

*Published in: Surveys in Combinatorics, 1995
(ed. P. Rowlinson), London Math. Soc. Lecture Notes Series
218, Cambridge Univ. Press, 119–154.*

Mixing of Random Walks and Other Diffusions on a Graph

László Lovász¹ and Peter Winkler²

Abstract

We survey results on two diffusion processes on graphs: random walks and chip-firing (closely related to the “abelian sandpile” or “avalanche” model of self-organized criticality in statistical mechanics). Many tools in the study of these processes are common, and results on one can be used to obtain results on the other.

We survey some classical tools in the study of mixing properties of random walks; then we introduce the notion of “access time” between two distributions on the nodes, and show that it has nice properties. Surveying and extending work of Aldous, we discuss several notions of mixing time of a random walk.

Then we describe chip-firing games, and show how these new results on random walks can be used to improve earlier results. We also give a brief illustration how general results on chip-firing games can be applied in the study of avalanches.

1 Introduction

A number of graph-theoretic models, involving various kinds of diffusion processes, lead to basically one and the same issue of “global connectivity” of the graph. These models include: random walks on graphs, especially their use in sampling algorithms; the “avalanche” or “sandpile” model of catastrophic events, which is mathematically equivalent to “chip-firing” games; load balancing in distributed networks; and, somewhat more distantly but clearly

¹Dept. of Computer Science, Yale University, New Haven, CT 06510

²AT&T Bell Laboratories 2D-147, Murray Hill, NJ 07974

related, multicommodity flows and routing in VLSI. In this paper we survey some recent results on the first two topics, as well as their connections.

Random walks. The study of random walks on finite graphs, a.k.a. finite Markov chains, is one of the classical fields of probability theory. Recently interest has shifted from asymptotic results to inequalities and other quantitative properties involving a finite, possibly even very small number of steps. Much of this was motivated by applications to computer science. Perhaps the most important of these (though certainly not the only one) is *sampling* by random walk (see, e.g., [21], [33], [25]). This method is based on the fact that (at least for connected non-bipartite undirected graphs, which is easy to guarantee), the distribution of the current node after t steps tends to a well-defined distribution π , called the *stationary distribution* (which is uniform if the graph is regular). So to draw an (approximately) uniformly distributed random element from a set V , it suffices to construct a regular, connected, non-bipartite graph on V , and run a random walk on this graph for a large fixed number of steps.

A good example to keep in mind is shuffling a deck of cards. Construct a graph whose nodes are all permutations of the deck, and whose edges lead from each permutation to those obtainable from a single shuffle. Then repeated shuffle moves correspond to a random walk on this (directed) graph.

A crucial issue for this algorithm is the choice of the number of steps. Informally, let us call the necessary number of steps the *mixing time*. The surprising fact, allowing these algorithmic applications, is that this mixing time may be much less than the number of nodes. For example, it takes only 7 moves [11] to shuffle a deck of 52 cards quite well, using the standard “dovetail” shuffle—even though the graph has $52!$ nodes. On an expander graph with n nodes, it takes only $O(\log n)$ steps to mix.

At the same time, proving good bounds on the mixing time even in quite special cases is a difficult question. Various methods have been developed for this. Using eigenvalues, it is easy to find the *mixing rate*, i.e., the quantity

$$\lim_{t \rightarrow \infty} d(\sigma^t, \pi)^{1/t},$$

where σ^t is the distribution of the node we are at after t steps, and d is the total variation (ℓ_1 -)distance (or any other reasonable distance function). But this result does not tell the whole story for two reasons. First, the underlying graph in the cases of interest is exponentially large, (cf. the example of card shuffling), and the computation of the eigenvalues by the tools of linear algebra is hopeless. Second, the mixing rate tells us only the asymptotic behavior of the distance $d(\sigma^t, \pi)$ as $t \rightarrow \infty$, while we are interested in relatively small values of t (7 shuffle moves, for example). To be sure, eigenvalue methods can provide very sharp estimates, but for this, detailed information on the

spectrum, and even on the eigenvectors, is needed (see Diaconis [21] or Chung and Yau [17]). This kind of spectral information can be derived, it seems, only in the presence of some algebraic structure, e.g. a large automorphism group.

Therefore, combinatorial techniques that yield only bounds on the mixing rate and mixing time are often preferable. Two main techniques that have been used are *coupling* and *conductance*. We only give a brief discussion of the second; see [35] for more details.

Recent work by the authors provides a further method to prove bounds on the mixing time. Our work was motivated by the following observation. There is no particular reason why a walk used in a sampling algorithm must be run for a fixed number of steps; in fact, more general stopping rules which “look where they are going” are capable of achieving the stationary distribution exactly, and just as fast. Motivated by this, we have studied stopping rules that achieve any given distribution, when starting from some other given distribution. It turns out that there is a surprising variety of such rules, many of which are optimal in the sense that they entail the smallest possible expected number of steps; one of them also minimizes the *maximum* number of steps. These rules are related to important parameters of the random walk, like hitting times and conductance. The expected number of steps in an optimal rule serves as a natural (non-symmetric) distance between the initial and final distributions.

The most important special case arises when one wishes to generate a node from the stationary distribution, starting from a given node. The “distance” from a node to the stationary distribution, maximized over all nodes, provides a precise natural definition of *mixing time* (considered by Aldous [5], [6] in the reversible case). This notion agrees, up to a constant factor, with most of the usual definitions of mixing time, which depend on a specific choice of how nearness to the limit distribution is measured.

In these considerations, we assume that the graph is known, and we put no restriction on the computation needed to decide when to stop. This requirement makes direct use of our stopping rules as sampling mechanisms unlikely. (We show in [38] that it is possible to obtain the exact stationary distribution with an unknown graph, not in as efficient a manner, although still in time polynomial in the maximum hitting time.) However, one can describe a simple rule whose implementation requires no knowledge about the graph other than its mixing time, takes only a constant factor more time, and yields a node whose distribution is approximately stationary. The machinery we build to determine the mixing time may thus be considered as a tool for analyzing this simple and practical sampling mechanism.

A main tool in the analysis of random walks on graphs is the Laplacian of the graph. Mixing times, hitting times, cover times and many other important

parameters are closely related to the “eigenvalue gap” of this matrix, at least in the undirected case. A simpler but powerful tool is the “conservation equation” first noted by Pitman [42] (see Section 4).

Chip-firing and avalanches. Another diffusion process on graphs was introduced by Björner, Lovász and Shor [13] under the name of “chip-firing game”. We place a pile of chips on each node of a directed graph, and then change this arrangement of chips as follows: we select a node which has at least as many chips as its outdegree, and move one chip from this node to each of its descendants. We call this step *firing* a node. This step is repeated as often as we wish or until no node remains that can be fired.

Procedures equivalent to chip-firing games were introduced, independently, at least three times (not counting the obvious similarity to neural nets, which remains unexplored). Engel [26], [27] considered a procedure he called the “probabilistic abacus”, as a method of determining the limit distribution of certain Markov chains by combinatorial means. Spencer [47] introduced the special case when the underlying graph is a path, as a tool in analyzing a certain “balancing” game. In [4] Spencer’s process was analyzed in greater detail. The analysis of the procedure was extended to general (undirected) graphs in [13], and to directed graphs by Björner and Lovász [14].

