

Random Walks on Colored Graphs: Analysis and Applications

Diane Hernek

TR-95-045

August 1995

Abstract

This thesis introduces a model of a random walk on a colored undirected graph. Such a graph has a single vertex set and k distinct sets of edges, each of which has a color. A particle begins at a designated starting vertex and an infinite color sequence C is specified. At time t the particle traverses an edge chosen uniformly at random from those edges of color C_t incident to the current vertex.

The first part of this thesis addresses the extent to which an adversary, by choosing the color sequence, can affect the behavior of the random walk. In particular, we consider graphs that are covered with probability one on all infinite sequences, and study their expected cover time in the worst case over all color sequences and starting vertices. We prove tight doubly exponential upper and lower bounds for graphs with three or more colors, and exponential bounds for the special case of two-colored graphs. We obtain stronger bounds in several interesting special cases, including random and repeated sequences. These examples have applications to understanding how the entries of the stationary distributions of ergodic Markov chains scale under various elementary operations.

The random walks we consider are closely related to space-bounded complexity classes and a type of interactive proof system. The second part of the thesis investigates these relationships and uses them to obtain complexity results for reachability problems in colored graphs. In particular, we show that the problem of deciding whether a given colored graph is covered with probability one on all infinite sequences is complete for natural space-bounded complexity classes.

We also use our techniques to obtain complexity results for problems from the theory of nonhomogeneous Markov chains. We consider the problem of deciding, given a finite set $\mathcal{C} = \{C_1, \dots, C_A\}$ of $n \times n$ stochastic matrices, whether every infinite sequence over \mathcal{C} forms an ergodic Markov chain, and prove that it is PSPACE-complete. We also show that to decide whether a given finite-state channel is indecomposable is PSPACE-complete. This question is of interest in information theory where indecomposability is a necessary and sufficient condition for Shannon's theorem.

This work was supported in part by a Lockheed graduate fellowship and NSF grant CCR92-01092.

Contents

1	Introduction	1
1.1	Notation and Terminology	3
1.2	Markov Chain Background	4
1.2.1	Homogeneous Markov Chains	4
1.2.2	Nonhomogeneous Markov Chains	5
2	Cover Time	7
2.1	Introduction	7
2.2	Upper Bounds	8
2.3	Lower Bounds	12
2.4	Concluding Remarks	14
3	Special Cases and Applications	15
3.1	Introduction	15
3.2	Special Graphs	16
3.2.1	Proportional Colored Graphs	16
3.2.2	Graphs with Self-Loops	17
3.3	Special Sequences	17
3.3.1	Random Sequences	17
3.3.2	Repeated Sequences	17
3.3.3	Corresponding Homogeneous Markov Chains	18
3.4	Lower Bounds	19
3.5	An Application to Products and Weighted Averages	21
4	Colored Graphs and Complexity Classes	22
4.1	Introduction	22
4.2	One-way Interactive Proof Systems	23
4.2.1	Example: Coin Flipping Protocol	23
4.3	Two-colored Directed Graphs	24
4.3.1	Example: Coin Flipping Protocol Revisited	25
4.4	Polynomial Space	25
4.5	Colored Graph Connectivity	26
4.5.1	Space-bounded Algorithms	27

4.5.2	Hardness Results	29
5	Applications	32
5.1	Introduction	32
5.2	Information Theory	33
5.2.1	Preliminaries	33
5.2.2	Noisy Communication and the Finite-State Channel	35
5.3	Complexity Results	36
5.4	Concluding Remarks	38
	Bibliography	40

Acknowledgements

There are many people to thank for the role they played during my graduate school years.

First there is my advisor, Manuel Blum, whose enthusiasm and encouragement gave me the confidence to develop my independence and a sense of research taste and style. Alistair Sinclair also deserves special mention. I have relied heavily on his insight and advice. In addition to being a second advisor, Alistair is also a good friend. I would like to thank Yuval Peres for his suggestions which greatly helped to improve the clarity of this thesis. The work in this thesis was done jointly with Anne Condon at the University of Wisconsin. I have learned a great deal working with Anne and have enjoyed it tremendously. Thanks to Dick Karp and Umesh Vazirani for their excellent teaching and for useful discussions.

Berkeley has a wonderful group of graduate students and researchers and I have made some of my dearest friends here. Over long distances my friendships with Sandy Irani and Ronitt Rubinfeld have only grown stronger. To me they are like family. Graduate school would not have been the same without Dana Randall. I continue to be amazed by her generosity and her ability to read my mind. Mike Luby has also been very special and I thank him for his friendship and advice. I have learned and laughed a lot in many long conversations with Amie Wilkinson. Some of the best laughs I have ever had were shared with Nina Amenta and Will Evans; I have appreciated their warmth and humor. I have greatly enjoyed time spent with Madhu Sudan, Francesca Barrientos, Sara Robinson, Mike Mitzenmacher, Z Sweedyk, Deborah Weisser, Mike Schiff, Ramon Caceres and Dan Jurafsky.

Finally, I would like to thank my mother, Joan Moderes, for her love and support.

Chapter 1

Introduction

A k -colored graph G is a $k + 1$ -tuple (V, E_1, \dots, E_k) , where V is a finite set of vertices and each $E_i \subseteq V \times V$ is a set of edges. We will refer to the set E_i as the edges of color i . If, for all i , whenever (u, v) is in E_i (v, u) is also in E_i , then G is a k -colored *undirected* graph. In this case we will write $\{u, v\}$ to represent the undirected edge that connects vertices u and v . Otherwise, G is a k -colored *directed* graph. Unless otherwise specified the graphs considered in this thesis will be undirected. As we will see, undirected colored graphs are as general as their directed counterparts.

This thesis introduces a model of a random walk on a colored undirected graph. A random walk on a colored graph proceeds as follows. A particle begins at a designated starting vertex and an infinite color sequence C over the alphabet $\{1, \dots, k\}$ is specified. At time t the particle traverses an edge chosen uniformly at random from those edges of color C_t incident to the current vertex. The case of $k = 1$ corresponds to a simple random walk on an undirected graph.

This thesis investigates intrinsic properties of random walks on colored graphs, such as expected cover time, as well as applications in computational complexity, where there are direct applications to the theory of nonhomogeneous Markov chains and coding and information theory. Many of the results have appeared in the papers [9] and [8].

We begin in Chapter 2 with an investigation of the expected cover time of random walks on colored graphs. The cover time of the colored graph G is the number of steps until a random walk visits all of the vertices of G , as a worst case over all starting vertices and infinite color sequences. We consider only those graphs that are covered with probability one on all infinite sequences from all start vertices, since without this property there is no bound on the cover time. We show that the expected cover time of colored graphs with two colors is exponential in the number of vertices, and that graphs with three or more colors have doubly exponential expected cover time. Since it

is well-known that connected undirected graphs (the case of one color) have polynomial expected cover time, these results establish a three-level hierarchy of cover times in colored graphs.

In Chapter 3 we go on to prove tighter bounds on the expected cover time in a variety of interesting special cases. These cases are of two types: we consider both special classes of colored graphs and special types of color sequences. We show that if a colored graph is *proportional* then its expected cover time is polynomial. The proportionality property simply says that a random walk on each of the underlying graphs (V, E_i) is an ergodic Markov chain, and that, in addition, the Markov chains for random walks on all of the (V, E_i) share the same stationary distribution.

We also consider the case where each underlying graph (V, E_i) is connected and has a self-loop at every vertex; that is, $(j, j) \in E_i$ for all j . In this case, a random walk on (V, E_i) is again an ergodic Markov chain; however, the stationary distributions of the Markov chains corresponding to each of the (V, E_i) may differ. In this case, we give tight exponential upper and lower bounds on the expected cover time. Hence, when the stationary distributions of the underlying graphs coincide the expected cover time is polynomial, but when the stationary distributions differ the expected cover time is exponential.

Finally, we consider the behavior of random walks on colored graphs when the color sequence is chosen at random and when the color sequence consists of a finite sequence $C_1 \dots C_l$ repeated ad infinitum. In both of these cases the random walk corresponds to a homogeneous Markov chain, and we can show that the expected cover time is at most exponential. In the case that the corresponding homogeneous Markov chain is ergodic and all of the entries of its stationary distribution are inversely polynomial, the expected cover time is polynomial. We give an example of a colored graph for which the homogeneous Markov chains defined by random and repeated sequences is ergodic, but the expected cover time is still exponential. Hence, we prove tight exponential upper and lower bounds on random and repeated sequences. Moreover, the example shows that it is possible for an ergodic Markov chain that is composed of an average or product of random walks on connected undirected graphs to have exponentially small entries in its stationary distribution, even though the entries of the stationary distributions for the original random walks are only inversely polynomial.

Two-colored *directed* graphs were first studied by Condon and Lipton [10] in their investigation of *one-way interactive proof systems with space-bounded verifiers*. In an interactive proof system a *prover* P wishes to convince a *verifier* V that a given shared input string x is a member of some language L . The prover and the verifier share independent read-only access to the input string x . The verifier V also has a private read-write worktape and the ability to toss coins during its

computation. In our case, we are interested in verifiers V that are *space-bounded*; that is, verifiers that write on at most $s(n)$ tape squares on all inputs of length n . In particular, we will be interested in systems where the verifier uses space $O(\log n)$ on all inputs of length n .

In a general system, the computation proceeds in rounds. In each round, the verifier tosses a coin and asks a question of the more powerful prover. Based on the answers of the prover, the computation continues until eventually the verifier decides to accept or reject x and halts by entering an accepting or rejecting state. The systems we consider are *one-way* in the sense that all communication goes *from* the prover *to* the verifier. Since the system is one-way we can think of the prover as being represented by a *proof string* and the verifier as having one-way read-only access to the proof. We say that a language L has a one-way interactive proof system with a logspace verifier if there exists a probabilistic Turing machine V that on all inputs x of length n uses space $O(\log n)$ and satisfies the following one-sided error conditions:

1. If x is in L , then there is some finite proof string that causes V to accept with probability 1.
2. If x is not in L , then on any finite or infinite proof V rejects with probability at least $2/3$.

In Chapter 4 we further the study of one-way interactive proof systems with logspace verifiers by showing that every language in PSPACE, the class of languages recognized by polynomial space-bounded Turing machines, has a one-way interactive proof system with a logspace verifier. In [10] the authors show that the question of whether a logspace verifier V accepts or rejects its input corresponds to a reachability question in an appropriately defined two-colored *directed* graph. We use this correspondence in conjunction with the PSPACE result to prove PSPACE-completeness results for connectivity problems for colored graphs. In particular, we show that the problem of deciding, given a colored graph G with three or more colors, whether G is covered with probability one on all infinite sequences is PSPACE-complete. We also show that the analogous problem for two-colored graphs is complete for nondeterministic logspace.

As was noted earlier, the random walks of this thesis correspond to nonhomogeneous Markov chains. In a nonhomogeneous Markov chain the probability transition matrix can change in each time step. Natural complexity-theoretic questions arise when we think of the matrices that define the Markov chain as being drawn from a finite set $\mathcal{C} = \{C_1, \dots, C_A\}$ of $n \times n$ stochastic matrices. In Chapter 5 we use the machinery of colored graphs to prove PSPACE-completeness of several problems from the study of nonhomogeneous Markov chains. Every infinite product $\prod_{j=1}^{\infty} C_{i_j}$ over the set \mathcal{C} defines a finite nonhomogeneous Markov chain. We show that the problem

of deciding whether every infinite product over \mathcal{C} defines an *ergodic* Markov chain is PSPACE-complete. We also show that the related problem of deciding whether all finite words over \mathcal{C} are indecomposable is PSPACE-complete. This question has applications to coding and information of finite-state channels. In particular, it is a necessary and sufficient condition for Shannon's coding theorem for finite-state channels. Hence, we show that to decide whether a given finite-state channel has an optimal code is PSPACE-complete.

The application to Shannon's theorem for finite-state channels lead to a series of papers [25] [26] [21] investigating the complexity of deciding whether all words over a given set \mathcal{C} are indecomposable. This work resulted in several finite decision procedures, all of which are easily seen to be in PSPACE and EXPTIME (deterministic time 2^{n^c} for some constant c). Our PSPACE-completeness result gives strong evidence that the currently known algorithms are the best possible. They show that a subexponential time algorithm would imply a separation of PSPACE from EXPTIME, which would be a major breakthrough in complexity theory.

The remainder of this chapter is a brief description of the notation and terminology that will be used in this thesis, as well as a review of the necessary Markov chain background.

1.1 Notation and Terminology

Let $G = (V, E_1, \dots, E_k)$ be a k -colored undirected graph with n vertices. We will refer to the undirected graph (V, E_i) as the *underlying graph colored i* . For each color i and vertex v , the degree $d_i(v)$ is $|\{w : \{v, w\} \in E_i\}|$. For each color i , we will use A_i to denote the $n \times n$ adjacency matrix for the edge set E_i . The $n \times n$ stochastic matrix P_i is the probability transition matrix for a simple random walk on (V, E_i) , and is given by:

$$P_i(u, v) = \begin{cases} \frac{1}{d_i(u)}, & \text{if } (u, v) \in E_i; \\ 0, & \text{otherwise.} \end{cases}$$

Let $C = C_1 C_2 C_3 \dots$ be an infinite color sequence over the alphabet $\{1, \dots, k\}$ and let $s \in V$ be a vertex in G . A random walk starting from s on the color sequence C proceeds as follows. The walk begins at time 0 at the vertex s . Suppose that at time $t \geq 0$ the walk is at vertex u . Then, for all vertices v , at time $t + 1$ the walk moves to vertex v with probability $P_{C_t}(u, v)$.

