a stochastic approximation algorithm if $X_n \in [0, 1]$ and if there are adapted γ_n and U_n , and a function f such that

$$X_{n+1} = X_n + \gamma_{n+1}(f(X_n) + U_{n+1}), \quad (2.3)$$

where $|f|, |U_n|, n^2 \cdot |\mathbb{E}_n(\gamma_{n+1} \cdot U_{n+1})|$ and $n \cdot \gamma_n$ are bounded and $n \cdot \gamma_n$ is bounded away from zero, and $\mathbb{E}_n(\cdot)$ denotes conditional expectation with respect to the filtration, i.e. $\mathbb{E}_n(\cdot) = \mathbb{E}(\cdot | \mathcal{F}_n)$.

Some remarks concerning the definition.

- Our algorithm is bounded, specifically to the interval $[0, 1]$, although this is of less importance as one can scale any bounded process to this interval. Through relation (2.3), this naturally entails boundedness of $|f|, |U_n|$ and γ_n .
- Our step lengths are essentially $1/n$. However, in the applications we look at, these are typically random, although there are positive constants a and b such that $\gamma_n \in [a/n, b/n]$ a.s. for every n .
- The perturbation away from a deterministic system in (2.3) is

$$\gamma_{n+1} \cdot U_{n+1},$$

which is not required to have conditional expectation zero, but this expected value must tend to zero at the fast rate of $1/n^2$. In our applications, U_{n+1} typically has conditional mean zero, but is not independent of the step length γ_n , making the conditional expectation nonzero but vanishing quickly.

2.3 An urn model

Urn models in probability theory are common, a general reference to this subject is the book [JK77]. We are interested in a class of models derived as generalizations of the Pólya (or Pólya-Eggenberger) urn model that was described in a 1923 paper “Über die Statistik verketteter Vorgänge” [EP23]. The essence of these models is to consider repeated draws from an urn that contains balls of at least two colors. After each draw, the color is noted and the ball is replaced into the urn along with more balls, how many and of which colors is to depend on the color drawn according to some prescribed rule. The machinery of stochastic approximation algorithms is a suitable tool for analyzing the sequence of fractions of balls of a given color.

Our attention was brought to these models in connection to reinforcement learning. Consider a toy example; imagine that we repeatedly are faced with some situation that admits three different courses of action, which we label 1, 2 and 3. Each action is associated with a payoff, measured in nonnegative

integers. An adaptive algorithm that tries to learn the optimal payoff is the following (which can be implemented in a computer program): place a ball each of the colors white, black and turquoise, say, in an urn. Each time we are required to make a choice between the three options at our disposal, we draw a ball randomly from the urn with the rule that

- if a white ball is drawn, perform act 1,
- if a black ball is drawn, perform act 2, and
- if a turquoise ball is drawn, perform act 3.

Upon receiving the payoff k , which depends on the action taken, we return the drawn ball along with k additional balls of the same color. In this way the probability of repeating this chosen action at a later stage increases in relation to how beneficial the action was.

Consider e.g. the payoffs for actions 1, 2 and 3 to be 0, 10 and 50. Then the rational choice of action is that labeled 3 as it gives the highest payoff, and we certainly do not want to do action 1, as it gives no payoff whatsoever. Action 1 is not so much a threat to the optimal strategy, since we never increase the probability of this suboptimal choice. However, there is always the chance of “accidentally” choosing act 2 first. If this happens, the urn content is now 1 white, 11 black and 1 turquoise, so the probability of doing act 2 again the second time is $11/13 \approx 85\%$ and, if so, yet again $21/23 \approx 91\%$, and so on.

A natural question is thus, what is the probability of learning the optimal strategy? (It is in fact guaranteed.) One way to resolve this question is to monitor the sequence of fractions of turquoise balls – will this sequence tend to 1?

The example was given only to serve as some justification as to why it might be interesting to track the fraction of balls of some given color in an urn model that evolves by the repeated adding of balls.

To end this section, let us describe a recurring application of the results concerning stochastic approximation algorithms in Paper I and II. An urn has balls of two colors, white and black say. Let W_n and B_n denote the number of balls of each color, white and black respectively, after the n 'th draw and consider the initial values $W_0 = w_0 > 0$ and $B_0 = b_0 > 0$ to be fixed. After each draw we notice the color and replace it along with additional balls according to the *replacement matrix*

$$\begin{array}{cc} & \begin{array}{cc} \text{W} & \text{B} \end{array} \\ \begin{array}{c} \text{W} \\ \text{B} \end{array} & \left(\begin{array}{cc} a & b \\ c & d \end{array} \right), \end{array} \tag{2.4}$$

where a, b, c and d are nonnegative numbers (they need in fact not be integers, although if not it is somewhat more difficult to picture the draws). The replacement matrix (2.4) should be interpreted as; if a white ball is drawn it is replaced along with an additional a white and b black balls. If a black ball is drawn it is replaced along with an additional c white and d black balls.