Chip-firing turns out to be closely related to the “avalanche” or “sandpile” model of catastrophic events (also called self-organized criticality), introduced by Bak, Tang and Wiesenfeld [12] and Dhar [20]. The nodes of the digraph represent “sites” where snow is accumulating. There is a special node, the “outside universe”. Once the amount of snow on a site surpasses a given threshold, the site “breaks”, sending one unit of snow to each of its out-neighbors, which in turn can break again etc., starting an avalanche. After some easy reductions, avalanches can be considered as chip-firing games; even firing the special node can be viewed as a snowfall.

A key property of these games is that from a given position, all sequences of firings behave similarly: either they all can be extended infinitely, or they all terminate after the same number of moves, with the same final position (Church-Rosser property). This was observed in [4] and in [20].

Considering a chip-firing process on a given digraph, we can ask a number of natural questions: will this procedure be finite or infinite? If finite, how long can it last? If infinite, how soon can it cycle? How many chips are needed for an infinite procedure? How does one determine if a given position (distribution of chips) can be transformed into another one by firings?

In the case of undirected graphs, these questions are more-or-less fully answered in [13], [14] and the work of Tardos [48]. For example, a finite procedure terminates in $O(n^4)$ steps; the shortest period of a periodic game is n ; the minimum number of chips that allow an infinite game is m , the number of edges. There are polynomial time algorithms to determine if a

position starts a finite or infinite game, and also to determine if two positions can be reached from each other. The case of directed graphs is more difficult, and the complexity of some of these questions is still open.

There is a strong connection between chip firings, random walks on graphs, and the Laplace operator. In particular, the “conservation equation” plays an important role. This connection in the undirected case was observed in [13]; the extension to the directed case is due to [14], where it was used to show that no terminating firing sequence is longer than a polynomial times the length of the period of a periodic firing sequence. (This extends the result of [48], the directed case.) The new results on mixing times of random walks give improvements of these results. A converse inequality, conjectured in [14], will also be proved here, using the conservation equation.

There are a number of other diffusion processes on graphs, which we do not survey here in detail. *Load balancing* in distributed networks seems to be very closely related. In this model, every node of a (typically undirected and regular) graph corresponds to a processor, and each processor i is given a certain amount w_i of workload. The processors want to pass load to each other along the edges, so that eventually their loads should be (approximately) equal. This is quite similar in spirit to random walks on a regular graph, where “probability” is passed along the edges and eventually equalized. Indeed, upper and lower bounds on the time needed to equalize the loads ([2], [32]) involve parameters familiar from the theory of random walks: expansion rate, conductance, eigenvalue gap. On the other hand, there is a substantial difference: in chip-firing and random walks, load is distributed among the neighbors of a node evenly; in the load-balancing models, usually only one neighbor gets load in one step. Still, we hope that some of the ideas used in the analysis of random walks (and perhaps chip firing) might be applicable to a larger class of distribution processes.

2 Random walks, hitting and mixing times

Consider a strongly connected digraph $G = (V, E)$ with n nodes and m edges (we allow multiple edges and loops). We denote by a_{ij} or $a(i, j)$ the number of edges from i to j , and by d_i^- and d_i^+ the indegree and outdegree of node i , respectively. If the graph is undirected, then $d_v = d_v^+ = d_v^-$ is the degree of the node.

A *random walk* on G starts at a node w_0 ; if after t steps we are at a node w_t , we move to any node u with probability $a(w_t, u)/d^+(w_t)$. Clearly, the sequence of random nodes $(w_t : t = 0, 1, \dots)$ is a Markov chain. The node w_0 may be fixed, but may itself be drawn from some initial distribution σ . We

denote by σ^t the distribution of w_t :

$$\sigma_i^t = \Pr(w_t = i).$$

We denote by $M = (p_{ij})_{i,j \in V}$ the matrix of transition probabilities of this Markov chain. So

$$p_{ij} = \frac{a_{ij}}{d_i^+}.$$

The rule of the walk can be expressed by the simple equation

$$\sigma^{t+1} = M^T \sigma^t,$$

(the distribution of the t -th point is viewed as a vector in \mathbb{R}^V), and hence

$$\sigma^t = (M^T)^t \sigma.$$

It follows that the probability p_{ij}^t that, starting at i , we reach j in t steps is given by the ij -entry of the matrix M^t .

If G is undirected (which is viewed as a special case of directed graphs, with each edge corresponding to a pair of arcs oriented in opposite directions) then this Markov chain is *time-reversible*. Roughly speaking, this means that every random walk considered backwards is also a random walk (see below for a precise definition). If, in addition, G is regular, then the Markov chain is *symmetric*: the probability of moving to u , given that we are at node v , is the same as the probability of moving to node v , given that we are at node u .

The probability distributions $\sigma^0, \sigma^1, \sigma^2, \dots$ are of course different in general. We say that the distribution σ is *stationary* (or *steady-state*) for the graph G if $\sigma^1 = \sigma$. In this case, of course, $\sigma^t = \sigma$ for all $t \geq 0$. It is easy to see that there is a unique stationary distribution for every strongly connected digraph; we denote it by π . Algebraically, π is a left eigenvector of the transition matrix M , belonging to the eigenvalue 1.

A one-line calculation shows that for an undirected graph G , the distribution

$$\pi_i = \frac{d_i}{m} \tag{1}$$

is stationary (note that m is twice the number of undirected edges.) In particular, the uniform distribution on V is stationary if the graph is regular. An important consequence of this formula is that the stationary distribution is only a polynomial factor off the uniform (in terms of the number of edges, which we shall consider the input size of the graph. Loops and multiple edges are allowed.)

The stationary distribution for general directed graphs is not so easy to describe, but the following (folklore) combinatorial formula can be derived,

e.g., from Tutte’s “matrix-tree theorem”. Let A_i denote the number of all spanning in-arborescences in G rooted at i . Then

$$\pi_i = \frac{d_i^+ A_i}{\sum_i d_i^+ A_i} \quad (2)$$

The stationary distribution on a directed graph can be very far from the uniform; it is easy to find examples where the stationary probability of some nodes is exponentially small (in the number of edges). The value

$$\hat{\pi} = \min_i \pi_i$$

is an important measure of how “lopsided” the walk is. However, if the digraph is eulerian, then the stationary distribution is proportional to the degrees just like in the undirected case:

$$\pi_i = \frac{d_i^+}{m}.$$

Specifically, the uniform distribution is stationary for every regular eulerian digraph.

The most important property of the stationary distribution is that if the digraph is aperiodic, i.e., the cycle lengths in G have no common divisor larger than 1, then the distribution of w_t tends to the stationary distribution, as $t \rightarrow \infty$. (This is not true if the cycle lengths have a common divisor, in particular, for undirected bipartite graphs.)