Let $C_1 \dots C_l$ be a finite color sequence. We use $(C_1 \dots C_l)^\omega$ to denote the infinite sequence obtained by repeating $C_1 \dots C_l$ ad infinitum.

1.2 Markov Chain Background

In this section we review the Markov chain terminology and background that will be used in the chapters that follow.

1.2.1 Homogeneous Markov Chains

An $n \times n$ stochastic matrix M defines a *homogeneous Markov chain* \mathcal{M} whose state space is the set $[n] = \{1, \dots, n\}$, and for which the probability of going from state i to state j in one step is given by $M(i, j)$.

The Markov chain \mathcal{M} is said to be *ergodic* if the limit $\lim_{t \rightarrow \infty} M^t$ exists and has all rows equal. An equivalent condition for ergodicity is that the probability transition matrix M is both *indecomposable* and *aperiodic*.

In order to define indecomposable and aperiodic, consider the directed graph G induced by the nonzero entries of M . That is, consider the directed graph $G = ([n], E)$ with vertex set $[n] = \{1, \dots, n\}$ and edge set $E = \{(i, j) : M(i, j) > 0\}$. Let $G' = (V', E')$ be the directed graph whose vertices correspond to the strongly connected components of G . There is a directed edge (C, C') from component C to component C' if and only if there exists an $i \in C$ and a $j \in C'$ such that $(i, j) \in E$. The graph G' is called the *component graph* of G and is necessarily acyclic.

The matrix M is *indecomposable* if the component graph G' contains exactly one vertex that is a sink; that is, there is exactly one vertex with no non-loop edges leaving it. In the terminology of nonnegative matrices, each vertex in the component graph corresponds to a *communicating class* of indices of M . Sink vertices correspond to *essential classes*. Other vertices are *inessential classes*. The stochastic matrix M is indecomposable if it contains exactly one essential class of indices. For examples, see Figure 1.1 below. In the first example, $\{v_1, v_3\}$ is an inessential class and $\{v_2\}$ is an essential class, so the chain is indecomposable. In the second example, $\{w_1\}$ is an inessential class and $\{w_2\}, \{w_3\}$ are essential classes, so the chain is decomposable.

The greatest common divisor of the lengths of all cycles in G is called the *period* p of M . The matrix M is *aperiodic* if p is equal to one.

Notice that ergodicity is completely determined by the positions of the non-zero entries in the probability transition matrix M , and is independent of the actual values in those positions. We will define the *type* of M to be the $n \times n$ matrix $\langle M \rangle$ that has a 1 in position (i, j) if $M(i, j) > 0$, and a 0 otherwise. Stochastic matrices M_1 and M_2 are said to be of the same type if $\langle M_1 \rangle = \langle M_2 \rangle$; that is, if they have positive elements and zero elements in the same positions.




Figure 1.1: Example illustrating the definition of indecomposable

An ergodic Markov chain \mathcal{M} has a unique limiting or *stationary distribution* which is the n -dimensional row vector π corresponding to any row of the limit $\lim_{t \rightarrow \infty} M^t$. The vector π satisfies $\pi(i) \geq 0$ for all i , $\sum_i \pi(i) = 1$, and $\pi M = \pi$.

A stronger definition of ergodicity is that the limit $\lim_{t \rightarrow \infty} M^t$ exists, is *positive*, and has all rows equal. An equivalent set of conditions is that the matrix M is *irreducible* and aperiodic. The matrix M is irreducible if the graph G induced by the nonzero entries of M is strongly connected. That is, for every pair of vertices u and v , v is reachable from u and u is reachable from v . In this case M contains one communicating class of indices. Following Seneta [22] we will call such an ergodic Markov chain *regular*. In a regular Markov chain all entries in the stationary distribution are strictly positive.

A random walk on a connected undirected nonbipartite graph $G = (V, E)$ forms a regular Markov chain. It is easy to verify that its unique stationary distribution π is given by $\pi(v) = d(v)/2|E|$, for all $v \in V$.

1.2.2 Nonhomogeneous Markov Chains

A finite *nonhomogeneous Markov chain* \mathcal{M} is defined by an infinite sequence

$$M_1, M_2, M_3, \dots$$

of $n \times n$ stochastic matrices. Once again the state space of the Markov chain is $[n]$ but the transition probabilities can be different at different time steps. The matrix M_i is the probability transition

matrix for the i th time step. A homogeneous Markov chain with probability transition matrix M is the special case M, M, M, \dots

Let $M^{(t,t')}$ denote the product $\prod_{k=t}^{t+t'} M_k$. The nonhomogeneous Markov chain \mathcal{M} is said to be *ergodic* if, for each t , as $t' \rightarrow \infty$:

$$|M^{(t,t)}(i, j) - M^{(t,t')}(i', j)| \rightarrow 0, \text{ for all } i, i', j.$$

That is, \mathcal{M} is ergodic if, for all t , as t' tends to infinity the rows of the matrix $M^{(t,t')}$ tend to equality. If, in addition, for all t , $M^{(t,t')}$ tends to a limit as t' tends to infinity then the Markov chain \mathcal{M} is said to be *strongly ergodic*. Otherwise, \mathcal{M} is said to be *weakly ergodic*.

The following example illustrates the difference between weak and strong ergodicity for nonhomogeneous Markov chains. Consider the matrices A_1 and A_2 whose nonzero entries are represented by the directed graphs shown in Figure 1.2. All infinite products over $\{A_1, A_2\}$ are



Figure 1.2: Example illustrating the difference between weak and strong ergodicity

weakly ergodic since in both of the graphs the next state is independent of the previous state. However, the infinite product $A_1 A_2 A_1 A_2 A_1 \dots$ is not strongly ergodic.

Chapter 2

Cover Time

2.1 Introduction

In this chapter we investigate the expected cover time of colored graphs. We say that a colored graph G can be *covered from* s if, on every infinite sequence C of colors, a random walk on C starting at s visits every vertex with probability one. The *expected cover time* of G is defined to be the supremum, over all infinite sequences C and start vertices s , of the expected time to cover G on C starting at s . Throughout this chapter we only consider those graphs G that can be covered from all start vertices. This property is needed since without it there is no bound on the cover time.

The condition that G be covered from all its vertices makes it necessary for the underlying graphs of each color to be connected. This is because G must be covered with probability one on the sequence $(c)^\omega$ for all colors c . The condition that all of the underlying graphs are connected, however, is not a sufficient condition. For instance, consider the graph of Figure 2.1, where the solid lines are the edges colored R and the dotted lines are the edges colored B . Both of the underlying




Figure 2.1: Underlying graphs connected but *not* covered from all start vertices

graphs are connected; however, a random walk on the sequence $(RB)^\omega$ starting from s does not cover the graph.

The property that G be covered from all of its vertices is a generalization of the connectivity

property for undirected graphs. In this chapter we use the property as stated. In Chapter 4 we return to give an exact combinatorial characterization and to investigate the computational complexity of determining whether or not it is satisfied.

The expected cover time of a simple random walk on an undirected graph (the case of one color) has been well-studied, and various polynomial bounds on the expected cover time have been shown [1] [7]. In what follows we prove the following two main results on the expected cover time of colored graphs with n vertices:

Theorems 2.1 and 2.5 The expected cover time of colored graphs is bounded above by $2^{2^{O(n)}}$, and there are graphs with three colors that achieve this bound.

Theorems 2.2 and 2.3 The expected cover time of two-colored graphs is bounded above by $2^{\text{poly}(n)}$, and there are graphs with two colors that achieve this bound. More precisely, we prove an upper bound of $2^{O(n^2 \log n)}$ and a lower bound of $2^{\Omega(n)}$.

These results combined with known results about the one color case establish a three-level hierarchy of cover times in colored graphs.

2.2 Upper Bounds

Let G be a colored graph and let s and t be two vertices of G . We say that t is *reachable from s on the color sequence $C = C_1 \dots C_l$* , if there is a sequence of vertices $s = v_0, v_1, \dots, v_l = t$ such that G contains an edge of color C_i between v_{i-1} and v_i , for $1 \leq i \leq l$. We call v_0, v_1, \dots, v_l a *path from s to t on C* .

For any pair of vertices s and t , we define the *distance* $\text{dist}(s, t)$ to be the minimum l such that t is reachable from s on a prefix of *every* sequence of length l . Notice that since we assume that G is covered from all start vertices, $\text{dist}(s, t)$ is necessarily finite. The key to proving the upper bounds on the cover time is to obtain good bounds on the maximum distance between vertices in a colored graph.

Lemma 2.1 *Let G be a colored graph with n vertices, and let s and t be vertices in G . If G is covered from all of its vertices, then $\text{dist}(s, t)$ is at most 2^n .*

Proof. Let $C = C_1 \dots C_l$ be any color sequence of length $l = 2^n$. Assume that t is not reachable from s on any prefix of C . Let $S_0 = \{s\}$ and, for $1 \leq i \leq l$, let S_i be the set of vertices

reachable from s on the color sequence $C_1 \dots C_i$. By assumption, t is not in any of the sets S_i , but by the pigeonhole principle $S_j = S_k$ for some $j \neq k$. Hence, on the infinite sequence $C_1 \dots C_j (C_{j+1} \dots C_k)^\omega$, t is never reached from s , which is a contradiction. \square

We are now prepared to prove the following theorem:

Theorem 2.1 *Let G be a colored graph with n vertices that is covered from all vertices. The expected cover time of G is at most $2^{2^{O(n)}}$.*

Proof. Let $C = C_1 C_2 C_3 \dots$ be an infinite color sequence and let s be any vertex in G . Consider an arbitrary ordering $s = 1, \dots, n$ of the vertices of G . We will consider the random walk in intervals of length $l = 2^n$. Suppose that after the first i intervals vertices $1, \dots, t-1$ have been visited but t has not been visited. Let v_i be the current vertex after the first i intervals. Then, since G is covered from all start vertices, by Lemma 2.1, $\text{dist}(v_i, t)$ is at most l . Hence, t is visited in interval $i+1$ with probability at least $1/n^l$. Thus, the expected number of intervals until all vertices are visited is at most $(n-1)n^l$. Since each interval consists of $l = 2^n$ steps, the expected time to cover G is at most $(n-1)2^n n^{2^n} = 2^{2^{O(n)}}$. \square

The result in Theorem 2.1 is independent of the number of colors in G . In the case of graphs with two colors, however, the expected cover time is only singly exponential in n . In what follows we will assume that the two colors are red and blue, and denote them by R and B , respectively. The approach is to strengthen Lemma 2.1 as follows.

Lemma 2.2 *Let G be a two-colored graph with n vertices, and let s and t be vertices in G . If G is covered from all of its vertices, then $\text{dist}(s, t)$ is at most $(4n-3)(n-1)$.*

Once Lemma 2.2 is in place, the proof follows the same general outline as the proof of Theorem 2.1. However, in subsequent chapters we will need a slightly different statement from the one given in Lemma 2.2. Instead we will prove the following equivalent lemma.

Lemma 2.3 *Let G be a two-colored graph with n vertices, and let s and t be vertices in G . If t is reachable from s on a prefix of each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$, then $\text{dist}(s, t)$ is at most $(4n-3)(n-1)$.*

Notice that Lemma 2.2 follows easily from Lemma 2.3, since if a random walk from s visits t with probability one on all infinite sequences then t must be reachable from s on a prefix of each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$.

To prove Lemma 2.3 we will relate arbitrary color sequences to prefixes of the four sequences $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$ using the infinite alternating path P shown in Figure 2.2. Alternate edges of this graph are colored R and B . Thus any sequence of colors defines a unique path from any fixed starting point p on P . For clarity we will refer to the vertices of P as *points* to distinguish them from the *vertices* of G .




Figure 2.2: Alternating path P with fixed starting point p

We say that two finite color sequences C and C' are *similar* if, starting from any point p , the unique point reached on the color sequence C is the same as the unique point reached on C' . For instance, the sequences $RRBRRB$ and $BBRRRRBRR$ are similar. The following lemma is the key to proving Lemma 2.3.

Lemma 2.4 *Suppose that C is similar to C' , where C' is a prefix of $(RB)^\omega$ (or $(BR)^\omega$), and let u and v be vertices of G . If there is a path from u to v on C' , then there is a path from u to v on C .*

Proof. Let q be the unique point on P that is reached from p on sequences C and C' . Since there is a path from u to v on C' , C' defines a path from u to v in G along which the edges are colored the same as the edges from p to q in P . We will construct a path from u to v on C that wanders along this path in the same way that the path from p to q on C wanders along P . Of course the path from p to q on C may visit points that do not lie between p and q . In constructing our path from u to v we need to extend the path in G accordingly.

More precisely, let $C' = C'_1 \dots C'_{m'}$, and let $C = C_1 \dots C_m$. Let $p = p'_0, p'_1, \dots, p'_{m'} = q$ be the path from p to q on C' . Let $p = p_0, p_1, \dots, p_m = q$ be the path from p to q on C . Let $u = u'_0, u'_1, \dots, u'_{m'} = v$ be a path from u to v on C' . We will show how to construct a path $u = u_0, u_1, \dots, u_m = v$ from u to v on C .