2.4 A summary of Paper I:

Generalized Pólya urns via stochastic approximation

This paper is concerned with $\lim X_n$ as $n \rightarrow \infty$, where X_n throughout is a stochastic approximation algorithm according to Definition 1 of Section 2.2 on page 11. It should be noted that most results herein are fashioned after similar result from related models, these are pointed out in the article and will not be so in this summary, but let us just mention the three influential papers [Pem88], [Pem91] and [Pem07], all by Robin Pemantle. We will not talk about proofs in this summary, but most results are connected to martingale theory.

Let $Q_f = \{x : f(x) = 0\}$ denote the zero set of the drift function f .

A first thing to notice is that if f is continuous (at the boundary points at least) and if the process is free, in principle, to move about on the interval $(0, 1)$ then $f(0) \geq 0$ and $f(1) \leq 0$, as stated in Lemma 3 of Paper I. More precisely, this requires that neighborhoods of the points 0 and 1 are *attainable*, meaning that at any given entry X_n of the sequence X_0, X_1, \dots , there is a positive probability of a later entry X_{n+k} being arbitrarily close to any of these two points. This is a requirement that sensible applications will meet, and it shows that typically one should not expect Q_f to equal the empty set \emptyset (at least not for a continuous f).

The first theorem tells us where to look for limit points.

Theorem 1 *If f is continuous then $\lim_{n \rightarrow \infty} X_n$ exists almost surely and is in Q_f .*




Figure 2.1: Three different kinds of zeros. On the bottom, the arrows indicate the direction of the drift induced by the function plotted in the plane above.

However, one should not expect that all points of Q_f are possible limit points. We make the following categorization of points in Q_f , exemplified by Figure 2.1;

- A point x_u is called *unstable* if $f(x)(x - x_u) \geq 0$ whenever x is close to x_u . This means that the drift is locally pushing the process away from x_u (or not pushing at all), compare with Figure 2.1. If the inequality is strict when $x \neq x_u$, we say that x_u is *strictly unstable*.
- A point x_s will be called *stable* if $f(x)(x - x_s) < 0$ whenever $x \neq x_s$ is close to x_s . Locally, the drift pushes the process towards x_s from both directions, compare with Figure 2.1.
- A point x_t is called a *touchpoint* if $f(x) > 0$ for all $x \neq x_t$ close to x_t , or $f(x) < 0$ for all $x \neq x_t$ close to x_t . A touchpoint may be thought of as having one stable and one unstable side.

Heuristically, one should not expect the process to end up at an unstable point, since the noise will tend to push it out, and, once this has happened, it will tend to drift away. Indeed this is what happens, provided a lower bound on the conditional expected second moment of the noise terms exists.

Theorem 2 *If x_u is unstable and if $\mathbb{E}_n U_{n+1}^2 \geq C > 0$ holds whenever the process is close to x_u , then convergence to x_u is impossible.*

In our applications, the boundary $\{0, 1\}$ represents the points where the fraction of balls of a certain color is either 0 or 1. This means that the subsequent draw is deterministic and hence that there is no error term (the error term being the difference between what happens and what happens on average). In applications it is not impossible to have strictly unstable zeros on the boundary and, with no error term, Theorem 2 is inapplicable. The key ingredient in establishing non-convergence to such a point - when it can not be reached in a finite time - is an upper bound on the speed at which the second moment vanishes. Heuristically, one might think that as the error terms decay, there is an increasing tendency to follow the drift away from this point.

Theorem 3 *Suppose $x_u \in \{0, 1\}$ is strictly unstable, that $X_n > 0$, and*

$$\mathbb{E}_n U_{n+1}^2 \leq c_1 |X_n - x_u|, \quad [f(x)]^2 \leq c_2 |x - x_u|, \quad \text{and} \quad k \cdot |X_k - x_u| \rightarrow \infty.$$

Then, convergence to x_u is impossible.

Remark 1 *The assumption $k \cdot |X_k - x_u| \rightarrow \infty$ may look a bit strange, but is naturally fulfilled in our applications. In the case of $x_u = 0$, it simply translates to the assumption that balls of the color whose fraction we are monitoring should be drawn infinitely often. In the case $x_u = 1$ the same should be true of the opposite color.*

The bound on $f(x)$ – although satisfied in all the applications we have looked at – is superfluous, a stronger drift away from the unstable point should be even better, and could be dropped with a coupling argument.

Alas, we were unaware at the time¹ of related work by Tarrés, Pagès and Lambertson. In fact, Theorem 3 exists in a stronger form [Tar01], but with a different proof. The articles [LPT04], [Pag05] and [LP08] deal with an algorithm called the two-armed bandit, and contain interesting results in relation to unstable boundary points.

The next theorem confirms the intuition that stable points are possible limit points.

Theorem 4 *Suppose x_s is stable and that every neighborhood of x_s is attainable. Then there is a positive probability of convergence to x_s .*

The last result is that of touchpoints. It turns out that as long as the slope towards a touchpoint x_t is below a critical level, it may happen that the process converges towards x_t in such a way as to never exceed x_t . The theorem requires a technical condition similar to that of attainability, which we will omit here. Also, we state the result as if the touchpoint is stable from the left, as in Figure 2.1.