In terms of the stationary distribution, it is easy to formulate the property of time-reversibility of the random walk on an undirected graph: for every pair $i, j \in V$, $\pi_i p_{ij} = \pi_j p_{ji}$. This means that in a stationary walk, we step as often from i to j as from j to i . From (1), we have $\pi_i p_{ij} = 1/m$ if $ij \in E$, so we see that we move along every edge, in every given direction, with the same frequency. If we are sitting on an edge and the random walk just passed through it, then the expected number of steps before it passes through it in the same direction again is m .

There is a similar fact for nodes, valid for all digraphs: if we are sitting at a node and the random walk just visited this node i , then the expected number of steps before it returns is $1/\pi_i$. If G is a regular eulerian digraph (in particular, a regular undirected graph), then this “return time” is just n , the number of nodes.

The *mixing rate* is a measure of how fast the random walk converges to its limiting distribution. This can be defined as follows. If the digraph is aperiodic, then $p_{ij}^{(t)} \rightarrow d_i/(2m)$ as $t \rightarrow \infty$, and the mixing rate is

$$\mu = \limsup_{t \rightarrow \infty} \max_{i,j} \left| p_{ij}^{(t)} - \frac{d_i}{2m} \right|^{1/t}.$$

One could define the notion of “mixing time” as the number of steps before the distribution of w_t will be close to uniform (how long should we shuffle a deck of cards?). This number will be about $(\log n)/(1-\mu)$. However, the exact value depends on how (in which distance) the phrase “close” is interpreted. Another concern is that this definition excludes periodic digraphs, and is very pessimistic in the case of “almost periodic” digraphs. For example, if G is obtained from a complete bipartite graph by adding an edge, then after a single step the distribution will alternate between almost uniform on one color class, and the other, but it takes a (relatively) long time before this alternation disappears. In applications to sampling, simple averaging tricks take care of this problem. Soon we will be able to introduce a more sophisticated, but “canonical” definition of mixing time.

In this paper, we do not study other important parameters of random walks, like cover times, commute times and the like. But one “time” will play an important role in the analysis of mixing speed: the *hitting time* (or *access time*) $H(i, j)$ is the expected number of steps before node j is visited, starting from node i . We denote by $H(G)$ the largest hitting time between any two nodes of the graph G . For undirected graphs, hitting times are polynomial in the number of edges ([1]). Brightwell and Winkler [15] proved that for every simple graph, $H(G) \leq (4/27)n^3$, and determined the graph that provides the maximum.

For digraphs, hitting times are not bounded by any polynomial of the number of edges in general. In fact, they are closely tied to the smallest stationary probability $\hat{\pi}$. Björner and Lovász proved in [14] that

$$H(G) \leq \sum_{i \in V} \frac{d_i^+}{\pi_i}, \quad (3)$$

which, together with the trivial lower bound, implies that

$$\frac{1}{\hat{\pi}} - 1 \leq H(G) \leq \frac{m}{\hat{\pi}}. \quad (4)$$

Hitting times have many interesting combinatorial and algebraic properties; see [35] for several of these. We only state here two special properties, for later reference. The *random target identity* states that

$$\sum_j \pi_j H(i, j) = C \quad (5)$$

is independent of the choice of i ; in other words, the expected number of steps we have to walk to hit a node randomly chosen from the stationary distribution is C , independent of the starting point (see, e.g., the “right averaging principle” in Aldous [5]).

The hitting time from i to j may be different from the hitting time from j to i , even in an undirected regular graph. Still, one expects that time-reversibility should give some sort of symmetry of these quantities. One symmetry property of hitting times for undirected graphs was discovered by Coppersmith, Tetali and Winkler [19]:

$$H(i, j) + H(j, k) + H(k, i) = H(i, k) + H(k, j) + H(j, i) \quad (6)$$

for every three nodes.

3 Mixing, eigenvalues and conductance

In this section we give a brief account of the use of these two tools in estimating the speed of mixing of a random walk. A more detailed survey, at least in the case of undirected graphs, can be found in [35].

The matrix M has eigenvalue 1, with corresponding left eigenvector π and corresponding right eigenvector $\mathbf{1}$, the all-1 vector on V . It follows from the Frobenius-Perron Theorem that every other eigenvalue λ satisfies $|\lambda| \leq 1$ and if G is non-periodic, then in fact $|\lambda| < 1$. We denote by μ the largest absolute value of any eigenvalue different from 1.

Now the key fact in the use of eigenvalue techniques is the following. Let σ be any starting distribution, then

$$\sigma^t - \pi = (M^T)^t(\sigma - \pi)$$

and hence it is easy to derive the following:

Theorem 3.1 *For every starting distribution σ , every $t \geq 1$ and $A \subseteq V$,*

$$|\sigma^t(A) - \pi(A)| \leq \frac{1}{\sqrt{\pi}} \mu^t.$$

Conductance. Let G be a digraph and $S \subset V$, $S \neq \emptyset$. Let $e(S, T)$ denote the number of edges connecting a set S to a set T . We define the *conductance* of the set $S \subset V$, $S \neq \emptyset$ by

$$\Phi(S) = \frac{1}{\pi(S)\pi(V \setminus S)} \sum_{i \in S} \pi(i) \frac{e(i, V \setminus S)}{d_i^+}$$

and the conductance of the graph by

$$\Phi = \min_S \Phi(S),$$

where the minimum is taken over all non-empty proper subsets $S \subset V$. If the graph is a d -regular and undirected, then the conductance of S is

$$\Phi(S) = \frac{n e(S, V \setminus S)}{d |S| \cdot |V \setminus S|},$$

which is (up to normalization) the edge-density in the cut determined by S .

To digest this quantity a little, note that $\sum_{i \in S} \pi(i) e(i, V \setminus S) / d_i^+$ is the frequency with which a stationary random walk switches from S to $V \setminus S$; while $\pi(S)\pi(V \setminus S)$ is the frequency with which a sequence of independent random elements of V , drawn from the stationary distribution π , switches from S to $V \setminus S$. So Φ can be viewed as a certain measure of how independent consecutive nodes of the random walk are.

Sinclair and Jerrum [33] established a connection between the spectral gap and the conductance of an undirected graph. A similar result for the related, but somewhat different parameter called *expansion rate* was proved by Alon [3] and, independently, by Dodziuk and Kendall [23] (cf. also Diaconis and Stroock [22]). All these results may be considered as discrete versions of Cheeger's inequality in differential geometry.

Theorem 3.2 *If G is an undirected graph, then every eigenvalue $\lambda \neq 1$ of M satisfies*

$$\lambda \leq 1 - \frac{\Phi^2}{8}.$$

This result allows an eigenvalue near -1 , which means that the graph is almost bipartite. While such an eigenvalue prevents us from applying 3.1 right away, it is in fact easy to handle. For example, we may attach d_i loops at each node i ; for the random walk on this modified graph we get

Corollary 3.3 *For any starting distribution σ , any $A \subseteq V$ and any $t \geq 0$,*

$$|\sigma^t(A) - \pi(A)| \leq \frac{1}{\sqrt{\hat{\pi}}} \left(1 - \frac{\Phi^2}{8}\right)^t.$$

See Diaconis and Stroock [22], Mihail [41], Fill [29], Sinclair [45], and also Lovász and Simonovits [36] for sharper bounds, connections with multicommodity flows, and for extensions to the directed case.