The path is defined inductively. We let $u_0 = u$ and, for $j = 1, \dots, m$, define u_j as follows:

$$u_j = \begin{cases} u_i, & \text{if } p_j = p_i, \text{ for some } i < j \\ u'_i, & \text{if } p_j = p'_i, \text{ for some } i \\ w, & \text{otherwise, where } w \text{ is any vertex} \\ & \text{connected to } u_{j-1} \text{ by an edge of color } C_j \end{cases}$$

For example, suppose that $C' = RBR$ and $C = BBRBBBBR$, and that u, u_1, u_2, v is a path from u to v on C' . Then the path we construct on C is: $u, w, u, u_1, u_2, u_1, u_2, v$, where w is a vertex connected to u by an edge colored B . This example is shown in Figure 2.3.




Figure 2.3: Defining a path from u to v on the sequence $C = BBRBBBBR$

It is straightforward to verify that the sequence u_0, u_1, \dots, u_m is indeed a path from u to v on C by checking that $u_0 = u$, $u_m = v$ and, for all $1 \leq j \leq m$, vertices u_{j-1} and u_j are connected by an edge of color C_j .

□

We are now prepared to give the proof of Lemma 2.3.

Proof of Lemma 2.3. We begin by making a few simple observations. Since t is reachable from s on a prefix of $(R)^\omega$, there is a path from s to t on which all edges are colored R . The shortest such path is a simple path and has length at most $n - 1$. Hence, there is a path from s to t on a prefix of $(R)^\omega$ of length at most $n - 1$. Similarly, there is a path from s to t on a prefix of $(B)^\omega$ of length at most $n - 1$.

Since t is reachable from s on a prefix of $(RB)^\omega$, there is a path from s to t that begins with an edge colored R and alternates between R and B . The shortest such path has length at most $2n - 1$, since in a shortest path t appears exactly once and each other vertex appears at most once in an even numbered position and at most once in an odd numbered position. Similarly, there is a path from s to t on a prefix of $(BR)^\omega$ of length at most $2n - 1$.

In what follows we will use these simple observations to prove that, on *any* sequence $C = C_1 \dots C_l$ of length $l = (4n - 3)(n - 1)$, t is reachable from s on a prefix of C . Consider the

unique path from p on P on the sequence $C_1 \dots C_l$. By our choice of $l = (4n - 3)(n - 1)$ it must be the case that either:

1. Some point of P is visited n times on the sequence $C_1 \dots C_l$, or
2. $2n - 1$ distinct points to the right or left of p are visited on the sequence $C_1 \dots C_l$.

In either case we will show that t is reachable from s on $C_1 \dots C_{l'}$, for some $l' \leq l$.

We first consider the case where some point q on P is visited n times on $C_1 \dots C_l$. If q is visited n times on $C_1 \dots C_l$ then at least $2n - 2$ times we traverse one of the two edges incident to q . Hence, either we traverse the edge colored R adjacent to q at least $n - 1$ times, or we traverse the edge colored B adjacent to q at least $n - 1$ times. Without loss of generality assume that the edge colored R is traversed $n - 1$ times. (The argument in the case that the edge colored B is traversed $n - 1$ times is analogous.)

Let $s = v_0, v_1, \dots, v_{m-1}, v_m = t$ be a shortest path from s to t on a prefix of $(R)^\omega$, where $m \leq n - 1$. We will incorporate this path into a walk on C . Since the edge colored R adjacent to q is traversed at least $n - 1$ times, we can rewrite C as follows:

$$C = C^{(0)} R C^{(1)} R C^{(2)} R \dots C^{(m-1)} R C^{(m)},$$

where $C^{(0)}, C^{(1)}, \dots, C^{(m-1)}$ are (possibly empty) strings over $\{R, B\}$ that are similar to the empty string, and $C^{(m)}$ is a string over $\{R, B\}$. For $0 \leq i < m$, $C^{(i)}$ is similar to the empty string so, by Lemma 2.4, for any vertex v in G there is a path from v back to v on $C^{(i)}$. For $0 \leq i < m$, let p_i be a path from v_i back to v_i on $C^{(i)}$. Then $p_0, v_1, p_1, v_2, \dots, p_{m-1}, v_m$ is a path from s to t on $C^{(0)} R C^{(1)} R \dots C^{(m-1)} R$.

Now we consider the case where $2n - 1$ distinct points to the right (or left) of p are visited on the sequence $C_1 \dots C_l$. We do the proof for the case that $2n - 1$ distinct points to the right of p are visited and the edge from p to the point to its right is colored R . We know that on some prefix $C' = C'_1 \dots C'_m$ of $(RB)^\omega$, where $m \leq 2n - 1$, t is reachable from s in G . Let q be the point reachable from p in P on the color sequence $C'_1 \dots C'_m$. Since $2n - 1$ points to the right of p are visited on the sequence $C_1 \dots C_l$, the point q is reached from p on the sequence $C_1 \dots C_{l'}$, for some $l' \leq l$. Thus the sequences $C_1 \dots C_{l'}$ and C' are similar. So, by Lemma 2.4, t is reachable from s on $C_1 \dots C_{l'}$, as required. \square

We can now prove the upper bound on the expected cover time of graphs with two colors using Lemma 2.2 and a proof analogous to that of Theorem 2.1.

Theorem 2.2 *Let G be a two-colored graph with n vertices that is covered from all vertices. The expected cover time of G is at most $2^{O(n^2 \log n)}$.*

Proof. Let $C = C_1 C_2 C_3 \dots$ be an infinite color sequence and let s be any vertex in G . Consider an arbitrary ordering $s = 1, \dots, n$ of the vertices of G . We will consider the random walk in intervals of length $l = (4n - 3)(n - 1)$. Suppose that after the first i intervals vertices $1, \dots, t - 1$ have been visited but t has not been visited. Let v_i be the current vertex after the first i intervals. Then, since G is covered from all of its vertices, by Lemma 2.2, $d(v_i, t) \leq l$ and so t is visited with probability at least $1/n^l$ in the next interval. Thus, the expected number of intervals until all vertices are visited is at most $(n - 1)n^l$. Since each interval consists of $l = (4n - 3)(n - 1)$ steps, the expected time to cover G from s is at most $(n - 1)n^l l = 2^{O(n^2 \log n)}$. \square

Suppose that the colored graph G is not covered from all vertices, but satisfies the weaker condition that it is covered starting from s . It should be noted that the same techniques can be used to bound the expected cover time of a random walk starting from s , as a worst case over all color sequences. It follows from Lemmas 2.1 and 2.3 and the proofs of Theorems 2.1 and 2.2 that if a random walk, after some number of steps, reaches vertex v without visiting t , then $\text{dist}(v, t)$ is at most l , where l is bounded by 2^n in general, and by $(4n - 3)(n - 1)$ in the case of two-colored graphs.

2.3 Lower Bounds

In Theorems 2.3 and 2.5 we prove exponential and doubly exponential lower bounds on the expected cover time of colored graphs with two and three colors, respectively. The lower bounds are based on the following lemma.

Lemma 2.5 *Let G be a k -colored directed graph and let s be a vertex in G . There exists a $(k + 1)$ -colored undirected graph G' and a vertex s' in G' such that:*

1. *the number of vertices in G' is twice the number of vertices in G ,*
2. *G' is covered from all vertices if and only if G is covered from all vertices, and*
3. *for every k -color sequence C , there exists a $(k + 1)$ -color sequence C' such that the expected cover time of G' from s' on C' is at least twice the expected cover time of G from s on C .*

Proof. Let G be a k -colored directed graph with vertices $\{v_1, \dots, v_n\}$ and edge colors $\{1, \dots, k\}$. We will construct a $(k+1)$ -colored undirected graph G' with vertex set $L \cup R$, where $L = \{l_1, \dots, l_n\}$ and $R = \{r_1, \dots, r_n\}$. The graph G' will have an edge colored $k+1$ between l_i and r_i , for all i . There will also be an undirected edge colored c connecting l_j and r_i , for each directed edge (v_i, v_j) colored c in G . In addition, there will be a complete graph on L in each of the colors $1, \dots, k$, and a complete graph on R in the color $k+1$.

This construction is illustrated for an example with $k = 1$ in Figure 2.4 below.




Figure 2.4: Converting a directed graph into a two-colored undirected graph

Now, for every path $p = v_{i_0}, v_{i_1}, \dots, v_{i_m}$ in G on color sequence $C = C_1 C_2 \dots C_m$, there is a corresponding path $p' = l_{i_0} r_{i_0} l_{i_1} r_{i_1} \dots l_{i_m} r_{i_m}$ in G' on color sequence $C' = (k+1)C_1(k+1)C_2 \dots (k+1)C_1(k+1)C_2 \dots (k+1)C_1(k+1)$. Note that, for all j , the path p includes v_j if and only if the corresponding path p' includes l_j and r_j . Moreover, the probability that a random walk on G from v_{i_0} on C takes the path p is exactly the same as the probability that a random walk on G' from l_{i_0} on C' takes the path p' . Since every two steps of the random walk on G' correspond to one step of the random walk on G , the expected cover time of G' on C' from l_j is exactly twice the expected cover time of G on C from v_j . Hence, the expected cover time of G' is at least twice the expected cover time of G .

It remains to show that G' is covered from all its vertices if and only if G is covered from all its vertices.

For the *only if* direction, suppose that there exists a vertex v_i in G and an infinite color sequence C such that G is not covered from v_i on $C = C_1 C_2 C_3 \dots$. Then G' is not covered from l_i on $C' = (k+1)C_1(k+1)C_2(k+1)C_3 \dots$. This is because, for all j , the probability that the walk on G' visits l_j and r_j is exactly the same as the probability that the corresponding walk on G visits v_j .

For the *if* direction, suppose that G is covered from all start vertices. We must show that G' is also covered from all start vertices. First note that, since G is covered from all of its

vertices, for every color c in $\{1, \dots, k\}$ and every vertex v_i in G , v_i has at least one incoming edge of color c and at least one outgoing edge of color c . Hence, for every vertex v in G' and color c in $\{1, \dots, k+1\}$, v has at least one incident edge of color c that crosses the cut (L, R) . From this it follows that a random walk on G' on any infinite sequence visits the set L and the set R infinitely often with probability one.

Now suppose that the color sequence C has the property that colors from the set $\{1, \dots, k\}$ appear only a finite number of times in C . In this case, the sequence C can be written as $C'(k+1)^\omega$, where C' is a finite color sequence. Then, since the underlying graph colored $k+1$ is connected, a random walk on C covers G' with probability one. Similarly, if $k+1$ appears only a finite number of times in C , the graph G' is covered with probability one.

Assume now that colors from $\{1, \dots, k\}$ and the color $k+1$ appear infinitely often in C . Let E_L be the event that the random walk is at a vertex in L and the next color in the sequence is in the set $\{1, \dots, k\}$. Let E_R be the event that the random walk is at a vertex in R and the next color in the sequence is $k+1$. If on the random walk the events E_L and E_R occur infinitely often, the graph is covered with probability one. This is because there are cliques of each of the colors $1, \dots, k$ on the L vertices, and a clique of color $k+1$ on the R vertices.

On the other hand, if either of the events E_L or E_R happens only a finite number of times, then the sequence C must be of the form $C'(k+1)c_1(k+1)c_2(k+1)c_3 \dots$, where C' is a finite color sequence and each c_i is in $\{1, \dots, k\}$. Furthermore, the walk must be at some vertex $l_j \in L$ at the end of the walk on C' . In this case, the random walk on G' from l_j on $(k+1)c_1(k+1)c_2(k+1)c_3 \dots$ corresponds to a random walk on G from v_j on $c_1c_2c_3 \dots$. Since G is covered from all of its vertices, the graph G' is covered with probability one in this case. \square

Lemma 2.5 shows how to simulate a random walk on a k -colored directed graph with a random walk on a $(k+1)$ -colored undirected graph. We use the construction to prove the lower bounds that match our upper bounds on the expected cover time of colored undirected graphs.

By applying Lemma 2.5 to a family of strongly connected directed graphs with exponential expected cover time, we obtain Theorem 2.3. An example of such a family of graphs is given by a sequence of vertices numbered $1, \dots, n$ with a directed edge from vertex i to vertex $i+1$, for $1 \leq i \leq n-1$, and a directed edge from vertex i to vertex 1, for $2 \leq i \leq n$. Hence, we obtain the following theorem.

Theorem 2.3 *There are two-colored undirected graphs that are covered from all vertices and have expected cover time $2^{\Omega(n)}$.*

The doubly exponential lower bound for graphs with three or more colors is a consequence of Lemma 2.5 and the following theorem:

Theorem 2.4 (Condon and Lipton [10]) *There are two-colored directed graphs that can be covered from all vertices and have expected cover time $2^{2^{\Omega(n)}}$.*

On a particular sequence of colors a random walk on the n th graph in the family simulates 2^n tosses of a fair coin and reaches a designated state if and only if all outcomes were heads. In the paper by Condon and Lipton, the theorem is not stated as above but is instead stated in terms of proof systems with space-bounded verifiers. The result as stated is a consequence of the connection between two-colored directed graphs and proof systems, and the example is discussed in detail in Chapter 4.

By applying the construction of Lemma 2.5 to the family of graphs of Theorem 2.4, we obtain the following result:

Theorem 2.5 *There are three-colored undirected graphs that can be covered from all vertices and have expected cover time $2^{2^{\Omega(n)}}$.*

2.4 Concluding Remarks

There is a sizable gap between our upper and lower bounds on the expected cover time of two-colored graphs. The upper bound is obtained by proving that if G is a two-colored graph that is covered with probability one on all infinite sequences then, for all vertices s and t , $\text{dist}(s, t) \leq (4n - 3)(n - 1) = O(n^2)$. However, in the graph we construct for the lower bound, all pairs of vertices have distance $\text{dist}(s, t) = O(n)$.