Let $b = \sup_n (n \cdot \gamma_n)$. (Such a number b exists by Definition 1.)

Theorem 5 *Suppose that $0 < f(x) < K(p - x)$ for some $K < \frac{1}{2b}$ whenever $x < p$ is close to p . Under an attainability-like condition, there is a positive probability of convergence to x_t .*

The above results are applied to the urn model described at the end of Section 2.3, with replacement matrix given by (2.4), page 12. If we require

$$\min\{a + b, c + d\} > 0,$$

then the fraction $Z_n = W_n/T_n$ of white balls is a stochastic approximation algorithm according to Definition 1, with drift function given by

$$f(x) = (c + d - a - b)x^2 + (a - 2c - d)x + c. \quad (2.5)$$

(The cases where $\min\{a + b, c + d\} = 0$ are easily handled separately.) It turns out that often a unique zero of f exists, and then necessarily $Z = \lim_n Z_n$ must equal this point by Theorem 1. If two zeros exist, then one of these is unstable and on the boundary, and the other is necessarily stable and Z will equal the latter. When $a = d$ and $b = c = 0$, then $f \equiv 0$ and our method reveals nothing about the distribution of Z . (It is however well known that Z then has a Beta distribution with parameters W_0/a and B_0/a .)

We also apply the result to an urn model whose evolution is determined by a simultaneous draw of two balls. This results in a drift function being a polynomial of degree 3, which admits touchpoints and multiple stable points.

¹This manuscript was prepared in 2009.

2.5 A summary of Paper II:

Limit theorems for stochastic approximation algorithms

This paper investigates the asymptotic distribution of X_n , properly scaled, where X_n throughout is a stochastic approximation algorithm according to Definition 1 of Section 2.2 on page 11. Assume that X_n tends to a stable point p . We must assume that f is differentiable at this point, so – for the purpose of this section – we may define *stable* by saying that $f'(p) < 0$. Now, as $Q_n = X_n - p$ tends to zero, one typically expects there to be some way to *inflate* Q_n – i.e. multiply by some increasingly large real number w_n – so as to make $w_n \cdot Q_n$ tend neither to 0 nor $\pm\infty$, but rather to stabilize on \mathbb{R} (the set of real numbers) according to some probability distribution. When this limiting distribution is the Gaussian, i.e. normal, distribution, we say that X_n satisfies a *central limit theorem*. In this paper we find circumstances in which a central limit theorem applies, as well as the proper scaling for some other cases, in which we can not identify the limiting distribution. The results are extensions of methods found in [Fab68] and [MP73].

So, we imagine that $X_n \rightarrow p$, p being stable, and that the drift function f is differentiable at p . Then we can write $f(x) = -h(x)(x - p)$ where $h(x)$ is continuous at p and positive close to, but not at, p . What largely determines the asymptotic behavior and scaling turns out to be

$$\hat{\gamma}_n = n\gamma_n h(X_{n-1}), \quad \text{and secondarily} \quad \hat{U}_n = n\gamma_n U_n.$$

Now, we go through the three cases dealt with in the paper. All assume the existence of $\hat{\gamma} = \lim_n \hat{\gamma}_n$.

- If $\hat{\gamma} > 1/2$ and $\mathbb{E}_n \hat{U}_{n+1}^2 \rightarrow \sigma^2 > 0$, then

$$\sqrt{n} \cdot (X_n - p) \xrightarrow{d} \text{N}\left(0, \frac{\sigma^2}{2(\hat{\gamma} - 1/2)}\right),$$

where \xrightarrow{d} denotes convergence in distribution, and the distribution is the Gaussian with mean 0 and variance $\sigma^2/2(\hat{\gamma} - 1/2)$.

- If $\hat{\gamma} = 1/2$ and $\mathbb{E}_n \hat{U}_{n+1}^2 \rightarrow \sigma^2 > 0$, then

$$\sqrt{\frac{n}{\ln n}} \cdot (X_n - p) \xrightarrow{d} \text{N}(0, \sigma^2),$$

if also $\hat{\gamma}_n - 1/2$ tends to zero faster than logarithmically.

- If $\hat{\gamma} \in (0, 1/2)$ and $\hat{\gamma}_n - \hat{\gamma}$ is bounded by some multiple of $|X_n - p + 1/n|$, then $n^{\hat{\gamma}}(X_n - p)$ converges almost surely, but we cannot identify the limiting distribution.

We also discuss the application of these results to the urn model described at the end of Section 2.3, with replacement matrix given by (2.4), page 12.

3. Models of percolation

3.1 Percolation

Before saying something about *first-passage percolation*, let us very briefly say something about *percolation*. Percolation is a process related to a graph, as an example we use the “infinite grid” graph \mathbb{Z}^2 (i.e. the graph with vertex set \mathbb{Z}^2 and with edges joining vertices at distance 1) part of which is depicted in Figure 3.1(a). Given some $p \in [0, 1]$, we go through all the edges, one at the time, and we let it remain with probability p (and thus remove it with probability $1 - p$) independently of all other edges. In this way we are left with a subgraph of the original. The process is depicted in Figure 3.1.