4 Stopping rules and exit frequencies

Examples. There are several examples of “stopping rules” that can achieve specified distributions in an elegant or surprising manner. We consider two; several more are mentioned in [39].

Consider the following interesting fact from folklore. Let G be a cycle of length n and start a random walk on G from a node u . Then the probability that v is the last node visited (i.e., the a random walk visits every other node before hitting v) is the same for each $v \neq u$.

While this is not an efficient way to generate a uniform random points of the cycle, it indicates that there are entirely different ways to use random walks for sampling than walking a given number of steps. This particular method does not generalize; in fact, apart from the complete graph, the cycle is the only graph which enjoys this property (see [37]).

Consider another quite simple graph, the cube, which we view as the graph of vertices and edges of $[0, 1]^n$. Let us do a random walk on it as follows: at each vertex, we select an edge incident with the vertex at random, then flip a coin. If we get “heads” we walk along the edge; if “tails” we stay where we are. We stop when we have selected every direction at least once (whether or not we walked along the edge).

It is trivial that after we have selected an edge in a given direction, the corresponding coordinate will be 0 or 1 with equal probability, independently of the rest of the coordinates. So the vertex we stop at will be uniformly distributed over all vertices.

This method takes about $n \ln n$ coin flips on the average, thus about $\frac{1}{2}n \ln n$ actual steps, so it is a quite efficient way to generate a random vertex of the cube, at least if we insist on using random walks (of course, to choose the coordinates independently is simpler and faster). We will see that it is in fact optimal.

Stopping rules. To begin a systematic study of stopping rules, we first define them. A *stopping rule* Γ is a map that associates with every walk w in the digraph G a number $0 \leq \Gamma(w) \leq 1$. We interpret $\Gamma(w)$ as the probability of continuing given that w is the walk so far observed, each such stop-or-go decision being made independently. We can also regard Γ as a random variable with values in $\{0, 1, \dots\}$, whose distribution depends only on the w_0, \dots, w_Γ ; thus we stop at w_Γ .

The *mean length* $E\Gamma$ of the stopping rule Γ is its expected duration; if $E\Gamma < \infty$ then with probability 1 the walk eventually stops, and thus σ^Γ is a probability distribution. A stopping rule Γ for which $\sigma^\Gamma = \tau$ is also called a *stopping rule from σ to τ* .

For any strongly connected digraph G and any distribution τ on $V(G)$, there is at least one finite stopping rule Γ such that $\sigma^\Gamma = \tau$; namely, we select a target node j in accordance with τ and walk until we reach j . We call this the “naive” stopping rule $\Omega_{\sigma\tau}$. Obviously, the mean length of $\Omega_{\sigma\tau}$ is given by

$$E\Omega_{\sigma\tau} = \sum_{i,j} \sigma_i \tau_j H(i, j).$$

In the case when $\tau = \pi$ is the stationary distribution, this formula can be simplified using the “random target identity” (5), and we get that the mean length of the naive rule to reach π is C , independently of the starting distribution.

We often think of a stopping rule Γ as a means of moving from a starting distribution σ to a given target distribution $\tau = \sigma^\Gamma$. Such a Γ is said to be *mean-optimal* or simply *optimal* (for σ and τ) if $E\Gamma$ is minimal. The mean length of a mean-optimal stopping rule from σ to τ will be denoted $H(\sigma, \tau)$. We call this number the *access time* from σ to τ , and think of it as a generalized hitting time.

Trivially, $H(\sigma, \tau) = 0$ if and only if $\sigma = \tau$. It is easy to see that the following *triangle inequality* is satisfied for any three distributions σ , ρ and τ :

$$H(\sigma, \tau) \leq H(\sigma, \rho) + H(\rho, \tau). \quad (7)$$

(To generate τ from σ , we can first use an optimal rule to generate ρ from σ and then use the node obtained as a starting node for an optimal rule generating τ from ρ). We should warn the reader, however, that $H(\sigma, \tau) \neq H(\tau, \sigma)$ in general.

We have seen that the access time $H(\sigma, \tau)$ has the properties of a metric on the space of node-distributions, except for symmetry; the latter is of course too much to expect since the ordinary hitting time, even for an undirected graph, is not generally symmetric.

Clearly if τ is concentrated at j (for which we write, rather carelessly, “ $\tau = j$ ”) then

$$H(\sigma, j) = \sum_i \sigma_i H(i, j), \quad (8)$$

since the only optimal stopping rule in this case is Ω_j , “walk until node j is reached.”

By considering the naive rule Ω_τ , we get the inequality

$$H(\sigma, \tau) \leq \sum_{i,j} \sigma_i \tau_j H(i, j). \quad (9)$$

This may be quite far from equality; for example, $H(\sigma, \sigma) = 0$ for any σ .

We set $H_{\max}(\tau) = \max_\sigma H(\sigma, \tau) = \max_i H(i, \tau)$. From the point of view of applications, stopping rules generating nodes from the stationary distribution are of particular interest. The value $T_{\text{mix}} = \max_i H(i, \pi)$ (the mean time of an optimum rule, starting from the worst point) is a natural and very useful definition of the mixing time.

It turns out that for given target distribution τ there are at least four interesting optimal stopping rules: the *filling rule*, the *local rule*, the *shopping rule* and the *threshold rule*. We describe these rules, together with some important non-optimal stopping rules, a bit later.

The conservation law. Let us now fix the digraph, a starting distribution σ and a finite stopping rule Γ . The expected number x_i of times the walk leaves node i before stopping will be called the *exit frequency* of node i for Γ . Clearly

$$E\Gamma = \sum_i x_i.$$

Exit frequencies were considered by Pitman [42]; he gave the following simple but very powerful “conservation law”, relating them to the starting and ending distributions:

Lemma 4.1 *The exit frequencies of any stopping rule from σ to τ satisfy the equation*

$$\sum_i p_{i,j} x_i - x_j = \tau_j - \sigma_j .$$

The identity expresses the simple fact that the probability of stopping at node j is the expected number of times j is entered minus the expected number of times j is left. The first application of this identity is the following theorem ([39]), relating different rules leading from the same starting distribution to the same target distribution:

Theorem 4.2 *Fix σ and let Γ and Γ' be two finite stopping rules from σ to τ with exit frequencies x and x' respectively. Let $D = E\Gamma - E\Gamma'$ be the difference between their mean lengths. Then $x' - x = D\pi$.*

It follows from Theorem 4.2 that the exit frequencies of any mean-optimal stopping rule from σ to τ are the same. We denote them by $x_i(\sigma, \tau)$.

Let us determine the exit frequencies in some simple cases. The first result is from Aldous [5]. Several related formulas could be derived using relations to electrical networks, as in [18] or [48].