This leaves us with the following interesting combinatorial problem. Let

$$d(n) = \max_{\substack{G=(\{n\}, E_1, E_2) \\ s, t \in \{n\}}} \text{dist}(s, t),$$

where the maximum is taken over only those two-colored graphs that are covered with probability one on all infinite color sequences. Our analysis shows that $d(n)$ lies somewhere between $\Omega(n)$ and $O(n^2)$. It is an interesting open question to determine the true asymptotic behavior of the function $d(n)$.

Chapter 3

Special Cases and Applications

3.1 Introduction

In this chapter we obtain tighter bounds on the expected cover time of colored graphs in a variety of interesting special cases. In most of these cases the proofs are elementary applications of known results about Markov chains. However, in the end we are able to use these results to prove an interesting theorem about the stationary behavior of Markov chains that are averages or products of random walks on connected undirected graphs with n vertices. In particular, we address the question of how the stationary distributions of random walks on undirected graphs scale under the operations of multiplication and addition. We begin this chapter by describing this application in detail.

Let G_1 and G_2 be a pair of connected nonbipartite undirected graphs with n vertices. Let \mathcal{M}_1 and \mathcal{M}_2 denote the finite regular Markov chains that correspond to simple random walks on G_1 and G_2 , respectively, and let M_1 and M_2 be their corresponding probability transition matrices. Let π_1 and π_2 be the unique stationary distributions of \mathcal{M}_1 and \mathcal{M}_2 , respectively. Since \mathcal{M}_1 and \mathcal{M}_2 correspond to random walks on undirected graphs, we know that all entries in π_1 and π_2 are at least $1/n^2$. Consider the Markov chain $\mathcal{M}_{\text{average}}$ defined by the probability transition matrix $M_{\text{average}} = \frac{1}{2}(M_1 + M_2)$. Since M_1 and M_2 correspond to connected nonbipartite graphs, it follows that $\mathcal{M}_{\text{average}}$ is an ergodic Markov chain. Hence, $\mathcal{M}_{\text{average}}$ has a unique stationary distribution π_{average} . We are interested in bounding the values of the entries of π_{average} as a function of the values of the entries of π_1 and π_2 . We will show that probabilities in π_{average} can be exponentially small in n , even though the probabilities in π_1 and π_2 are all inversely polynomial in n .

Similarly, we consider the Markov chain $\mathcal{M}_{\text{product}}$ defined by the probability transition

matrix $M_{\text{product}} = M_1 \cdot M_2$. Suppose that $\mathcal{M}_{\text{product}}$ is a regular Markov chain (this is not always the case; for an example, see Figure 5.1 in Chapter 5) and let π_{product} be the unique stationary distribution of $\mathcal{M}_{\text{product}}$. Again we show that the probabilities in π_{product} can be exponentially small in n , even though all probabilities in π_1 and π_2 are inversely polynomial.

The organization of this chapter is as follows. In Section 3.2 we obtain upper bounds on the expected cover time for two special classes of graphs. In Section 3.3 we prove upper bounds on the expected cover time for two special types of color sequences. In Section 3.4 we give an example that shows that all of the bounds given in Sections 3.2 and 3.3 are tight. In Section 3.5 we use the results from earlier sections to derive the above results about weighted averages and products of random walks on graphs.

3.2 Special Graphs

3.2.1 Proportional Colored Graphs

In this section we prove polynomial bounds on the expected cover time of a special class of colored undirected graphs, which we call *proportional graphs*.

A *proportional* colored graph is one in which

$$\frac{d_i(v)}{|E_i|} = \frac{d_j(v)}{|E_j|},$$

for all colors i and j , and all vertices v .

Theorem 3.1 *Let G be a proportional colored graph with n vertices that is covered from all of its vertices. If each of the underlying graphs of G is connected and nonbipartite, then the expected cover time of G is polynomial in n .*

Proof. Let c be any color. Since the underlying graph colored c is connected and nonbipartite, a random walk on the sequence $(c)^\omega$ is a simple random walk on the underlying graph (V, E_c) , which has a unique stationary distribution given by $\pi_c(i) = d_c(i)/2|E_c|$ for all vertices i . Since G is proportional, the distribution π_c is independent of c . Thus, we will use π to denote π_c for all c .

We wish to bound the expected cover time for a random walk on color sequence C starting from vertex s . Let v_0 be the n -dimensional row vector with a 1 in the position corresponding to s and a 0 in all other positions. In general, let v_t be the probability distribution of the random walk at

time t . The vector v_t is given by:

$$v_t = v_0 P_{C_1} \cdots P_{C_t}.$$

We will show that, for t polynomial in n , the distribution v_t is very close to the distribution π . We will use *pointwise distance* as a measure of distance between two distributions. The pointwise distance between v_t and π is given by :

$$\|v_t - \pi\| = \sum_i |v_t(i) - \pi(i)|.$$

Since, for every color c , P_c is the probability transition matrix of a simple random walk on a connected nonbipartite undirected graph, its largest eigenvalue is 1 with multiplicity one, and all of the other eigenvalues are at most $1 - n^{-3}$ in absolute value [19]. So for $t = n^4$, the pointwise distance $\|v_t - \pi\|$ is at most e^{-n} . Since each $\pi(i)$ is at least $1/n^2$, $v_t(i)$ is at least $1/cn^2$ for all i , where c is a positive constant. We can now derive bounds on the expected cover time by viewing the process as a coupon collector's problem on cn^2 coupons, where sampling one coupon takes n^4 steps of a random walk. The resulting bound on the expected cover time is $O(n^6 \log n)$. \square

3.2.2 Graphs with Self-Loops

Suppose that every vertex in G has a self-loop of every color at every vertex. That is, for every color i and vertex v , $(v, v) \in E_i$. We refer to these as *graphs with self-loops*. If each of the underlying graphs in a graph with self-loops is connected, then the graph is covered with probability one from all vertices. This is because for every pair of vertices s and t , the distance $\text{dist}(s, t)$ is at most $k(n - 1)$. In fact, it follows from this reasoning that the expected cover time of graphs with self-loops is at most exponential in n . This gives us the following theorem.

Theorem 3.2 *Let G be a colored graph with self-loops with n vertices. If each of the underlying graphs is connected then the expected cover time of G is at most exponential in n .*

Notice that graphs with self-loops satisfy the nonbipartite condition of Theorem 3.1, but in general the stationary distributions of the underlying graphs may be different. In fact, we will show in Section 3.4 that the bound of Theorem 3.2 is tight.

3.3 Special Sequences

In this section we assume, as usual, that the graph is covered from all start vertices, but will make no other assumptions about the graphs themselves. Instead we consider the behavior of random walks on special types of color sequences. The sequences we will consider are *random sequences* and *repeated sequences*.

3.3.1 Random Sequences

In this case, instead of analyzing the expected cover time on the worst case sequence, we will assume that at each time step the color is chosen randomly from the set $\{1, \dots, k\}$. If each of the underlying graphs is connected then the graph is covered from all its vertices. This is because for every pair of vertices s and t , a walk beginning at s visits t within $n - 1$ steps with probability at least $1/(nk)^{n-1}$. In fact, it follows from this reasoning that the expected cover time is at most exponential in this case. Notice that here the expectation is taken over both the random choices in the steps of the walk and the random choice of the color sequence.

Theorem 3.3 *Let G be a colored undirected graph with n vertices. If each of the underlying graphs is connected then the expected cover time on a randomly chosen color sequence is at most exponential in n .*

In Section 3.4 we will show that this bound is tight.

3.3.2 Repeated Sequences

We now consider the behavior of a random walk on sequences $(C_1 \dots C_l)^\omega$, where $C_1 \dots C_l$ is a fixed length color sequence. Again it is not difficult to see that the expected cover time is at most exponential in n . Since G is covered from all start vertices, for all vertices s and t , t is reachable from s on some prefix of $(C_1 \dots C_l)^\omega$. Let p be a shortest path from s to t on a prefix of $(C_1 \dots C_l)^\omega$. On a shortest path t appears once and every other vertex appears at most once in a position whose number is congruent to i modulo l , where $0 \leq i < l$. Hence, $\text{dist}(s, t)$ is at most $(n - 1)l$. This gives us the following theorem.

Theorem 3.4 *Let G be a colored undirected graph that is covered from all its n vertices and let $C_1 \dots C_l$ be a fixed length color sequence. The expected cover time of G on the sequence $(C_1 \dots C_l)^\omega$ is at most exponential in n .*

In Section 3.4 we will show that this bound is tight.

3.3.3 Corresponding Homogeneous Markov Chains

Random sequences and repeated sequences are similar because in both cases a random walk corresponds to a homogeneous Markov chain \mathcal{M} . In the case of a random sequence, the relevant Markov chain \mathcal{M} has probability transition matrix $\frac{1}{k} \sum_{i=1}^k P_i$, where P_i is the probability transition matrix for a simple random walk on the underlying graph colored i . In the case of a repeated sequence $(C_1 \dots C_l)^\omega$, every l steps of the random walk correspond to a single step with probability transition matrix $P_{C_1} \dots P_{C_l}$.

We can use the following lemma about homogeneous Markov chains to obtain a polynomial bound on the cover time for random and repeated sequences in a large number of special cases.

Lemma 3.1 *Let \mathcal{M} be an n -state homogeneous Markov chain with probability transition matrix M and let δ, ϵ be in the interval $(0, 1]$. Suppose that (1) M is irreducible and aperiodic, (2) all nonzero entries of M are at least δ , and (3) all entries of the stationary distribution of \mathcal{M} are at least ϵ . Then the expected time for the Markov chain \mathcal{M} to visit every state is at most $2n^2\delta^{-1}\epsilon^{-1}$.*

Proof. Consider the directed graph induced by the nonzero entries of M . That is, consider the graph $G = ([n], E)$, where $E = \{(i, j) : M(i, j) > 0\}$. Since M is irreducible there is a directed walk on G from any starting vertex that visits every vertex at least once and has length at most n^2 . We will bound the expected time for the process to complete such a walk on G .

Let i and j be a pair of adjacent vertices in the walk. We will bound the expected time for the process to traverse the edge from i to j . Each time the walk is at vertex i it traverses the edge from i to j with probability $M(i, j)$. Hence, the expected number of returns to i until the edge from i to j is traversed is $1/M(i, j)$. If $M(i, j) = 1$, the expected time to traverse the edge from i to j is 1. In what follows we will assume that $0 < M(i, j) < 1$.

Let $T(i, i)$ denote the mean recurrence time of vertex i . Then the expected time to return to i , given that the edge from i to j is not traversed, is at most $T(i, i)/(1 - M(i, j))$. Hence, the expected time for the walk to traverse the edge from i to j is at most $T(i, i)/M(i, j)(1 - M(i, j))$.

Since each non-zero entry of M is at least δ , $M(i, j)$ and $1 - M(i, j)$ are both at least δ . Hence, $M(i, j)(1 - M(i, j)) \geq \delta/2$, and the expected time for the walk to traverse the edge from i to j is at most $2\delta^{-1}T(i, i)$. Then, from the fact that the mean recurrence time of state i is the

reciprocal of its stationary probability $\pi(i)$, we get that the expected time for the walk to traverse the edge from i to j is at most $2\delta^{-1}\epsilon^{-1}$. It follows that the expected time for \mathcal{M} to visit every state is at most $2n^2\delta^{-1}\epsilon^{-1}$. \square

We can use Lemma 3.1 to obtain polynomial bounds for repeated and random sequences whenever the product and weighted average matrices satisfy its three conditions with δ and ϵ inversely polynomial in n . Conditions (1) and (2) are not particularly strong conditions. For example, the weighted average matrix satisfies condition (1) if the underlying graphs are connected and nonbipartite. The product matrix satisfies condition (1) if, for instance, the underlying graphs are connected and there is a self-loop of every color at every vertex. Products and weighted averages always satisfy condition (2) with δ inversely polynomial in n . Thus our question about polynomial expected cover time in graphs with self-loops on repeated sequences, and, in general, on randomly chosen color sequences becomes a question about the behavior of the stationary distributions of products and weighted averages, respectively. We state this formally below.

Theorem 3.5 *Let G be a colored undirected graph with n vertices such that each underlying graph is connected and nonbipartite. Suppose that the stationary distribution of the Markov chain with probability transition matrix $\frac{1}{k} \sum_{i=1}^k P_i$ has all entries bounded below by an inverse polynomial. Then the expected cover time of G on a randomly chosen color sequence is polynomial in n .*

Theorem 3.6 *Let G be a colored undirected graph with n vertices that is covered from all its vertices. Let $C_1 \dots C_l$ be a fixed length color sequence. Suppose that the Markov chain with probability transition matrix $P_{C_1} \dots P_{C_l}$ is irreducible and aperiodic, and its stationary distribution has all entries bounded below by an inverse polynomial. Then the expected cover time of G on $(C_1 \dots C_l)^\omega$ is polynomial in n .*

3.4 Lower Bounds

In this section we prove that the exponential upper bounds of Theorems 3.2, 3.3, and 3.4 are tight by constructing a two-colored graph with self-loops that has exponential expected cover time on a randomly chosen sequence of colors and on the sequence $(RB)^\omega$. The graph is shown in Figure 3.1. The solid lines represent edges colored R and the dotted lines represent edges colored B .