We may think of this subgraph as the structure of some porous medium, through which some liquid – coffee is suggested by the name of the process – might flow along the remaining edges.




Figure 3.1: Percolation on the graph \mathbb{Z}^2 at three different values of p .

A natural question to ask is whether an infinitely large component exists in the resulting subgraph, or rather; what is the probability thereof. This is essentially the probability that a liquid can “percolate” through the medium. The answer is obviously *no* when $p = 0$, since there are no edges at all. The answer is obviously *yes* when $p = 1$, as in 3.1(a), since the entire graph remains. It is intuitive that the probability that there exists an infinite component is increasing in the parameter p . Also, this is a tail event whose probability necessarily equals 0 or 1. Hence, there must be a critical value p_c (which is $1/2$ for this graph) such that an infinite component exists when $p > p_c$ and does not when $p < p_c$. This radically different behavior occurring below and above p_c is said to constitute a *phase transition*. A reference to percolation is [Gri99].

3.2 First-passage percolation

First-passage percolation is also a process related to a graph, e.g. \mathbb{Z}^2 considered in the previous section. This process does not remove any edges but considers each to take some random time to traverse. We can keep the idea of a liquid flowing along the edges of the graph; think of an inflow at the origin $(0,0)$. We start the clock when this flow begins. The liquid proceeds to spread out to adjacent vertices. At each point of time we keep track of wet vertices. In Figure 3.2 we depict this for three different time points.




Figure 3.2: Wet nodes marked as black. The first-passage percolation process on \mathbb{Z}^2 observed at three distinct time points.

The first time liquid reaches a vertex v is said to be the *first-passage time* of v . A natural question to ask is how fast the liquid is moving through the graph. We can measure this e.g. by calculating the time, T_n , to reach the vertex $(n,0)$ as a function of n . Typically, one assumes that the times associated to edges are independent and identically distributed. Despite this, it is, in general, extremely difficult to calculate exact properties of T_n . This is due to the fact that contributions come from every possible path between the origin and $(n,0)$, and there are infinitely many such paths with a complex dependence structure.

The quantity that we shall be most interested in is, essentially, the limit describing the long-term average “speed” of the process,

$$\lim_{n \rightarrow \infty} \frac{n}{T_n},$$

but on a different graph.

The analysis of first-passage percolation is intertwined with that of *subadditive processes*. A general reference to this subject is [SW78].

3.3 A summary of Paper III:

First-passage percolation with exponential times on a ladder

We consider first-passage percolation on *the ladder*, depicted in Figure 3.3, i.e. the graph with vertex set $\mathbb{N} \times \{0, 1\}$ with edges between any pair of vertices at distance 1 from each other.

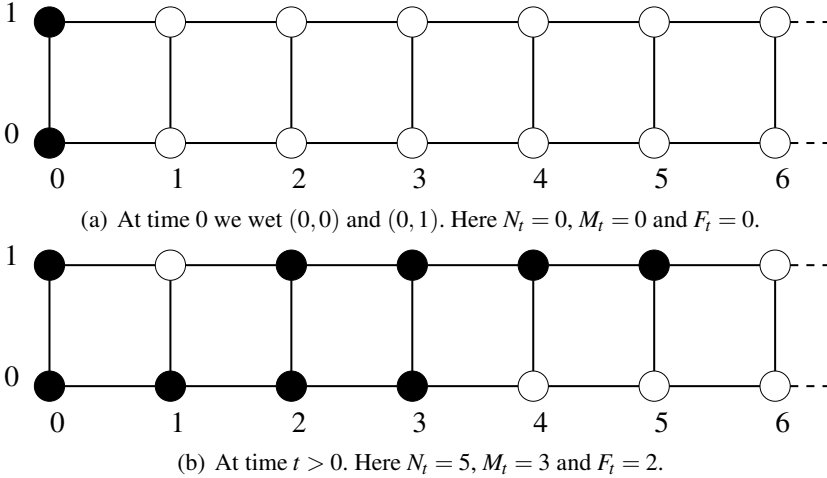


Figure 3.3: The ladder. Wet nodes indicated as black.

Each edge is associated with a random time that has the exponential distribution with mean 1 and times associated with different edges are independent. At time zero we start the process with an inflow at both vertices at “height” zero, i.e. $(0,0)$ and $(0,1)$, as in 3.3(a). The liquid will start to spread to adjacent vertices, and the time it takes to spread along a specific edge is determined by the random time that is associated with that particular edge. At some time $t > 0$ it might look like the situation in 3.3(b). We introduce some notation:

- N_t denotes the x -coordinate of the highest wet node at time t .
- M_t denotes the highest x -coordinate such that both $(x,0)$ and $(x,1)$ are wet at time t .
- F_t denotes the difference $N_t - M_t$.

These quantities are exemplified in Figure 3.3.