Lemma 4.3 *The exit frequencies \tilde{x} for the naive stopping rule Ω_j in reaching node j from node i are given by*

$$\tilde{x}_k = \pi_k(H(i, j) + H(j, k) - H(i, k)).$$

More generally, the exit frequencies for the naive stopping rule Ω_τ from initial distribution σ are given by

$$\tilde{x}_k = \pi_k \sum_{i,j} \sigma_i \tau_j (H(i, j) + H(j, k) - H(i, k)) = \pi_k (E\Omega_{\sigma\tau} + H(\tau, k) - H(\sigma, k)) .$$

Combining this lemma with Theorem 4.2, we get the following general formula for exit frequencies:

Theorem 4.4 *The exit frequencies of a mean-optimal stopping rule from σ to τ are given by*

$$x_k(\sigma, \tau) = \pi_k \left(H(\sigma, \tau) + H(\tau, k) - H(\sigma, k) \right).$$

Any node j for which $x_j = 0$ is called a *halting node*. By definition we stop immediately if and when any halting node is entered. (But of course we may stop in other nodes too, just not all the time.) The following theorem from [39] gives an extremely useful characterization of optimality.

Theorem 4.5 *A stopping rule Γ is optimal if and only if it has a halting node.*

The “if” part is a trivial consequence of Theorem 4.2. The “only if” part is more difficult: we have to prove that from every σ to every τ there is a stopping rule that has a halting node. There are several ways to specify such a rule. Later on we shall describe four optimal stopping rules. Any of these could be used to prove this theorem, but none of the proofs is really straightforward, and we don’t give any of them here.

This theorem shows that from the two stopping rules on the cycle and the cube, discussed as introductory examples, the first is not optimal, but the second is (the node of the cube opposite the starting node is a halting node).

From Theorem 4.5, a formula for the access times follows easily. Consider an optimum stopping rule from σ to τ . Let j be any node, and consider the triangle inequality:

$$H(\sigma, j) \leq H(\sigma, \tau) + H(\tau, j)$$

The right hand side can be interpreted as the expected number of steps in a stopping rule that consists of first following an optimal rule from σ to τ and then following the naive rule (which is clearly the only optimal rule) from τ to j . Now if j is the halting node of the optimum rule from σ to τ then, trivially, it is a halting node for this composite rule, and so the composite rule is optimal. Thus for at least one j , equality holds. Rearranging, we get that

$$H(\sigma, \tau) = \max_j (H(\sigma, j) - H(\tau, j)). \quad (10)$$

Note that the access times on the right hand side can be expressed by the hitting times, using (8):

$$H(\sigma, \tau) = \max_j \sum_i (\sigma_i - \tau_i) H(i, j).$$

There is, in fact, a more general formula for the exit frequencies, which can be derived by similar arguments:

$$x_k(\sigma, \tau) = \pi_k \left(\sum_i (\tau_i - \sigma_i) H(k, i) - \min_j \sum_i (\tau_i - \sigma_i) H(j, i) \right)$$

In the special case of undirected graphs and target distribution π (which is perhaps the most common in applications of random walk techniques to sampling), we can use the cycle-reversing identity (6) and the random target identity (5) to obtain the following formula for the exit frequencies of an optimal rule:

$$x_k = \pi_k(\max_j H(j, i) - H(k, i)) \tag{11}$$

and

$$H(i, \pi) = \max_j H(j, i) - H(\pi, i) . \tag{12}$$

We have thus identified the halting node in an undirected graph, in attaining the stationary distribution from node i , as the node j from which the hitting time to i is greatest. This seems slightly perverse in that we are interested in getting from i to j , not the other way 'round!

Examples. Consider the classic case of a random walk on the path of length n , with nodes labeled $0, 1, \dots, n$. We begin at 0, with the object of terminating at the stationary distribution.

The hitting times from endpoints are $H(0, j) = H(n, n - j) = j^2$ and the stationary distribution is

$$\pi = \left(\frac{1}{2n}, \frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n}, \frac{1}{2n}\right).$$

The naive stopping rule has a halting node, namely n , and hence it is optimal. From (11) we have

$$x_k = \pi_k(H(n, 0) - H(k, 0)) = \pi_k H(n, k) = \pi_k(n - k)^2$$

and

$$H(0, \pi) = \sum_{i=0}^n x_i = \frac{1}{2n}n^2 + \sum_{i=1}^{n-1} \frac{1}{n}(n - i)^2 = \frac{n^2}{3} + \frac{1}{6} .$$

From this it is not difficult to derive that for a cycle of length n , n even,

$$H(s, \pi) = \frac{n^2}{12} + \frac{1}{6}$$

as compared with expected time

$$\frac{n-1}{n} \frac{n(n-1)}{2} = \frac{(n-1)^2}{2}$$

for staying at 0 with probability $1/n$ else walking until the last new vertex is hit as in the example discussed earlier.

The random walk on the following digraph is sometimes called the *winning streak*. Let $V = \{0, 1, \dots, n - 1\}$, and connect i to $i + 1$ by an edge for

$i = 0, \dots, n - 2$; also connect i to 0 for $i = 1, \dots, n - 1$. It is easy to check that the stationary distribution is

$$\pi_i = \frac{2^{n-i-1}}{2^n - 1}.$$

Hence the exit frequencies x_i for an optimal rule from 0 to π can be determined using Lemma 4.1, working backwards from $i = n - 1, n - 2, \dots$, obtaining

$$x_i = (n - i - 1) \frac{2^{n-i-1}}{2^n - 1}.$$

Summing over all nodes, we get

$$H(0, \pi) = \frac{(n - 2)2^n + 2}{2^n - 1} = n - 2 + O(n2^{-n}).$$

Next we describe four optimal stopping rules.

The filling rule. This rule is the discrete version of the “filling scheme,” introduced by Chacon and Ornstein [16] and shown by Baxter and Chacon [10] to minimize expected number of steps. We call it the *filling rule* (from σ to τ), and define it recursively as follows. Let p_i^k be the probability of being at node i after k steps (and thus not having stopped at a prior step); let q_i^k be the probability of stopping at node i in *fewer* than k steps. Then if we are at node i after step k , we stop with probability $\min(1, (\tau_i - q_i^k)/p_i^k)$.

Thus, the filling rule stops myopically as soon as it can without overshooting the target probability of its current node. One can prove that it is a finite stopping rule and thus it does in fact achieve τ when started at σ . One can also prove that it has a halting node.

The filling rule has a “now-or-never” property that once a node is exited, we never stop there later. In fact, it can be described in terms of “deadlines” g_j : we stop at node j if we hit it *before* this time; if we hit the node j at time t where $g_j < t \leq g_j + 1$, then we stop with probability $g_j + 1 - t$; we don’t stop if we hit it after time $g_j + 1$. A halting node j gets $g_j = \infty$.

The threshold rule. Every “threshold vector” $h = (h_1, \dots, h_n)$, $h_i \in [0, \infty]$ gives rise to a stopping rule in a manner opposite to the “deadlines” mentioned in connection with the filling rule: we stop at node j if we hit it *after* time $h_j + 1$; if we hit the node j at time t where $h_j < t \leq h_j + 1$, then we stop with probability $t - h_j$; we don’t stop if we hit it before time h_j . A rule obtained this way is called a *threshold rule*.

The threshold vector may not be uniquely determined by a threshold rule (e.g. all possible thresholds h_i smaller than the time before any possible walk reaches i are equivalent), but by convention we always consider the vector each of whose coordinates is minimal. Then in view of Theorem 4.5, the threshold

rule is optimal just when some coordinate of the associated threshold vector is zero.