Figure 3.1: Graph for lower bounds

In what follows we prove that the expected cover time of the graph in Figure 3.1 on a randomly chosen sequence of colors is exponential in n . Our claim is that, on a randomly chosen color sequence, the expected time for a random walk that begins at vertex 1 to reach vertex n is exponential in n .

We refer to $1, \dots, n$ as the *primary vertices*, and $1', \dots, n'$ as the *secondary vertices*. Suppose a random walk from vertex i is performed on a randomly chosen sequence of colors until a primary vertex other than i is reached. We will call such a path a *primitive path*. The end of any primitive path from vertex i must be either $i + 1$ or $i - 1$. Let $p(i, i - 1)$ be the probability that the next primary vertex reached is $i - 1$, and let $p(i, i + 1)$ be the probability that the next primary vertex reached is $i + 1$. We will show that, for $2 \leq i \leq n - 1$, $p(i, i - 1)$ exceeds $p(i, i + 1)$ by a constant factor. Hence, the walk is biased backwards by a constant, and it is a routine calculation (see, for example, [15]) to show that the expected time to reach vertex n is exponential in n .

Let \mathcal{P}_i^+ be the set of primitive paths from i to $i + 1$, and let \mathcal{P}_i^- be the set of primitive paths from i to $i - 1$. Associated with each path p in \mathcal{P}_i^+ and \mathcal{P}_i^- is a probability, which is simply the product of the probabilities on the edges of p . We will establish a bijection ϕ from \mathcal{P}_i^+ to \mathcal{P}_i^- , with the property that, for every path p in \mathcal{P}_i^+ , the probability of p is strictly less than the probability of its image $\phi(p)$ in \mathcal{P}_i^- . It follows from this that $p(i, i - 1) > p(i, i + 1)$. Figure 3.2 shows the relevant transition probabilities for this argument.

Let p be a path $i = p_0, p_1, \dots, p_{l-1}, p_l = i + 1$ in \mathcal{P}_i^+ . The vertex p_{l-1} must be either i or i' . Suppose that $p_{l-1} = i$. Then we define $\phi(p)$ to be the path $i = p_0, p_1, \dots, p_{l-1}, i - 1$. The probability of the path $\phi(p)$ divided by the probability of the path p is equal to $4/3 > 1$. On the other hand, if $p_{l-1} = i'$ then let j be the largest index such that $p_j = i$. We define $\phi(p)$ to be the




Figure 3.2: Transition probabilities when color chosen at random

path of length l given by $i = p_0, p_1, \dots, p_j, (i - 1)', \dots, (i - 1)', i - 1$. The probability of the path $\phi(p)$ divided by the probability of the path p is equal to $15/14 > 1$.

This argument shows the existence of a sequence on which the expected cover time is exponential. A similar type of analysis can be used to show that $(RB)^\omega$ is one such sequence. The calculation, however, is tedious and is omitted.

3.5 An Application to Products and Weighted Averages

The construction given in Figure 3.1 has the following interesting application to the question posed at the beginning of this chapter. Let P_R and P_B be the probability transition matrices of the graphs colored R and B , respectively. Recall from the discussion in Section 3.3.3 that the matrices $P_R \cdot P_B$ and $(P_R + P_B)/2$ satisfy conditions (1) and (2) of Lemma 3.1 with δ inversely polynomial in n . So the fact that the expected cover time of this graph is exponential shows that the stationary distributions of $P_R \cdot P_B$ and $(P_R + P_B)/2$ each contain at least one entry that is exponentially small in n . But P_R and P_B correspond to undirected graphs, so all entries in their stationary distributions are inversely polynomial. So the example shows that, in general, it is possible for the stationary distribution of a product or weighted average of random walks on graphs to contain exponentially small entries, even though all entries of the stationary distributions of the original random walks are inversely polynomial.

Chapter 4

Colored Graphs and Complexity Classes

4.1 Introduction

Two-colored *directed* graphs were first studied by Condon and Lipton [10] in their investigation of the power of interactive proof systems with space-bounded verifiers.

In an interactive proof system a *prover* P wishes to convince a *verifier* V that a given shared input string x is a member of some language L . The prover and the verifier share independent read-only access to the input string x . The verifier V also has a private read-write worktape and the ability to toss coins during its computation.

In a general system, the computation proceeds in rounds. In each round, the verifier tosses a coin and asks a question of the more powerful prover. Based on the answers of the prover, the computation continues until eventually the verifier decides to accept or reject x and halts by entering an accepting or rejecting state.

Interactive proof systems in which the verifier is a probabilistic polynomial time Turing machine have been studied extensively in the literature. Results such as $IP = PSPACE$ [23], and $NEXPTIME \subseteq MIP$ [4] in the case of multiple provers, have characterized the class of languages recognized by such systems. Interactive proof systems have also been used to prove hardness of approximation for a class of combinatorial optimization problems known as MAX SNP in a series of papers [14], [3], [2] and others.

The systems considered by Condon and Lipton [10] and in this chapter differ from the standard ones in two ways. The first is that they are *one-way*, meaning that all communication goes *from* the prover *to* the verifier. Secondly, we are interested in verifiers V that are *space-bounded*; that is, verifiers that write on at most $s(n)$ tape squares on all inputs of length n . In particular, we

will be interested in systems where V uses space $O(\log n)$. We will use the term $\text{IP}_1(\text{SPACE}(\log n))$ to denote the class of languages with one-way interactive proofs with logspace verifiers. Related systems have also been studied in [13]. Since the system is one-way we can think of the prover as being represented by a *proof string* and the verifier as having one-way read-only access to the proof. As we will see, colored graphs are closely related to the class $\text{IP}_1(\text{SPACE}(\log n))$.

In this chapter we define the class $\text{IP}_1(\text{SPACE}(\log n))$. Our definition differs slightly from that used by Condon and Lipton, but the differences are purely technical. Once we have defined $\text{IP}_1(\text{SPACE}(\log n))$ we will review the correspondence between this class and two-colored directed graphs. We will prove that every language in PSPACE has a one-way interactive proof system with a logspace verifier. This result will be used at the end of this chapter and throughout Chapter 5 to prove that certain problems about colored graphs and from the theory of nonhomogeneous Markov chains are PSPACE-complete.

4.2 One-way Interactive Proof Systems

A *verifier* for language L is a three-tape probabilistic Turing machine V that takes as input a pair (x, π) , where x and π are strings over the alphabet $\{0, 1\}$. The string π is called a *proof*, and can be infinitely long. The proof π is stored on a one-way infinite, read-only tape. The verifier is constrained to read π in one direction; in fact, for technical reasons we will require that the head on π begins on its leftmost symbol and moves to the right in every step. We will also assume, without loss of generality, that V flips one coin per time step. The string x is stored on a second read-only tape, but its length is finite, and the head on x can move in both directions. The third tape of V is a worktape, which is initially inscribed with blanks. We will assume without loss of generality that V has exactly two halting states, an accepting state q_{accept} and a rejecting state q_{reject} , and that V erases its entire worktape and returns its input and worktape heads to the leftmost square before it accepts or rejects.

Let x be any string in $\{0, 1\}^n$. A language L is in $\text{IP}_1(\text{SPACE}(\log n))$ if there exists a verifier V that on input x uses $O(\log n)$ space on its worktape and satisfies the following halting and one-sided error conditions:

1. If x is in L , there exists a (finite) proof $\pi \in \{0, 1\}^*$ such that V accepts (x, π) with probability 1.
2. If x is not in L , then on any proof π , V rejects (x, π) with probability at least $2/3$.

3. V halts (accepts or rejects) with probability 1 on all inputs (x, π) . In fact, starting from any possible configuration of its worktape, state and tape heads, V halts with probability 1.

4.2.1 Example: Coin Flipping Protocol

Condon and Lipton [10] give the following example for $L = \emptyset$ to show that there exist one-way interactive proof systems with logspace verifiers that halt on all inputs and take doubly exponential time to halt on some input. We have adapted their example to satisfy our technical condition that the verifier read one bit of the proof in every step.

The verifier V behaves as follows on any input x of length n . Let $k = \lceil \log n \rceil$. Let i be an integer in the range 0 to $2^n - 1$. Consider the encoding of i as an $(n + k)$ -bit binary string. In this encoding the first k bits are zero and the remaining n bits are the usual binary encoding of i . Let C denote the $2^n(n + k)$ -bit string that consists of the encodings of the numbers 0 through $2^n - 1$.

On any proof string the verifier V flips one coin for each $(n + k)$ -bit disjoint substring, and maintains a single bit which tells whether all the coin flips so far were heads. Whenever V encounters the encoding of the number $2^n - 1$, it halts and rejects if all coin flips were heads. Otherwise, it resets the bit and repeats the process.

On the proof C^ω , V repeatedly flips 2^n coins and halts if and only if all 2^n outcomes were heads. Hence, the expected time for V to halt on the proof C^ω is doubly exponential in n . The verifier, however, does not halt with probability one on all inputs. In fact, if the encoding of $2^n - 1$ never appears in the proof, then V will never halt.

For this reason the verifier V must check that the proof string consists of the encodings of the numbers 0 through $2^n - 1$. Since V has only logarithmic space, it must do this probabilistically. While V scans the string of k zeros that begins the i th substring it flips k coins. The outcome of the k coin flips selects a random position b in the i th substring to check for consistency with the $(i + 1)$ st substring. When the proof is advanced to bit b of the i th substring the verifier checks whether the bit is a zero or a one. It then counts and advances through to the b th position in the $(i + 1)$ st substring. As it does this it remembers the logical AND of all of the lower order bits of the i th substring. If all of the lower order bits are one, it looks for the corresponding bit in the $i + 1$ st substring to be the flip of bit b in i . Otherwise, it looks for the two bits to be equal. If the test fails, the verifier V halts and rejects. Otherwise, it continues. The consistency check of the $(i + 1)$ st substring with the $(i + 2)$ nd substring overlaps with this check in the obvious way.

If the proof contains the encoding of $2^n - 1$ an infinite number of times in π , then the

verifier V halts with probability one. If the encoding of $2^n - 1$ appears only a finite number of times, then we can write the proof π as $\pi_1\pi_2$, where π_1 consists of all of the $(n + k)$ -bit substrings up to the last occurrence of $2^n - 1$, and π_2 consists of the rest of π . Then each subsequence of π_2 of length $2^n(n + k)$ contains at least one inconsistency, and V detects the inconsistency and halts with positive probability 2^{-k} . Hence, V halts with probability one in this case.

4.3 Two-colored Directed Graphs

Two-colored directed graphs were introduced by Condon and Lipton in their study of proof systems with space-bounded verifiers. We review the correspondence between proof systems with logspace verifiers and two-colored directed graphs here.

Let V be a logspace verifier and let x be an input of length n for V . A *configuration* of V is a quadruple (q, w, h_w, h_i) , where q is the state of V , w is a string representing the contents of the $O(\log n)$ bit worktape, h_w is the position of the head on the worktape, and h_i is the position of the head on the input tape, all encoded in binary. Notice that on inputs of length n , the number of possible distinct configurations of V is polynomial in n .

Consider the graph G_x defined as follows. The vertices of G_x correspond to the configurations of V on input x . If the verifier in configuration C responds to reading a 0 on the proof string by moving randomly to a configuration in $\{C_1, C_2\}$, then there is an edge colored R from the vertex corresponding to C to the vertices for configurations C_1 and C_2 . The edges colored B encode the actions of the verifier when it reads a 1 in the proof analogously.

The verifier V has a unique starting configuration $v_0 = (q_0, \bar{b} \cdots \bar{b}, 0, 0)$, a unique accepting configuration $v_{\text{accept}} = (q_{\text{accept}}, \bar{b} \cdots \bar{b}, 0, 0)$, and a unique rejecting configuration $v_{\text{reject}} = (q_{\text{reject}}, \bar{b} \cdots \bar{b}, 0, 0)$. Since we have assumed that q_{accept} and q_{reject} are halting states of V , configurations v_{accept} and v_{reject} have no outgoing edges in G_x . In fact, v_{accept} and v_{reject} are the only sinks in G_x since condition 3 says that on any proof, from any configuration V reaches a halting state with probability one.

4.3.1 Example: Coin Flipping Protocol Revisited

We can now describe in detail the construction of a two-colored directed graph that is covered with probability one on all infinite sequences and has doubly exponential expected cover time. This example was used in Section 2.3 of Chapter 2 for the lower bound for undirected graphs

with three or more colors. The example is based on the coin flipping protocol of Section 4.2.1.

Let V be the $O(\log n)$ space verifier of Section 4.2.1. Let x be any string of length n and let G_x be the graph of configurations of V on input x . We will augment G_x with an edge colored R and an edge colored B from v_{accept} and v_{reject} to every vertex in G_x . We will call the resulting graph G'_x . Since V halts on all proofs, the graph G'_x is covered with probability one on all infinite sequences. However, on the color sequence which corresponds to the encoding of the numbers 0 through $2^n - 1$ repeated ad infinitum, the expected time to reach v_{reject} is doubly exponential in n .

4.4 Polynomial Space

In this section we will show that every language in PSPACE has a one-way interactive proof system of the type defined above. This result will be used later in this chapter to prove PSPACE-completeness for reachability problems in colored graphs and in Chapter 5 to prove PSPACE-completeness of problems from the theory of nonhomogeneous Markov chains. The technique used is similar to that used in the construction of Example 4.2.1.

Theorem 4.1 $PSPACE \subseteq IP_1(\text{SPACE}(\log n))$

Proof. Let L be any language in PSPACE, and let M be a binary Turing machine that accepts L using $p(n)$ space on inputs of length n , where p is a polynomial. Without loss of generality, assume that M counts its steps and halts and rejects if it detects that it has looped by repeating a configuration.