Our main concern is if N_t/t will stabilize at some value as t tends to infinity. If so, this may be regarded as the (asymptotic) *speed* \mathcal{V} of the process. The fact that this number will stabilize actually follows from a more general theory of subadditive processes. We can calculate this value \mathcal{V} in a roundabout way, namely via the seemingly uninformative process F_t . The processes N_t and M_t are increasing (to ∞), but since $N_t \geq M_t$ for all t , the process F_t is jumping back and forth between its possible states $\{0, 1, 2, \dots\}$. Maybe, as t increases,

F_t will stabilize over these numbers, i.e. if we were to calculate $\mathbb{P}(F_t = k)$, it might converge, as $t \rightarrow \infty$, to a number π_k , for any $k \geq 0$. If so, we call $\Pi = (\pi_0, \pi_1, \dots)$ the *stationary distribution* and it determines the average behavior of F_t ; averaged over all time, F_t will be in state k the proportion π_k of the time.

Now, think of how F_t relates to N_t . If $F_t = 0$, as in 3.3(a), then there are two edges along which liquid may flow with the effect of an increase in the N -process. If F_t is in any other state $\{1, 2, 3, \dots\}$, there is only 1 edge along which the liquid may flow which results in an increase in the N -process. Thus, by knowing the average behavior of the F -process, we can calculate the average speed to $\mathcal{V} = 2\pi_0 + 1(1 - \pi_0) = 1 + \pi_0$, which is what N_t/t will converge to.

Now, aiming at an exact formula for the speed \mathcal{V} , we had no choice in the distribution of the times associated to edges. These have to be exponential (not with mean 1, but this is just a matter of scaling) in order for the stationary distribution $\Pi = (\pi_0, \pi_1, \pi_2, \dots)$ to be calculable. Chosen as they are, F_t becomes a continuous time Markov chain, which is a well studied process. Such a process has an *intensity matrix* Q , that governs the probabilities of how long the process spends in each state and to where it jumps. This matrix is readily available from the description of the problem.

The problem is how to calculate the distribution Π . From Markov theory we know that such a distribution must satisfy the system of equations $\Pi Q = 0$ (a shorthand notation for what is in fact an infinite set of equations). These equations yield a certain relationship between any π_n and π_0 , namely;

$$\pi_n = a_n \pi_0 - b_n,$$

where a_n and b_n are (integer) sequences that satisfies the recursion

$$c_n = (n+3)c_{n-1} - (n+1)c_{n-2} + c_{n-3}, \quad n \geq 4, \quad (3.1)$$

and the initial sequences a_1, a_2, a_3 and b_1, b_2, b_3 are calculable.

We discover that a transformation of the sequences a_n and b_n is related to the *Bessel functions*, from which an exact expression for every π_n , and thus \mathcal{V} , can be derived. It turns out that

$$\mathcal{V} = 1 + \frac{J_0(2)}{2J_3(2) + J_0(2)} \approx 1.46,$$

where $J_\nu(x)$ is the Bessel function of the first kind,

$$J_\nu(x) = \sum_{k=0}^{\infty} \frac{(-1)^k}{k!(\nu+k)!} \left(\frac{x}{2}\right)^{\nu+2k}. \quad (3.2)$$

The speed of ≈ 1.46 should be compared with the fact that the average speed across an edge is 1.

3.4 A summary of Paper IV:

First-passage percolation on ladder-like graphs with heterogeneous exponential times

This paper generalizes the results of Paper III in two ways.

First, we allow horizontal and vertical edges to have different speeds or *intensities* (the name of the parameter that governs the exponential distribution). We do this by letting the horizontal intensity (more exact: the common intensity associated with the exponential distribution associated with the horizontal edges) be the unit by which we measure the vertical intensity, which we denote λ – see Figure 3.4(a).

Second, we allow diagonals in the graph, see Figure 3.4(b). In this case, the diagonal intensities must be the same as the horizontal, else the method breaks down, but the vertical can still be arbitrary. Both generalizations make use of the stationary distribution of the F -process discussed in the summary of Paper III.




Figure 3.4: Ladder-like graphs where the vertical edges have an arbitrary intensity, denoted λ (as measured in units of the intensities associated to the other edges).

The former case can be solved almost exactly in the same way as in Paper III. The connection to Bessel functions is still true, although the computations are more involved. The latter case can not be solved so easily. Instead we employ results from [Jan10], which deal with general recursions that cover those of the form (3.1). In both cases exact expressions are derived for the speed as a function of λ .

In the first case, Figure 3.4(a), the speed is given by

$$\nu_1(\lambda) = 1 + \frac{(2\lambda^2 + 4\lambda + 1)J_{2+2/\lambda}(2/\lambda) - (\lambda + 1)J_{3+2/\lambda}(2/\lambda)}{(2\lambda^2 + 8\lambda + 5)J_{2+2/\lambda}(2/\lambda) - (\lambda + 3)J_{3+2/\lambda}(2/\lambda)},$$

where the Bessel function $J_\nu(x)$ is defined in (3.2).