Theorem 4.6 *For every target distribution τ there is a mean-optimal threshold rule.*

The threshold rule has a couple of further properties that distinguish it among other rules. First, if τ has sufficient support then it is bounded:

Theorem 4.7 *Suppose that every directed cycle contains a node i with $\tau_i > 0$. Then there is a $K > 0$ such that, for every starting distribution, the threshold rule always stops in fewer than K steps.*

The condition of Theorem 6 is also necessary in the sense that if a cycle with target probability 0 exists, then starting at a node of this cycle, no bound can be given on the number of steps in the threshold (nor on the number of steps of any other stopping rule).

The threshold rule is special among all rules in the following sense:

Theorem 4.8 *The maximum number of steps taken by the threshold rule is not larger than the maximum number of steps taken by any other rule from the same starting distribution to the same target.*

The local rule. Let x_i be the exit frequencies for an optimal stopping rule from σ to τ , i.e., solutions of the conservation law with $\min_i x_i = 0$. (An easy algebraic argument shows that for any σ and τ , there is a unique solution of the conservation equation with this property.) Consider the following “local” rule: if we are at node i , we stop with probability $\tau_i/(x_i + \tau_i)$, and move on with probability $x_i/(x_i + \tau_i)$ (if $x_i + \tau_i = 0$ the stopping probability does not need to be defined). Thus the probability of stopping depends only on the current node, not the time.

One can prove that the local rule generates τ . It is mean-optimality is clear since the node j with $x_j = 0$ is a halting node.

The shopping rule. Any probability distribution on the subsets of the node set V provides a stopping rule: “choose a subset U from ρ , and walk until some node in U is hit.” The naive rule is of course a special case, with ρ concentrated on singletons. The special case when ρ is concentrated on a chain of subsets is more efficient:

Theorem 4.9 *For every target distribution τ , there exists a unique distribution ρ which is concentrated on a chain of subsets and gives a stopping rule for generating τ . This stopping rule is optimal.*

The chain supporting the distribution ρ can be constructed recursively, starting from V and going down. Once we know that such a rule from σ to τ exists, its optimality is obvious, since a node in the smallest member of the chain is never exited.

Another rather neat way to think of this rule is to assign the real “price” $r(i) = \sum\{\rho(U) : i \in U\}$ to each node i . The “shopping rule” is then implemented by choosing a random real “budget” r uniformly from $[0, 1]$ and walking until a node j with $r(j) \leq r$ is reached.

The shopping rule shares with the filling rule Φ the “now-or-never” property that once a node is exited, it can never be the node at which the rule stops.

It is interesting to point out that the four stopping rules described above have a lot of common features. Of course, they all have the same exit frequencies and halting nodes. Each is described in terms of a numerical vector on V (deadlines, thresholds, exit frequencies, prices). Each of these vectors can be calculated from the starting and target distribution, by an algorithm that is polynomial in the number of nodes (which is unfortunately not good enough in a typical application of these techniques to sampling, where the number of nodes is exponential).

Each of these rules (or the corresponding vector) defines an ordering (with ties—technically, a “preorder”) of the nodes for every σ and τ . These orderings are in general different.

On the other hand, the four rules described above are different, and have in fact quite different properties. The threshold rule is bounded if, say, the target distribution has full support; the other three are, in general, not. The filling and shopping rules have the “now or never” property, but the other two rules do not. Finally, the filling rule has the “inverse boundedness” property that there is a time K so that it never stops after time K except in a halting node, which is not shared by any of the others.

5 Mixing times

We can define *the mixing time* of a random walk as $T_{\text{mix}} = \max_s H(s, \pi)$. This is not quite in line with the usual definition of mixing time, which is the smallest t such that, for every initial distribution σ , the distribution σ^t of the t -th element is “close” to π in one sense or another. To be specific, say we want $\sigma_i^t \geq (9/10)\pi_i$ for all i . (In [37], the dependence on a parameter c in place of $9/10$ is also studied, but here we simplify our discussion by fixing this value.)

It is not immediately clear how to compare these two definitions. On the one hand, the traditional definition requires only approximate mixing, so it

could be much less than our mixing time. On the other hand, the traditional definition is restricted to a trivial stopping rule (stop after t steps), and so it could be lead to much larger stopping times.

To be precise, we have to make one more point. If the graph is periodic (i.e., the lengths of its cycles have a common divisor larger than 1, say we have a bipartite graph), then σ^t may never be close to π . The way out is to do some kind of averaging: the (somewhat improperly named) “continuous time” model corresponds to choosing t from a Poisson distribution, while the “lazy walk” trick (see e.g. Lovász and Simonovits [36]) corresponds to choosing t from a binomial distribution.

It turns out that none of these differences mean too much, at least if we allow averaging. In fact, the following value is a lower bound on both versions of mixing time; on the other hand, both versions are at most a constant factor larger.

Let T_{fill} denote the smallest T such that for every starting distribution σ , there is a stopping rule Φ with mean length at most T such that $\sigma_i^\Phi \geq (9/10)\tau_i$ for all i . We can modify this definition, by using a different notion of approximation, the so-called total variation distance: let T_{tv} denote the smallest T such that for every starting distribution σ , there is a stopping rule Φ with mean length at most T such that $|\sigma^\Phi(A) - \pi(A)| \leq 1/10$ for every set A of nodes.

Obviously,

$$T_{\text{tv}} \leq T_{\text{fill}} \leq T_{\text{mix}}.$$

The last two quantities are always close to each other (this is a consequence of a simple folklore argument):

Theorem 5.1

$$T_{\text{mix}} \leq \frac{10}{9}T_{\text{fill}}.$$

On the other hand, T_{tv} and T_{fill} may be far apart, as the “winning streak” example shows. Aldous (see e.g [5]) proved a converse inequality in the time-reversible case. Adapted to our case (and improving the constant a little), this implies:

Theorem 5.2 *If the graph G is undirected, then*

$$T_{\text{mix}} \leq 4T_{\text{tv}}.$$

In the general case, the following inequality can be proved (the “winning streak” graph shows that it is tight).

Theorem 5.3

$$T_{\text{mix}} \leq O(\log(1/\hat{\pi}))T_{\text{tv}}.$$

Now we turn to the issue of how to implement optimal or near-optimal rules, to generate a node from the stationary distribution. It turns out that there exist simple, easily implementable rules that give a good approximation of the stationary distribution, while having a mean length only a constant factor more than the mixing time.

The *uniform averaging rule* $\Phi = \Phi_t$ ($t \geq 0$) is defined as follows: choose a random integer Y uniformly from the interval $0 \leq Z \leq t - 1$, and stop after Y steps. (To describe this as a stopping rule: stop after the j -th step with probability $1/(t - j)$, $j = 0, \dots, t - 1$).