A *configuration* of M is an encoding of the tape contents, the head position and the state at a given time during the computation. Let Q be the state set of M . We will assume the states in Q are numbered 1 through $|Q|$. We will encode a tape square of M as a $(\lceil \log |Q| \rceil + 2)$ -bit binary string. The last bit of the string will be used to encode the contents (zero or one) of the tape square. The other $\lceil \log |Q| \rceil + 1$ bits will be used to encode the index of the current state of M if the head is currently scanning the square, and will contain all zeros otherwise. Let $k > 0$ be the smallest integer such that $2^k \geq p(n)$. We will represent a configuration using the encodings of the first 2^k tape squares.

We can now represent an accepting computation of M on x by the sequence of configurations in the computation. Since M detects when it loops and rejects, the number of configurations in an accepting computation is bounded. The sequence of configuration encodings will be preceded

by a string of k ones, and each pair of consecutive configuration encodings will be separated by a string of k ones.

An $O(\log n)$ space-bounded verifier V can check that a given position in a configuration is consistent with the next configuration. The verifier must simply remember $O(1)$ symbols of the configuration and then count to $2^k + k$, advancing through the encoding as it counts. When V has finished counting, it can check the corresponding positions in the next configuration.

The verifier can choose a random position in the configuration to check by tossing k coins while it reads the k ones that precede the configuration. The verifier will overlap the consistency check of configurations j and $j + 1$ with the consistency check of configurations $j - 1$ and j in the obvious way.

The verifier can check that the first configuration is correct; that is, that the computation of M begins in the start state with x on its tape. If this test fails, or if the rejecting configuration ever appears, then V rejects. The verifier can recognize when the accepting configuration appears. If the computation contains an inconsistency in any of the intermediate steps, V detects it with probability at least 2^{-k} and rejects.

To reduce the probability of error, we concatenate 2^{k+1} copies of the encoding of the computation of M on x . The verifier can count the copies as it does the consistency checks. If V checks 2^{k+1} computations and no consistency check fails, then V accepts. If π is finite in length and V reaches the end of π without accepting, then V rejects.

If x is in L , then on the proof π which is the encoding of an accepting computation of M on x repeated 2^{k+1} times, V accepts with probability one.

Suppose that x is not in L and let π be any proof. If the first $2^k + k$ symbols of π do not encode the starting configuration of M on x preceded by k ones, then V rejects. Assume that the starting configuration is correctly encoded, and suppose that the accepting configuration appears 2^{k+1} times in π . Consider π parsed into $\pi_1\pi_2 \dots \pi_{2^{k+1}}\pi'$. The string π_1 is the initial portion of π , up to and including the first occurrence of the accepting configuration. For $2 \leq i \leq 2^{k+1}$, π_i is the portion of π that follows π_{i-1} , up to and including the i th occurrence of the accepting configuration. The string π' is everything that follows the (2^{k+1}) st occurrence of the accepting configuration in π . Since x is not in L , for all $1 \leq i \leq 2^{k+1}$, there is an inconsistency in the computation encoded by π_i . So, for all $1 \leq i \leq 2^{k+1}$, V detects an inconsistency in π_i and rejects with probability at least 2^{-k} . Hence, the probability that V accepts is at most $(1 - 2^{-k})^{2^{k+1}} < 1/3$.

Suppose that the accepting configuration appears fewer than 2^{k+1} times in π . Let π' be all of π after the last occurrence of the accepting configuration. If π' is finite or if π' contains the

rejecting configuration, then V rejects. Suppose that π' is infinite and does not contain the rejecting configuration. Consider π' in pieces of length $(2^{2^k} + 1)(2^k + k)$. Since M counts its steps and rejects if it loops, each such piece contains an inconsistency. In each piece the verifier detects an inconsistency and rejects with probability at least 2^{-k} . Hence, V rejects with probability one in this case. \square

4.5 Colored Graph Connectivity

In Chapter 2 we gave upper and lower bounds on the expected cover time of colored undirected graphs that are covered from all start vertices. We now investigate the complexity of determining whether a given colored undirected graph satisfies this condition. This condition is a generalization of the connectivity property for undirected graphs, and we will show that it is complete for natural space-bounded complexity classes. And again, as in Chapter 2, the complexity of the problem differs significantly in the case of two colors versus three or more colors. More formally, we consider the following decision problem:

COLORED GRAPH CONNECTIVITY

INSTANCE: Colored undirected graph G

QUESTION: Is G covered from all start vertices with probability 1 on all infinite sequences?

and show that COLORED GRAPH CONNECTIVITY for graphs with two colors is complete for nondeterministic logspace (NL), and for graphs with three or more colors it is PSPACE-complete.

In general, there is a close relationship between space-bounded complexity classes and problems of reachability in graphs. For instance, associated with any $s(n)$ space-bounded Turing machine M and input x of length n there is a directed graph G_x with a vertex for each of the $O(n2^{s(n)})$ configurations of M on x , and an edge from C_i to C_j if configuration C_i yields configuration C_j in one step on M . The question of whether M accepts x is equivalent to the question of whether there is a path from the starting configuration to an accepting configuration in G_x . In the case that $s(n)$ is equal to $\log n$, the graph G_x has $O(n^2)$ vertices. This demonstrates that s - t CONNECTIVITY (i.e., given a directed graph G and vertices s and t , is there a path from s to t in G ?) is complete for NL.

Another example is the correspondence between one-way interactive proof systems with space-bounded verifiers and two-colored directed graphs described in Section 4.3 of Chapter 4. The results of this section generalize these ideas.

The organization of the rest of this section is as follows. In Section 4.5.1, we show that, in general, COLORED GRAPH CONNECTIVITY is in PSPACE, and that when restricted to graphs with two colors the problem is in NL. In Section 4.5.2 we show that COLORED GRAPH CONNECTIVITY is hard for NL, and that the problem on graphs with three or more colors is PSPACE-hard.

4.5.1 Space-bounded Algorithms

We begin by proving combinatorial conditions that are equivalent to the connectivity property for colored graphs. These conditions will be used to obtain algorithms that work within the space bounds stated above.

Lemma 4.1 *Let G be a colored undirected graph with n vertices. The following conditions are equivalent:*

- (1) G is covered from all start vertices with probability one on all infinite sequences.
- (2) For all vertices s and t , the distance $\text{dist}(s, t)$ is at most 2^n .

Proof. That (1) \Rightarrow (2) is simply Lemma 2.1. To see that (2) \Rightarrow (1), notice that since, for all s and t , $\text{dist}(s, t)$ is at most 2^n , a random walk of length $(n - 1)2^n$ on any color sequence from any starting vertex covers the graph with positive probability. It follows that any infinite random walk covers the graph with probability one. \square

We can now use condition (2) above to obtain an algorithm for colored graph connectivity that uses polynomial space. Given a colored graph G with n vertices, Lemma 4.1 tells us that G is *not* covered from all starting vertices if and only if there exists a pair of vertices s and t , and a color sequence C of length 2^n such that t is not reachable from s on any prefix of C .

We will demonstrate that a nondeterministic polynomial space-bounded Turing machine, given G , can recognize that G is *not* covered from all vertices. Then, since PSPACE is closed under complement and under the addition of nondeterminism, it follows that COLORED GRAPH CONNECTIVITY is in PSPACE.

A nondeterministic polynomial space-bounded Turing machine can simply guess vertices s and t and count to 2^n , guessing the sequence C one character at a time and verifying that t is not reachable from s on each successive prefix of C . For this verification a single $n \times n$ boolean matrix M must be stored. This algorithm is given in detail in Figure 4.1. Throughout, we use A_j to denote the $n \times n$ adjacency matrix for edges of color j . The algorithm in Figure 4.1 uses space that is polynomial in the size of its input and so we have that COLORED GRAPH CONNECTIVITY is in PSPACE.

```

guess distinct vertices  $s$  and  $t$ 

guess a color  $C_1$  and set  $M \leftarrow A_{C_1}$ 
for  $i = 2$  to  $2^n$  do
    if  $M(s, t) \neq 0$  then reject
    guess a color  $C_i$  and set  $M \leftarrow M \times A_{C_i}$ 

accept

```

Figure 4.1: PSPACE algorithm for COLORED GRAPH CONNECTIVITY

The connectivity problem for two-colored graphs can be solved in NL. This would appear to be an easy extension of the result above. The approach would be to prove a lemma analogous to Lemma 4.1 with 2^n replaced by $(4n - 3)(n - 1)$. However, in the algorithm of Figure 4.1 an $n \times n$ matrix is stored and this would violate the logarithmic space restriction. Instead we use the following equivalence:

Lemma 4.2 *Let G be a two-colored undirected graph with n vertices. The following conditions are equivalent:*

- (1) G is covered from all start vertices with probability one on all infinite sequences.
- (2) For all vertices s and t , t is reachable from s on a prefix of each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$.

Proof. Suppose that G is covered from all start vertices. Then, for any pair of vertices s and t , a random walk from s on any sequence of colors visits t with probability one. It follows that t is reachable from s on each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$.

For the converse, suppose that for all s and t , t is reachable from s on a prefix of each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$. Then, by Lemma 2.3, $\text{dist}(s, t)$ is at most $(4n - 3)(n - 1) = l$, for all s and t . It follows that a random walk of length $(n - 1)l$ on any sequence from any starting vertex covers the graph with positive probability. Hence, the graph is covered with probability one on all infinite sequences. \square

Now we are prepared to show that COLORED GRAPH CONNECTIVITY for two-colored graphs is in NL. A nondeterministic logspace machine can simply run through all vertices s and t and verify that there is a path from s to t on a prefix of each of $(R)^\omega$, $(B)^\omega$, $(RB)^\omega$ and $(BR)^\omega$. Since such paths, if present, have length bounded by either $n - 1$ or $2n - 1$, the machine can nondeterministically guess and check these paths using only logarithmic space. The overall algorithm is given in Figure 4.2. Throughout we use A_0 to denote the adjacency matrix for the edges colored R , and A_1 to denote the adjacency matrix for the edges colored B . The algorithm uses space that is logarithmic in n and so we have shown that COLORED GRAPH CONNECTIVITY for graphs with two colors is in NL.

4.5.2 Hardness Results

In this section we prove the following two main results about COLORED GRAPH CONNECTIVITY:


```

for all distinct vertices  $s$  and  $t$ 

    set  $v_0 \leftarrow s$ 

    /* Check for a path from  $s$  to  $t$  on  $(R)^\omega$  */
    guess a length  $l$  such that  $0 < l < n$ 
    for  $i = 1$  to  $l - 1$  do
        guess vertex  $v_i$  and if  $A_0(v_{i-1}, v_i) = 0$  then reject
    if  $A_0(v_{l-1}, t) = 0$  then reject

    /* Check for a path from  $s$  to  $t$  on  $(B)^\omega$  */
    guess a length  $l$  such that  $0 < l < n$ 
    for  $i = 1$  to  $l - 1$  do
        guess vertex  $v_i$  and if  $A_1(v_{i-1}, v_i) = 0$  then reject
    if  $A_1(v_{l-1}, t) = 0$  then reject

    /* Check for a path from  $s$  to  $t$  on  $(RB)^\omega$  */
    guess a length  $l$  such that  $0 < l < 2n$ 
    for  $i = 1$  to  $l - 1$  do
        guess vertex  $v_i$  and if  $A_{(i-1)\bmod 2}(v_{i-1}, v_i) = 0$  then reject
    if  $A_{(l-1)\bmod 2}(v_{l-1}, t) = 0$  then reject

    /* Check for a path from  $s$  to  $t$  on  $(BR)^\omega$  */
    guess a length  $l$  such that  $0 < l < 2n$ 
    for  $i = 1$  to  $l - 1$  do
        guess vertex  $v_i$  and if  $A_{i\bmod 2}(v_{i-1}, v_i) = 0$  then reject
    if  $A_{l\bmod 2}(v_{l-1}, t) = 0$  then reject

accept

```

Figure 4.2: NL algorithm for two-colored graphs

Theorem 4.2 *COLORED GRAPH CONNECTIVITY for graphs with two colors is NL-complete.*

Theorem 4.3 *COLORED GRAPH CONNECTIVITY for graphs with three or more colors is PSPACE-complete.*

We have already shown that the problem is in PSPACE in general, and in NL for graphs with two colors. We now complete the proofs of Theorems 4.2 and 4.3 by giving proofs of hardness.

Proof of Theorem 4.2. We have already shown that COLORED GRAPH CONNECTIVITY for graphs with two colors is in NL in Section 4.5.1. Here we prove that every problem in NL can be reduced to COLORED GRAPH CONNECTIVITY on a two-colored graph.

We will use the fact that STRONG CONNECTIVITY (i.e., given a directed graph G , is G strongly connected?) is complete for NL. The proof of this is a straightforward reduction from s - t CONNECTIVITY and can be found as an exercise Hopcroft and Ullmans' book [17] on the theory of computation.

Lemma 2.5 shows how to construct a two-colored undirected graph G' that is covered from all vertices if and only if G is strongly connected. Since the construction of Lemma 2.5 can be carried out by a logspace Turing machine transducer, this completes the proof. \square

Next we show that COLORED GRAPH CONNECTIVITY for graphs with at least three colors is PSPACE-complete. For this we will use the connection between two-colored directed graphs and one-way proof systems with logspace verifiers, along with the fact that every language in PSPACE has a one-way proof system with a logspace verifier (Theorem 4.1) and the construction of Lemma 2.5.