For the second case, Figure 3.4(b), the speed is given by

$$\mathcal{V}_2(\lambda) = 2 + 2 \cdot \frac{\frac{(2\alpha-1)e^{-2\alpha}}{\Gamma(\hat{\gamma}+1)} + \frac{1}{2}(\sqrt{2\alpha})^{1-\hat{\gamma}}J_{1+\hat{\gamma}} + \mathcal{S}_\lambda}{\frac{(2\alpha-1)e^{-2\alpha}}{\Gamma(\hat{\gamma}+1)} + \frac{1}{2}(\sqrt{2\alpha})^{1-\hat{\gamma}}[J_{1+\hat{\gamma}} + 2\sqrt{2}J_{2+\hat{\gamma}}] + \mathcal{S}_\lambda},$$

where

$$\alpha = \frac{1}{2+\lambda} \quad \text{and} \quad \hat{\gamma} = -\frac{\lambda}{2+\lambda},$$

the argument $2\sqrt{2\alpha}$ has been suppressed from the Bessel function (i.e. $J_\nu = J_\nu(2\sqrt{2\alpha})$) and

$$\mathcal{S}_\lambda = \sum_{j=0}^{\infty} (-2\alpha)^j (j+1+\alpha) \sum_{m=0}^{\infty} \frac{(-1)^m (2\alpha^2)^m}{\Gamma(m+1+\hat{\gamma})\Gamma(m+1+j)}.$$

4. Sammanfattning

Denna avhandling berör två tämligen disparata delar inom den del av sannolikhetsteorin som benämns stokastiska processer, nämligen *stokastiska approximationsalgoritmer* och *förstapassage-perkolation*.

Stokastiska approximationsalgoritmer uppfanns 1951 i [RM51] som en iterativ procedur för att finna nollstället till en funktion m som man bara kan göra "brusiga" mätningar av. Man tänker sig att $m(x)$ är responsen från ett system då x ges som input. Dock, närhelst man mäter $m(x)$ så observerar man istället värdet

$$M(x) = m(x) + \text{"brus"}$$

där "brus" är en störning som i genomsnitt är noll men typiskt aldrig noll i en enskild mätning. Vi söker nu finna x^* så att $m(x^*) = 0$ (vill vi i själva verket lösa $m(x^*) = \theta$ så kan vi lätt transformera detta problem genom att betrakta $\hat{m}(x) = m(x) - \theta$ istället för m i problemformuleringen). Den i [RM51] föreslagna metoden var att – om m kan antas vara en ökande funktion – starta med någon godtycklig gissning X_0 och sedan skapa en följd "förbättringar" genom att låta påföljande värde skapas ur det föregående via

$$X_{n+1} = X_n - M(X_n)/n. \quad (4.1)$$

Då kommer $M(X_n)$ i genomsnitt vara positiv om $X_n > x^*$ och i genomsnitt negativ om $X_n < x^*$. Således kommer X_{n+1} i genomsnitt minska när $X_n > x^*$ och öka när $X_n < x^*$. Under restriktioner på bruset visade författarna till [RM51] att följden X_0, X_1, X_2, \dots konvergerar i sannolikhet mot x^* .

Denna typ av algoritm har fått många tillämpningar inom kontrollteori, systemidentifikation, maskininlärning, m.m. Typiskt för dessa är att man själv skapar algoritmen, dvs. det handlar ofta om att implementera mjukvara att kunna hantera en dataström. I sådan fall finns möjligheten att variera vissa parametrar, t.ex. är den så kallade steglängden "1/n" i (4.1) helt godtyckligt vald, måhända skulle ett annat val ge snabbare konvergens.

Vår ingångspunkt är en annan. Vi är primärt intresserade av en typ av processer som låter sig beskrivas som dylika algoritmer. Det finns en klass av urnmodeller som kallas generaliserade Pólya-urnor som faller under denna kategori. Detta är modeller där man tänker sig en urna med kulor av minst två olika färger. En kula dras och läggs tillbaka i urnan tillsammans med ett antal

extra kulor vars antal och färger beror på den dragna kulans färg. Så utvecklar sig urninnehållet på ett slumpmässigt sätt.

För att få ett endimensionellt problem – vilket är det enda vi studerar – så låter vi urnan ha kulor av blott två färger. Om vi tittar på följderna Z_0, Z_1, Z_2, \dots av *andelar* av en av dessa färger (andelen av den andra är då 1 minus den första) så kommer dessa typiskt att låta sig beskrivas iterativt via

$$Z_{n+1} = Z_n + \gamma_{n+1}(f(X_n) + U_{n+1}), \quad (4.2)$$

där γ_{n+1} är en ”steglängd”, U_{n+1} är ”brus” (avvikelse från genomsnittligt beteende) och f är den s.k. driftfunktionen.