Theorem 5.4 *Let Y be chosen uniformly from $\{0, \dots, t\}$. Then for any starting distribution σ , any $0 \leq c \leq 1$, and any $A \subseteq V$,*

$$|\sigma^Y(A) - \pi(A)| \leq \frac{1}{t}H(\sigma, \pi).$$

In particular, if $t \geq (1/\varepsilon)T_{\text{mix}}$ then

$$|\sigma^Y(A) - \pi(A)| \leq \varepsilon.$$

The contents of this (rather simple) theorem is that the averaging rule does as well as any sophisticated stopping rule, at least up to an arbitrarily small error and a constant factor in the running time. To illustrate how an “arbitrary” stopping rule can be related to the averaging rule, we sketch the proof.

Let Ψ be an optimal stopping rule from σ to π . Consider the following rule: follow Ψ until it stops at v^Ψ , then generate $Z \in \{0, \dots, t - 1\}$ uniformly and independently from the previous walk, and walk Y more steps. Since Ψ stops with a node from the stationary distribution, $\sigma^{\Psi+t}$ is also stationary for every $t \geq 0$ and hence so is $\sigma^{\Psi+Z}$.

On the other hand, let $Y = \Psi + Z \pmod{t}$, then Y is uniformly distributed over $\{0, \dots, t - 1\}$, and so

$$\begin{aligned} \sigma_i^Y &= \Pr(v^Y = i) \geq \Pr(v^{\Psi+Z} = i) - \Pr(v^{\Psi+Z} = i, \Psi + Z \geq t) \\ &= \sigma_i^{\Psi+Z} - \Pr(\Psi + Z \geq t). \end{aligned}$$

Hence for every set A of states,

$$\pi(A) - \sigma^Y(A) = \pi(A) - \sigma^{\Psi+Z}(A) + \Pr(v^{\Psi+Z} \in A, \Psi + Z \geq t) \leq \Pr(\Psi + Z \geq t).$$

Now for any fixed value of Ψ , the probability that $\Psi + Z \geq t$ is at most Ψ/t , and hence

$$\Pr(\Psi + Z \geq t) \leq \mathbb{E}(\Psi/t) = \frac{H(\sigma, \pi)}{t},$$

which proves the theorem.

Theorem 5.4 only asserts closeness in the total variation distance, not pointwise. Also, one would like that the error diminishes faster: it should be enough to choose t proportional to $\log(1/\varepsilon)$ rather than proportional to $1/\varepsilon$. We can give a slightly more complicated rule that satisfies these requirements. Choose $M = \lceil \log \varepsilon \rceil$, and let Y be the sum of M independent random variables, each being uniform over $\{0, \dots, \lceil 8T_{\text{mix}} \rceil\}$. Clearly $EY \approx 4T_{\text{mix}} \log(1/\varepsilon)$. Furthermore, stopping after Y steps gives a distribution very close to the stationary:

Theorem 5.5 *For any starting distribution σ ,*

$$\sigma^Y \geq (1 - \varepsilon)\pi.$$

(One little drawback in comparison with Theorem 5.4: we have to use the worst-case bound on the mixing time, not the access time from the given starting distribution.)

Blind rules. The averaging rules discussed above have an important property, which makes them practical but, at the same time, somewhat contrary to the philosophy of intelligent stopping rules: they don't look where they are. More exactly, let us call a stopping rule Γ *blind* if $\Gamma(w)$ depends only on the length of the walk w . Another way of describing a blind rule is to choose a non-negative integer Y from some specific distribution, and stop after Y steps.

The simplest blind stopping rule is the stopping rule used most often: "stop after t steps." Several other practical methods to generate elements from the stationary distribution (approximately) can also be viewed as blind rules. Stopping a lazy or continuous time random walk after a fixed number of steps corresponds to a blind rule for the original (discrete time) walk.

Our results above say that if we only want approximate mixing, then blind rules do essentially as well as any more sophisticated rule. The situation is very different if we want *exact* sampling. One cannot generate any distribution by a blind stopping rule; for example, starting from the stationary distribution, every blind rule generates the stationary distribution itself. We shall restrict our attention to stopping rules generating the stationary distribution (or at least approximations of it). Even this distribution cannot always be generated by a blind rule. The next theorem gives a characterization for the existence of a blind stopping rule for the stationary distribution.

Theorem 5.6 *Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of M , $\lambda_1 = 1$.*

(a) *If λ_k is positive real for some $k \geq 2$, then there exists a node s from which no blind stopping rule can generate π .*

(b) *If every λ_k , $k \geq 2$, is either non-real, negative or zero, then from any starting distribution σ there is a finite blind stopping rule that generates π .*

Interestingly, the condition formulated in the theorem is most restrictive for undirected graphs; then all the eigenvalues are real, and typically many of them are positive. If there are no multiple edges, only complete multipartite graphs give a spectrum with just one positive eigenvalue.

Almost blind rules for an unknown graph. Suppose that we do a random walk on a digraph that we do not know. We are told the number of vertices, and we are able to recognize a node if we have seen it before. It is easy to argue that no matter how long we observe the walk, it is impossible to compute the stationary distribution exactly. Thus it is a bit surprising that one can *achieve* it exactly. Nonetheless that is what is done by Asmussen, Glynn and Thorisson [9]: they give a stopping rule where the probability of stopping after a walk $w_0w_1w_2\dots$ depends only on the repetition pattern of nodes, and which produces a node from *exactly* the stationary distribution. The algorithm employed is complex and the expected number of steps required appears to be super-polynomial in the maximum hitting time $H(G)$, although no bound or estimate is given in the paper.

Note the emphasis on “exactly”. If we only require that the last node be approximately from the stationary distribution, then a natural thing to do is to stop, say, after Y steps, where Y is chosen, say, randomly and uniformly from a sufficiently long interval. It is not at all obvious how to know (just by observing the walk) how long is “sufficiently long”. But Aldous [7] describes a way to do so, and comes within total variation ε of the stationary distribution in time polynomial in $1/\varepsilon$ and linear in the maximum hitting time of the graph.

In [39] we describe a simple stopping rule which can reach the stationary distribution exactly, in any strongly connected digraph G . The rule requires only coin-flips for its randomization and can even be made deterministic unless the digraph is a single cycle (possibly with multiple edges). The expected number of steps is bounded by a polynomial in the maximum hitting time $H(G)$ of the graph.

The idea of the construction is to use formula (2) for the stationary distribution. Choose a node v *uniformly* from the set of all nodes. While observing the walk, mark the first exit from each node other than v . The edges we mark can be viewed as independent choices of one edge out of each node different from v . Hence given $v = i$, the probability that the $n - 1$ edges we marked form a spanning tree is $A_i / \prod_{j \neq i} d_j$. So by (2), the probability of getting an in-arborescence rooted at i , conditional on getting an arborescence at all, is just π_i .

Thus if the edges of first exits form an arborescence, we can walk until we hit v and stop; else, we start again.

Unfortunately, the probability of getting an arborescence may be exponentially small, which would result in an exponentially long algorithm (in

expected time). The trick is to replace the digraph by one whose adjacency matrix is a sufficiently large power of $I + A$; we omit the details here.

Other mixing measures. Theorems 5.1 and 5.2, and, in a weaker way, 5.3 are special cases of a surprising phenomenon, first explored by Aldous ([5], [6]). Mixing parameters of a random walk, that are only loosely related by their definition, are often very close. In fact, there seem to be three groups of parameters; within each group, any two are within (reasonably small) absolute constant factors to each other. For the time-reversible case (where these results are due to Aldous), the number of groups reduces to 2. We give a little “random walk” through some of these mixing measures.