Proof of Theorem 4.3. We have already shown that COLORED GRAPH CONNECTIVITY is in PSPACE in Section 4.5.1. We now prove that every problem in PSPACE can be reduced to COLORED GRAPH CONNECTIVITY.

Recall from Theorem 4.1 that every language in PSPACE has a one-way proof system with a logspace verifier V . Let G_x be the two-colored *directed* graph defined in Section 4.3. Recall that the vertices of G_x correspond to the configurations of V on input x . The edges colored R encode the transitions of V when the next proof bit read is 0, and the edges colored B encode the transitions when the next proof bit is 1. A pair of edges of the same color leaving a vertex correspond to a random coin flip of the verifier. Recall that G_x has vertices v_0 , v_{accept} and v_{reject} ,

which correspond to the unique starting, accepting, and rejecting configurations of V , respectively. Recall also that v_{accept} and v_{reject} have no outgoing edges.

We will augment G_x with the following edges. There will be an edge (v_{accept}, v) colored c , for each vertex v for which there is an edge (v_0, v) of color c . There is also an edge colored R and an edge colored B from v_{reject} to every vertex in G_x . We will call the augmented graph G'_x . We now claim that G'_x is covered from all start vertices if and only if x is not in L . Since PSPACE is closed under complement, this gives the desired result.

For the *if* direction, suppose that x is not in L and let π be any proof. Since V halts from all starting configurations, a random walk on G'_x from any starting vertex on color sequence π reaches v_{accept} or v_{reject} with probability one. The probability that the walk reaches v_{reject} , given that it has reached one of these two vertices, is at least $2/3$. If v_{accept} is reached, by construction of G'_x the remainder of the walk simulates V from its starting configuration, so again v_{accept} or v_{reject} is reached with probability one, and v_{reject} is reached with probability at least $2/3$. Hence, a random walk on G'_x from any start vertex, on any infinite sequence, repeatedly reaches v_{reject} . Since there is an edge of each color from v_{reject} to every other vertex in G'_x , G'_x is covered with probability one.

For the *only if* direction, suppose that x is in L . Then there is a finite proof π that takes V from the starting configuration to the accepting configuration with probability one. On the sequence of colors corresponding to repeating π ad infinitum, a random walk on G'_x from v_0 never visits v_{reject} . This is because v_{accept} is repeatedly reached on each copy of π with probability one.

The remainder of the proof comes from converting the two-colored directed graph G_x to a three-colored undirected graph using the construction of Lemma 2.5. \square

Chapter 5

Applications

5.1 Introduction

In this chapter we use the machinery of colored graphs to prove complexity theoretic results about nonhomogeneous Markov chains. The questions that we consider are fundamental in the theory of nonhomogeneous Markov chains and have applications to the theory of coding and information of finite-state channels.

Recall that a finite nonhomogeneous Markov chain \mathcal{M} is defined by an infinite sequence M_1, M_2, M_3, \dots of $n \times n$ stochastic matrices, where M_i is the probability transition matrix for time step i . Natural complexity theoretic questions arise when we think of the matrices that define the nonhomogeneous Markov chain \mathcal{M} as being drawn from a finite set $\mathcal{C} = \{C_1, \dots, C_A\}$ of $n \times n$ stochastic matrices.

In this chapter we consider the problem of deciding, given such a set \mathcal{C} , whether all finite products, or *words*, over \mathcal{C} are indecomposable. In order for all words to be indecomposable, each of the individual matrices C_1, \dots, C_A must be indecomposable. It is also necessary that each of the individual matrices be aperiodic; if there is a word W of any length of period $p > 1$, then the word W^p is decomposable.

The condition that each of the matrices C_1, \dots, C_A be indecomposable and aperiodic, however, is not a sufficient condition. For example, consider the product of the matrices whose nonzero entries are represented by the directed graphs pictured in Figure 5.1. Although the individual matrices are indecomposable and aperiodic, their product is decomposable.

We show that the problem of deciding whether all words are indecomposable is PSPACE-complete. This problem is fundamental in information theory, as it is a necessary and sufficient




Figure 5.1: Individual matrices that are irreducible and aperiodic, but whose product is decomposable

condition for optimal coding over finite-state indecomposable channels.

In addition, we show that the related problem of deciding whether all infinite products are weakly ergodic is PSPACE-complete, and that to decide whether all infinite products are strongly ergodic is PSPACE-hard.

In Section 5.2 we motivate the results of this chapter by giving some background for the application to information theory. For more details a good source is the book by Cover and Thomas [11] or Shannon's original 1948 paper [24]. In Section 5.3 we give the proofs for the two main theorems of this chapter, described above.

5.2 Information Theory

5.2.1 Preliminaries

Information theory is concerned with the problem of transmitting messages or signals over a device known as a *channel*. We begin this section by defining some of the basic notions of information theory. For now we will be concerned only with those channels which transmit signals with no possibility of loss or corruption.

The *capacity* C of such a channel is defined to be $C = \lim_{t \rightarrow \infty} \frac{\log N(t)}{t}$, where $N(t)$ is the number of possible signals of duration t . In a simplified situation where the channel can transmit one of n possible messages per unit time, the capacity C is equal to $\log n$. In general, channel capacity is a measure of the maximum number of bits of information that can be transmitted per unit of time.

We can think of a discrete source as generating its message symbol by symbol, where successive symbols depend probabilistically on previous symbols. This setup is modeled by an ergodic Markov chain \mathcal{M} described by an $n \times n$ stochastic matrix M , and is powerful enough to model natural languages and continuous information sources discretized by a quantizing process.

The rate at which rate information is produced by the source \mathcal{M} is defined using entropy.

Entropy was defined by Shannon in his original 1948 paper [24]. The entropy of a discrete random variable X which takes on value $x \in \mathcal{X}$ with probability $p(x)$ is defined to be:

$$H(X) = - \sum_x p(x) \log p(x).$$

Intuitively, $H(X)$ measures the amount of uncertainty in the random variable X . Alternatively, entropy can be interpreted as the number of bits of information contained in the random variable X ; that is, the number of bits required, on average, to describe X . The entropy function $H(X)$ takes on values in the interval $[0, \log |\mathcal{X}|]$.

For instance, suppose that X is a random variable that is either 0 or 1, each with equal probability. Then the entropy $H(X)$ is equal to 1, which is the maximum value of the entropy in this case. On the other hand, suppose that X always takes the value 0. It is not surprising that in this case $H(X)$ is 0, since the random variable X contains no information.

The *joint entropy* of random variables X and Y , which take on values x and y with probability $p(x, y)$ is defined to be:

$$H(X, Y) = - \sum_{x, y} p(x, y) \log p(x, y).$$

In general, $H(X_1, \dots, X_n) = - \sum_{x_1, \dots, x_n} p(x_1, \dots, x_n) \log p(x_1, \dots, x_n)$.

The *conditional entropy of Y given X* is defined to be:

$$H(Y|X) = \sum_x p(x) H(Y|X = x).$$

The joint and conditional entropies of X and Y are related by the following identity:

$$H(X, Y) = H(X) + H(Y|X).$$

This identity has the following natural interpretation. It says that the amount of uncertainty in the pair of random variables X and Y is equal to the amount of uncertainty in X plus the amount of uncertainty in Y when X is known. Put another way, the number of bits required to express both X and Y is equal to the number of bits required to express X plus the number of bits required to express Y when X is known.

Let $\{X_i\}$ be a stochastic process. Then the entropy rate of $\{X_i\}$ is defined to be:

$$\lim_{n \rightarrow \infty} \frac{1}{n} H(X_1, \dots, X_n),$$

provided the limit exists. If $\{X_i\}$ is a stationary ergodic process then there is a theorem which says that the limit exists. In the special case that $\{X_i\}$ is an ergodic Markov chain \mathcal{M} with stationary distribution π and probability transition matrix M , the entropy rate is given by the following simple formula:

$$H(\mathcal{M}) = - \sum_{i,j} \pi(i) M(i,j) \log M(i,j).$$

Note that if \mathcal{M} is generating i.i.d. random variables $X_i = X$ then $H(\mathcal{M}) = H(X)$.

An analogue of the law of large numbers known as the Asymptotic Equipartition Property (AEP) says that for large N there is a typical set (i.e., a set of probability approaching 1) of about $2^{NH(\mathcal{M})}$ sequences of length N , each with probability about $2^{-NH(\mathcal{M})}$. This means that typical sequences of length N can be represented using approximately $NH(\mathcal{M})$ bits. Hence, the entropy rate is a measure of the average number of bits of information produced by \mathcal{M} per unit of time. Shannon [24] proved the AEP in the i.i.d. case and stated it for stationary ergodic processes. Later McMillan [20] and Breiman [6] proved the AEP for stationary ergodic processes. This classical result is known as the Shannon-McMillan-Breiman Theorem.

In the case of noiseless communication the *rate of information transmission* is defined to be $\min\{C, H\}$ where C is the capacity of the channel and H is the entropy or information rate of the source. When information is transmitted at a rate equal to the capacity C of the channel, the source and channel are said to be *properly matched*.

5.2.2 Noisy Communication and the Finite-State Channel

In *noisy communication* the input to the channel is subject to random noise during transmission. In general, the output of the channel is a function of the input to the channel, the state of the channel at the time of transmission and random noise. This model of a *finite-state channel* was formalized by Blackwell, Breiman and Thomasian [5].

Formally, a finite-state channel is defined by a *source* and a *channel*. The *source* is a pair (M, ϕ) , where M is a $D \times D$ stochastic matrix corresponding to an ergodic Markov chain, and ϕ is a function from $[D] = \{1, \dots, D\}$ to $[A] = \{1, \dots, A\}$. The *channel* is a set of $A R \times R$ stochastic matrices C_1, \dots, C_A , and a function ψ from $[R] = \{1, \dots, R\}$ to $[B] = \{1, \dots, B\}$.

The elements of $[D]$ are considered the states of the source, and the elements of $[R]$ are the states of the channel. The set $[A]$ is the input alphabet and the set $[B]$ is the output alphabet. Suppose that d and r are the states of the source and channel, respectively, at the beginning of a cycle. The source moves into a new state d' according to transition matrix M (i.e., $M(d, d')$ is the

probability that the new state is d' and emits $\phi(d')$, which is fed into the channel. The channel then moves into state r' according to the transition matrix $C_{\phi(d')}$ and emits $\psi(r')$, completing the cycle. In the next cycle d' and r' are the initial states of the source and channel.

The joint motion of the source and channel is described by the *source-channel matrix* \widehat{M} , a $DR \times DR$ stochastic matrix whose rows and columns are indexed by pairs (d, r) , where $d \in [D]$ and $r \in [R]$. The entry of \widehat{M} in the (d, r) th row and the (d', r') th column is given by $M(d, d')C_{\phi(d')}(r, r')$. A channel is called *indecomposable* if for every source the source-channel matrix is indecomposable.

Let $\{(d_n, r_n)\}$ be the Markov chain with probability transition matrix \widehat{M} . Consider the ergodic processes $\{x_n = \phi(d_n)\}$, $\{y_n = \psi(r_n)\}$ and $\{(x_n, y_n)\}$, and denote their entropies by $H(X)$, $H(Y)$ and $H(X, Y)$, respectively. The *capacity* of a finite-state indecomposable channel is defined to be the upper bound H over all sources M of $H(X) + H(Y) - H(X, Y)$. Recall that the joint entropy $H(X, Y) = H(X) + H(Y|X)$ measures the amount of information in X plus the amount of information in Y when X is known. Hence, $H = H(Y) - H(Y|X)$ can be interpreted as the amount of information received, less the amount of information that is due to noise in the channel. Intuitively, the capacity is the maximum possible rate of transmission of information; that is, the rate when the source is properly matched to the channel.

Let ϵ be an error probability in the interval $(0, 1]$. We say that it is *possible to transmit information at rate G* if, for all sufficiently large N , there exist $J = 2^{GN}$ distinct sequences a_1, \dots, a_J in $[A]^N$ and J disjoint subsets B_1, \dots, B_J of $[B]^N$ satisfying the following condition. For all $j \in [J]$ and $r \in [R]$, the probability that the output sequence is in B_j when the channel starts in state r with input a_j is at least $1 - \epsilon$. The rate G measures the number of bits of information that are effectively transmitted per unit of time.

The collection of pairs (a_j, B_j) is called a *code*. The sequences a_1, \dots, a_J are the *codewords*. These are the only sequences of length N transmitted by the sender. If the receiver receives a message $b_j \in B_j$, then he interprets the original message as having been a_j . This is called *decoding*. The probability that the receiver decodes incorrectly is at most ϵ .

Shannon's coding theorem states that it is possible to transmit information with arbitrarily small (but positive) error probability at any rate less than the channel capacity but at no greater rate. In [5] Blackwell, Breiman and Thomasian give a proof of Shannon's theorem for finite-state indecomposable channels.

Theorem 5.1 (Blackwell, Breiman and Thomasian [5]) *For any indecomposable channel it is*

possible to transmit at any rate less than the capacity of the channel but not at any greater rate.

To verify that this result is valid for a particular finite-state channel we must know that the channel is indecomposable. Towards this end the authors give the following necessary and sufficient condition for channel indecomposability.