Arbete I handlar om gränsvärdet till (4.2). Finns det och i vilken mån kan man avläsa det från f ? Det visar sig nämligen – under rimliga villkor – att gränsvärdet för endimensionella algoritmer alltid finns och är ett nollställe till f . Finns det då ett unikt nollställe så är saken avgjord. Finns det flera så kan man börja med att skilja mellan *stabila* och *instabila* nollställen. Ett nollställe är stabilt om man – närhelst man befinner sig tillräckligt nära – i genomsnitt rör sig *mot* detta. Det finns en positiv sannolikhet (under ytterligare villkor) att algoritmen fastnar i ett sådant nollställe. Ett nollställe är instabilt om man – närhelst man befinner sig tillräckligt nära – i genomsnitt *inte* rör sig mot detta. Sannolikheten är noll (under ytterligare villkor) att algoritmen fastnar i ett sådant nollställe. Arbete I diskuterar också en annan typ av nollställen som kallas *tangeringspunkter* som också kan ingå i gränsfördelningen.

Arbete II fortsätter att studera samma typ av algoritmer. Om vi vet att processen konvergerar till en punkt p så är det, inom sannolikhetsteorin, naturligt att fråga sig om man kan ”förstora” följderna $Z_n - p$ (som i sig bara går mot noll) så att den konvergerar mot en sannolikhetsfördelning. Detta bestäms till stor del av gränsvärdet

$$\hat{\gamma} = - \lim_{n \rightarrow \infty} \frac{n \cdot \gamma_n \cdot f(X_n)}{X_n - p}$$

(om det existerar) eller, lättare uttryckt, av $\hat{\gamma} = -\gamma \cdot f'(p)$ om f är deriverbar (vid p) och $n\gamma_n \rightarrow \gamma$ (båda dessa villkor är uppfyllda i de tillämpningar vi är intresserade av).

Närhelst den asymptotiska variansen hos bruset inte försvinner så uppvisar algoritmen asymptotisk normalitet; om $\hat{\gamma} > 1/2$ så konvergerar $\sqrt{n}(Z_n - p)$, om $\hat{\gamma} = 1/2$ så konvergerar $\sqrt{n/\ln n}(Z_n - p)$. I båda fallen är konvergensen *i fördelning* mot en normalfördelning vars parametrar går att bestämma (de är olika i de två fallen).

Då $0 < \hat{\gamma} < 1/2$ så kommer $n^{\hat{\gamma}}(Z_n - p)$ att konvergera med sannolikhet 1 men vi kan inte säga något om gränsfördelningen.

Förstapassage-perkolation är en modell för hur en vätska i realtid sprider sig längs med sprickor i ett berg eller gångar i ett poröst material. Från en matematisk synvinkel så kan vi tänka oss följande; låt oss förenklat anta att den underliggande strukturen där vätskan kan färdas är såsom ett (oändligt) rutat papper. Alla skärningspunkter mellan linjer på detta papper kallar vi *noder* och alla små linjesegment mellan noder kallar vi *kanter*. Vi mäter längder i denna struktur i ”kantenheter”, dvs. vi bestämmer att en kant är en enhet lång. Strukturen verkar kanske väl enkel men vi ska tänka oss att varje kant tar en slumpmässig tid att passera. Typiskt så kan alla tider (i princip) vara olika men från *samma fördelning* samt *oberoende*.

Nu tänker vi oss ett inflöde av vätska i någon punkt. Eftersom alla punkter i strukturen är ekvivalenta kan vi utse denna punkt till origo i ett kartesiskt koordinatsystem. Från denna nod kommer vätskan sprida sig längs kanterna till närliggande noder. Första gången vätskan når en viss nod kallas *förstapassagetiden* för denna nod. Tänk på noden $(1, 0)$ som ligger direkt till höger om origo $(0, 0)$. Det är *inte* så enkelt att denna nods förstapassagetid alltid är lika med tiden det tar för vätskan att sprida sig via kanten som förbinder dessa noder. Måhända gjorde slumpen att detta var en ovanligt tidskrävande passage och att vätska dessförinnan hunnit rinna först *upp* till $(0, 1)$ sedan *höger* till $(1, 1)$ och sedan *nedåt* till $(1, 0)$ (eller varför inte någon annan av de oändligt många vägarna som finns att tillgå?). På detta sätt är varje nods förstapassagetid den i tid snabbaste vägen av *samtliga* tänkbara vägar från origo.

Den primära frågan vi är intresserade av är *hur snabbt* vätskan flödar genom grafen. Det finns förvisso flera sätt att mäta detta. Ett sätt är att försöka räkna ut förstapassagetiden T_n till punkten $(n, 0)$ som en funktion av n . Om man sedan, som vi, är intresserade av den asymptotiska hastigheten så kan man titta på huruvida n/T_n tycks konvergera mot något visst värde.

Arbete III undersöker hastigheten på en delgraf till det oändligt stora rutade pappret, nämligen *stegen* – denna stegen (liggande) syns i Figur 3.3(a) på sidan 19 – då tiden för vätskan att passera en enskild kant kommer från en särskild fördelning som kallas exponentialfördelningen. På denna graf och med dessa tider finns det nämligen ett sätt att exakt räkna ut den asymptotiska hastigheten (essentiellt såsom den definierades i föregående stycke). Vi kan inte räkna direkt på T_n , men det visar sig väl värt att studera avståndet mellan den högsta våta noden och den högsta höjd på vilken man har ett par våta noder bredvid varandra. Detta avstånd fluktuerar mellan sina tillstånd $\{0, 1, 2, \dots\}$ på ett slumpmässigt sätt *men* så att det stabiliserar sig över tiden. Denna stabilisering blir, i asymptotiskt avseende, det genomsnittliga beteendet för detta avstånd som i sin tur låter oss beräkna den i samma avseende genomsnittliga hastigheten *exakt*. Det visar sig att om hastigheten över en enskild kant är 1 så är perkolationshastigheten lika med ungefär 1.46.

Arbete IV fortsätter undersökningen av stegar (och andra steg-liknande grafer) då kanthastigheten tillåts bero på huruvida kanten är horisontell eller vertikal.

5. Acknowledgements

I am grateful to both my supervisors for their support and advice. However a reader of this thesis – if such exists – judges the exposition of ideas and clarity of proofs herein, these are immensely better as compared to my original manuscripts, thanks to the patient guidance of Sven Erick Alm and Svante Janson. My sincerest thanks to you both.

Thanks also to my friends whose time at the department of mathematics has overlapped my five years here. In particular;

- ~ Cecilia Holmgren, with whom I shared an office for the first couple of years, it was a lot of fun (especially learning all about your hometown ... Gothenburg, was it?),
- ~ Mikael Foghelin, whose integral estimations in connection to large deviations were always adequate (yet could be perfected upon, by me – it's a fact), *and*
- ~ Niclas Peterson, whose rendition of *ABC* is a treasured memory (it seems that what happens in Ödängla does not necessarily stay in Ödängla).

A special thanks is due to Persi Diaconis and Brian Skyrms for getting me interested in urn models back in 2006 and to Ann McLeod who so kindly allowed me to use her photo on the cover of this thesis.

On several occasions I have received financial aid for traveling via The Royal Swedish Academy of Sciences and *Stiftelsen G. S. Magnusons fond*, for which I am most grateful.

Finally, Emma and Ida; you are amazing and wonderful and I feel like the luckiest person in the world to have you as my family. Also, Edvin – if that is your name – we are waiting... is it not time to come out soon?

Bibliography

- [Ben99] M. Benaïm: *Dynamics of stochastic approximation algorithms*. Séminaire de Probabilités. Lectures Notes in Mathematics, **1709**, Springer (1999), 1–68.
- [Bor08] V. Borkar: *Stochastic approximation. A dynamical systems viewpoint*. Cambridge University Press (2008).
- [EP23] F. Eggenberger, G. Pólya: Über die Statistik verketteter Vorgänge. *Zeit. Angew. Math. Mech.*, **3** (1923), 279–289.
- [Fab68] V. Fabian: On asymptotic normality in stochastic approximation. *Ann. Math. Statist.*, **39** (1968), 1327–1332.
- [Gri99] G. Grimmett: *Percolation*. Second edition. Springer-Verlag, Berlin (1999).
- [Jan10] S. Janson: A divergent generating function that can be summed and analysed analytically. *Disc. Math. Theor. Comput. Sci.*, **12** No. 2 (2010) 1–22.
- [JK77] N. Johnson, S. Kotz: *Urn models and their application. An approach to modern discrete probability theory*. Wiley (1977).
- [KW52] J. Kiefer, J. Wolfowitz: Stochastic estimation of the maximum of a regression function. *Ann. Math. Statist.*, **23** (1952) 462–466.
- [LPT04] D. Lamberton, G. Pagès, P. Tarrès: When can the two-armed bandit algorithm be trusted? *Ann. Appl. Probab.*, **14** No. 3 (2004) 1424–1454.
- [LP08] D. Lamberton, G. Pagès: How fast is the bandit? *Stoch. Anal. Appl.*, **26** No. 3 (2008) 603–623.
- [MP73] P. Major, P. Révész: A limit theorem for the Robbins-Monro approximation. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, **27** (1973) 79–86.
- [Pag05] G. Pagès: A two armed bandit type problem revisited. *ESAIM Probab. Stat.*, **9** (2005) 277–282.
- [Pem88] R. Pemantle: *Random processes with reinforcement*. Doctoral Dissertation. M.I.T. (1988).
- [Pem91] R. Pemantle. When are touchpoints limits for generalized Pólya urns? *Proc. Amer. Math. Soc.*, **113** No. 1 (1991), 235–243.
- [Pem07] R. Pemantle: A survey of random processes with reinforcement. *Probab. Surv.*, **4** (2007), 1–79.

- [RM51] H. Robbins, S. Monro: A stochastic approximation method. *Ann. Math. Statist.*, **22** (1951) 400–407.
- [SW78] R. T. Smythe, J. C. Wierman: *First-Passage Percolation on the square lattice*. Lecture notes in mathematics **671**, Springer (1978).
- [Tar01] P. Tarrès: *Algorithmes stochastiques et marches aleatoires renforcées*. Thèse de l'ENS Cachan (2001).