Theorem 5.2 (Blackwell, Breiman and Thomasian [5]) *A channel C_1, \dots, C_A is indecomposable if and only if every finite word $C_{i_1} \cdots C_{i_k}$ is an indecomposable stochastic matrix, where $k = 1, 2, \dots$ and $i_j \in [A]$.*

5.3 Complexity Results

In this section we investigate the complexity of deciding, given a finite set $\mathcal{C} = \{C_1, \dots, C_A\}$ of $n \times n$ stochastic matrices, whether all words over \mathcal{C} are indecomposable. Several authors, motivated by the coding theorem, studied this question during the 1960s. Thomasian gave the first finite criterion for channel indecomposability in the following theorem.

Theorem 5.3 (Thomasian [25]) *Let $\mathcal{C} = \{C_1, \dots, C_A\}$ be a set of $n \times n$ stochastic matrices. All finite words over \mathcal{C} are indecomposable if and only if all words of length at most 2^{n^2} are indecomposable.*

Interestingly, the proof of Theorem 5.3 uses a similar idea to the one used in Chapter 2 in the proof of the doubly exponential upper bound on expected cover time. We include the proof here.

Proof of Theorem 5.3. Assume that there is a decomposable word over \mathcal{C} and let $W = C_{i_1} \cdots C_{i_l}$ be the shortest decomposable word. Suppose, for contradiction, that $l > 2^{n^2}$. Then, since there are only 2^{n^2} different types of $n \times n$ matrices, for some $j < k$, the word $C_{i_1} \cdots C_{i_j}$ is of the same type as the word $C_{i_1} \cdots C_{i_k}$. Hence, $C_{i_1} \cdots C_{i_j} C_{i_{k+1}} \cdots C_{i_l}$ is of the same type as W , and thus is a decomposable word of length strictly less than l , which is a contradiction. \square

As Thomasian points out in his paper, the result of Theorem 5.3 gives an immediate algorithm for channel indecomposability. The algorithm simply enumerates all words of length up to 2^{n^2} and checks that each one is indecomposable. The running time of this algorithm is doubly exponential in n . However, by eliminating the need to repeatedly examine matrices of the same type, we can solve this problem in singly exponential time as follows.

Consider the directed graph G whose vertices correspond to the 2^{n^2} different $n \times n$ zero-one matrices A_i . For every ordered pair of vertices A_i and A_j , there is a directed edge from A_i to A_j in G if, for some $C_k \in \mathcal{C}$, $\langle A_i \cdot C_k \rangle = A_j$. For every vertex A_i other than the identity matrix, mark A_i if it is decomposable. Since we can determine whether the matrix A_i is decomposable in $O(n^2)$ time using graph searching, we can construct and mark the graph G in time $O(2^{n^2}(n^2 + A))$. Now, there is a decomposable word over \mathcal{C} if and only if there is a path in G from the identity matrix I to some decomposable matrix $A_i \neq I$. We can determine whether such a path exists by performing a depth-first search of G from I . This takes time linear in the size of G . Hence, the total running time of this algorithm is $O(2^{n^2}(n^2 + A))$.

Even this exponential time algorithm is impractical for modest values of n . Several authors worked on improving Thomasian's procedure by reducing the length of the words that are examined. Using ideas from Hajnal [16], Wolfowitz [26] proposed the following improvement to Thomasian's procedure. A matrix M is *scrambling* if, for every pair of indices i_1 and i_2 , there exists an index i such that $M(i_1, i) > 0$ and $M(i_2, i) > 0$; that is, every pair of states share a common consequent. Wolfowitz observed that any word with a scrambling matrix as a factor is indecomposable; therefore, when running Thomasian's procedure one could disregard any word that is scrambling or contains a scrambling word as a subword.

In a subsequent paper, however, Paz [21] showed that even when scrambling matrices are discarded, Thomasian's procedure could be made to examine words of length as large as 2^{n^2-n} . In the same paper, Paz proposed an alternative decision procedure that examines words of length at most $\frac{1}{2}(3^n - 2^{n+1} + 1)$. Nevertheless, in the worst case algorithms based on any of these criteria take exponential time when the graph searching strategy is employed.

The result of Theorem 5.4 is two-fold. It first improves upon the exponential upper bound given above by showing that the problem can be solved in PSPACE. Secondly, it shows that it is unlikely that these exponential time algorithms will be substantially improved, by showing that the problem is PSPACE-hard.

The first part of the result is a simple observation based on Thomasian's criterion. Suppose that there is a decomposable word $W = C_{i_1} \cdots C_{i_l}$, where $i_j \in [A]$ and $l \leq 2^{n^2}$. A nondeterministic polynomial space-bounded Turing machine can generate the indices i_j , for $j = 1, \dots, l$, one at a time and incrementally compute $\langle M_j \rangle$, where $M_j = C_{i_1} \cdots C_{i_j}$. Once $\langle W \rangle$ has been computed, the algorithm can verify in polynomial time that W is indeed a decomposable word. Since PSPACE is closed under the addition of nondeterminism and under complement, this shows that Thomasian's criterion can be carried out in PSPACE.

For the proof of hardness we use the characterization of PSPACE by the class $IP_1(\text{SPACE}(\log n))$ from Chapter 4.

Theorem 5.4 *Given a set $\mathcal{C} = \{C_1, \dots, C_A\}$ of two or more $n \times n$ stochastic matrices, it is PSPACE-complete to decide whether all words over \mathcal{C} are indecomposable.*

Proof. We have already described how a polynomial space Turing machine can decide, on input $\mathcal{C} = \{C_1, \dots, C_A\}$, whether all words over \mathcal{C} are indecomposable. It remains to show that the problem is PSPACE-hard.

Let L be any language in PSPACE and let x be an input of length n for which we wish to determine whether $x \in L$. By Theorem 4.1, L has one-way proofs of membership that can be checked by an $O(\log n)$ space-bounded verifier V . As in the previous chapter, let G_x be the two-colored directed graph of the computation of V on x . Recall that the vertices of G_x correspond to configurations of V , and that v_0 , v_{accept} and v_{reject} correspond to the unique starting, accepting and rejecting configurations of V , respectively. Recall also that vertices v_{accept} and v_{reject} are the only sinks since they correspond to the two halting configurations. We will augment G_x with an edge (v_{accept}, v) of color c , for each vertex v for which there is an edge (v_0, v) of color c . We will also add a self-loop $(v_{\text{reject}}, v_{\text{reject}})$ in each of the two colors. We will call the resulting graph G'_x . Note that every vertex in G'_x has at least one outgoing edge of each color.

Let $\mathcal{C} = \{C_1, C_2\}$, where C_1 is the probability transition matrix for a random walk on the edges colored R , and C_2 is the probability transition matrix for a random walk on the edges colored B . We claim that x is not in L if and only if all words over \mathcal{C} are indecomposable. Since PSPACE is closed under complement, the result follows from this claim.

Suppose that $x \in L$. Then there is a finite proof π for which V accepts with probability one. Let l denote the length of π and let $W = C_{\pi_1} \cdots C_{\pi_l}$ be the word corresponding to π . Since the verifier V accepts with probability 1 on π , by construction of G'_x the entry of W whose row and column correspond to v_{accept} contains a 1. By construction of G'_x the entry of W whose row and column correspond to v_{reject} also contains a 1. Hence W is a decomposable matrix with at least two essential classes, one containing v_{accept} and another containing v_{reject} .

Suppose that $x \notin L$ and let π be any proof. Let $C_{\pi_1} C_{\pi_2} \dots$ be the infinite sequence of matrices corresponding to π . This sequence of matrices has a corresponding random walk on G'_x . Since V halts on all proofs, such a random walk eventually reaches either v_{accept} or v_{reject} . The probability that the walk reaches v_{reject} , given that it has reached one of these two vertices, is at least $2/3$. Suppose that the walk reaches v_{reject} . Then, by construction of G'_x , the walk stays in

v_{reject} forever. On the other hand, suppose that the walk reaches v_{accept} . Then, by construction of G'_x , the remainder of the walk simulates the computation of V from its starting configuration, so again one of v_{accept} or v_{reject} is reached, and v_{reject} is reached with probability at least $2/3$. Hence, v_{reject} is reached with probability 1 on $C_{\pi_1} C_{\pi_2} \dots$, and once it is reached the walk stays there forever. Hence, all words have a single essential class which contains only the index v_{reject} . \square

Theorem 5.5 *Given a set $\mathcal{C} = \{C_1, \dots, C_A\}$ of two or more $n \times n$ stochastic matrices, it is PSPACE-complete to decide whether all infinite products over \mathcal{C} are weakly ergodic. To decide whether all infinite products over \mathcal{C} are strongly ergodic is PSPACE-hard.*

The part of Theorem 5.5 concerning weak ergodicity is obtained as a corollary to Theorem 5.4 using the following result of Wolfowitz.

Theorem 5.6 (Wolfowitz [26]) *Let $\mathcal{C} = \{C_1, \dots, C_A\}$ be a set of $n \times n$ stochastic matrices. All infinite products over \mathcal{C} are weakly ergodic if and only if all finite words over \mathcal{C} are indecomposable.*

This equivalence shows that the problem of deciding whether all infinite products over \mathcal{C} are weakly ergodic is also PSPACE-complete.

The part of Theorem 5.5 concerning strong ergodicity is obtained by observing that in the proof of Theorem 5.4, if x is not in L then all infinite products converge to the $n \times n$ matrix in which all rows have a one in the column corresponding to v_{reject} and zeros elsewhere. On the other hand, if x is in L then there is an infinite product that is not weakly ergodic.

5.4 Concluding Remarks

In this chapter we have addressed the computational complexity of deciding, given a finite set $\mathcal{C} = \{C_1, \dots, C_A\}$ of $n \times n$ stochastic matrices, whether all nonhomogeneous Markov chains defined as products over \mathcal{C} are ergodic. We have shown that deciding whether all products are weakly ergodic is PSPACE-complete. We have also shown that the related problem of deciding whether all finite words over \mathcal{C} are indecomposable is PSPACE-complete, and have discussed the application of this question to coding and information of finite-state channels. Our results show that these are hard problems and give strong evidence that the known polynomial space (exponential time) algorithms are the best possible.

We have also shown that to decide whether all infinite products over \mathcal{C} are strongly ergodic is PSPACE-hard. It is unclear how close this result comes to capturing the true computational complexity of the problem. Although recent work has addressed related questions [12] [18], no effectively computable algorithm is known.

Bibliography

- [1] R. Aleliunas, R. Karp, R. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proc. 20th Symposium on Foundations of Computer Science*, pages 218–223, 1979.
- [2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. In *Proc. 33rd Symposium on Foundations of Computer Science*, pages 14–23, 1992.
- [3] S. Arora and S. Safra. Probabilistic checking of proofs; A new characterization of NP. In *Proc. 33rd Symposium on Foundations of Computer Science*, pages 2–13, 1992.
- [4] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [5] D. Blackwell, L. Breiman, and A.J. Thomasian. Proof of Shannon’s transmission theorem for finite-state indecomposable channels. *Ann. Math. Stat.*, 29:1209–1220, 1958.
- [6] L. Breiman. The individual ergodic theorems of information theory. *Ann. Math. Stat.*, 28:809–811, 1957.
- [7] A. Broder and A. Karlin. Bounds on the cover time. *Journal of Theoretical Probability*, 2(1):101–120, 1989.
- [8] A. Condon and D. Hernek. Random walks on colored graphs. In *Proc. 2nd Israel Symposium on Theory and Computing Systems*, pages 134–140, 1993.
- [9] A. Condon and D. Hernek. Random walks on colored graphs. *Random Structures and Algorithms*, 5(2):285–303, 1994.

- [10] A. Condon and R. Lipton. On the complexity of space-bounded interactive proofs. In *Proc. 30th Symposium on Foundations of Computer Science*, pages 462–467, 1989.
- [11] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley and Sons, 1991.
- [12] I. Daubechies and J.C. Lagarias. Sets of matrices all infinite products of which converge. *Lin. Alg. Appl.*, 161:227–263, 1992.
- [13] C. Dwork and L. Stockmeyer. Finite state verifiers I: The power of interaction. *JACM*, 39(4):800–828, 1992.
- [14] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. In *Proc. 32nd Symposium on Foundations of Computer Science*, pages 2–12, 1991.
- [15] W. Feller. *An Introduction to Probability Theory and its Applications*. Wiley, 1950.
- [16] J. Hajnal. Weak ergodicity in non-homogeneous Markov chains. *Proc. of the Cambridge Philos. Soc.*, 54:233–246, 1958.
- [17] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [18] J.C. Lagarias and Y. Wang. The finiteness conjecture for the generalized spectral radius of a set of matrices. *Lin. Alg. Appl.*, 214:17–42, 1995.
- [19] H.J. Landau and A.M. Odlyzko. Bounds for eigenvalues of certain stochastic matrices. *Lin. Alg. Appl.*, 38:5–15, 1981.
- [20] B. McMillan. The basic theorems of information theory. *Ann. Math. Stat.*, 24:196–219, 1953.
- [21] A. Paz. Definite and quasi-definite sets of stochastic matrices. *AMS Proceedings*, 16(4):634–641, 1965.
- [22] E. Seneta. *Non-negative Matrices and Markov Chains*. Springer-Verlag, 1981.
- [23] A. Shamir. $IP = PSPACE$. In *Proc. 22nd ACM Symposium on Theory of Computing*, pages 11–15, 1990.
- [24] C.E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27:379–423, 1948.

- [25] A.J. Thomasian. A finite criterion for indecomposable channels. *Ann. Math. Stat.*, 34:337–338, 1963.
- [26] J. Wolfowitz. Products of indecomposable, aperiodic, stochastic matrices. *AMS Proceedings*, 14:733–737, 1963.