Hitting times to sets. Let S denote a set of nodes and let $H(i, S)$ denote the expected number of steps before a random walk starting at i hits the set S . Of course, this number is larger if S is smaller, so it makes sense to scale by the stationary probability of S and define $T_{\text{set}} = \max_{s \in V, S \subseteq V} \pi(S)H(s, S)$. The upper bound in the following theorem is (in a somewhat different setting) due to Aldous, who also proved the lower bound for undirected graphs. The lower bound follows by an analysis of the shopping rule.

Theorem 5.7

$$\frac{1}{10}T_{\text{tv}} \leq T_{\text{set}} \leq 5T_{\text{tv}}.$$

We remark that sometimes the following upper bound may be stronger:

$$T_{\text{set}} \leq T_{\text{mix}}$$

(but here a reverse inequality can only be claimed in the undirected case).

Forget time and reset time. From the point of view of statistics, the following measure of mixing is important. The “forget time” T_{forget} of a random walk is defined as the minimum mean length of any stopping rule that yields a distribution τ from the worst-case starting distribution for τ . In other words,

$$T_{\text{forget}} = \min_{\tau} \max_{\sigma} H(\sigma, \tau) = \min_{\tau} \max_s H(s, \tau)$$

(since the worst starting distribution for any given target distribution τ is clearly concentrated on a single node). This notion is central to the modern theory of Harris-recurrent chains; see e.g. [8].

In applications to sampling algorithms, we almost always have to draw repeated samples; are later samples cheaper than the first sample? More exactly, suppose that we have a node j from the stationary distribution; how long do we have to walk to generate another node, also from the stationary distribution, independent of the first? It is clear that the optimum stopping

rule for this task is to follow an optimal stopping rule from j to the stationary distribution; so this stopping rule has mean length

$$T_{\text{reset}} = \sum_j \pi_j H(j, \pi),$$

which we call the *reset time* of the random walk. Trivially, $T_{\text{reset}} \leq T_{\text{mix}}$. The following result is proved in [40].

Theorem 5.8 *If the graph is undirected, then $T_{\text{forget}} = T_{\text{reset}}$.*

In the case of directed graphs, these two values may be arbitrarily far apart. But the theorem can be generalized to arbitrary digraphs in the form of an explicit formula for the forget time:

Theorem 5.9 *For every digraph,*

$$T_{\text{forget}} = \sum_j \pi_j \max_i H(i, j) - \sum_j \pi_j H(\pi, j).$$

($H(\pi, j)$ on the right hand side could be replaced by $H(k, j)$ with any $k \in V$ by (5).)

Using this formula, one can prove the following inequalities (for the case of undirected graphs, they were proved by Aldous [5]).

Theorem 5.10 *For every digraph,*

$$T_{\text{set}} \leq T_{\text{forget}} \leq 6T_{\text{tv}}.$$

(Hence $(1/10)T_{\text{tv}} \leq T_{\text{forget}} \leq 6T_{\text{tv}}$.) We conjecture that there is a constant c such that for any digraph, $T_{\text{mix}} \leq cT_{\text{reset}}$.

Maximum time and pointwise mixing. We have seen that the threshold rule was also optimal from the point of view that it minimizes the *maximum* number of steps needed to achieve the target distribution. If the target distribution is the stationary distribution, then we denote this maximum by T_{max} . This value may be quite different from the mean length of optimal stopping rules, even for undirected graphs. For example, let G be K_2 with N loops added on one of the nodes a single loop added on the other. It is easy to compute that $T_{\text{mix}} = (2N + 2)/(N + 3) \approx 2$, while (starting from the node with one loop), we need about $\log N$ steps to decrease the probability of staying there to the stationary value $2/(N + 3)$. Thus $T_{\text{max}} \approx \log N$.

A little unexpectedly, this value is also tied to mixing properties of the walk. Suppose that we want to generate any distribution τ such that $(9/10)\pi_i \leq \tau_i \leq (10/9)\pi_i$. If we allow an arbitrary stopping rule, then the

time needed for this is clearly between T_{fill} and T_{mix} , and since these two values are close by Theorem 5.1, we don't get anything new.

However, the situation changes if we use a *blind rule*. Let T_{pw} denote the smallest T such that there exists a blind rule with maximum length T that produces (from every starting distribution) a distribution τ such that $(9/10)\pi_i \leq \tau_i \leq (10/9)\pi_i$.

Theorem 5.11

$$T_{\text{max}} \leq 2T_{\text{pw}}$$

In particular, it follows that if we take, say, the uniform averaging rule then we have to average over the first $\Omega(T_{\text{max}})$ steps to achieve pointwise mixing (while in the filling sense, we only need $O(T_{\text{mix}})$ steps, and to achieve mixing in total variation distance, we only need $O(T_{\text{forget}})$ steps.

We conjecture that a reverse inequality also holds, in fact, averaging over $O(T_{\text{max}})$ steps yields a distribution that is pointwise close to the stationary.

6 Chip-firing

Let G be a strongly connected directed graph (many of the results below extend to general digraphs, but for simplicity of presentation we restrict our attention to the strongly connected case). Let us place a pile of s_i chips on each node i of G . Recall that *firing a node* means to move one chip from this node to each of its children. Clearly a node can be fired iff $s_i \geq d_i^+$. If no node can be fired, we call the vector $s = (s_i)$ a *terminal configuration*. A (finite or infinite) sequence of firings is called a *chip-firing game*. The sequence of points fired is called the *record* of the game. We denote by \mathcal{L}_s the set of all records of finite games starting from the same initial configuration s , where the digraph $G = (V, E)$ is fixed. For $\alpha \in \mathcal{L}_s$, we denote by $|\alpha|$ the length of α . The multiset of nodes occurring in α is called the *score* of α .

The following properties of \mathcal{L} have been proved in [13] for the undirected case, and extended to the directed case in [14]; they are also closely related to properties of abelian sandpiles proved by Dhar [20].

Proposition 6.1 *The set \mathcal{L}_s of records of chip-firing games starting with the same configuration s has the following properties:*

- (a) *Left-hereditary: whenever it contains a string, it contains all initial segments of the string.*
- (b) *Permutable: whenever $\alpha, \beta \in \mathcal{L}_s$ have the same score, and $\alpha x \in \mathcal{L}_s$ for some $x \in V$, we also have $\beta x \in \mathcal{L}_s$.*
- (c) *Locally free: whenever $\alpha x \in \mathcal{L}_s$ and $\alpha y \in \mathcal{L}_s$ for two distinct nodes $x, y \in V$, we also have $\alpha xy \in \mathcal{L}_s$.*

It turns out that these three simple properties have rather strong consequences. For example, it implies the following “antimatroid exchange property” (cf Korte, Lovász and Schrader [34]):

Proposition 6.2 *If $\alpha, \beta \in \mathcal{L}$ then there exists a subword α' of α such that $\beta\alpha' \in \mathcal{L}$ and the multiplicity of any v in $\beta\alpha'$ is the maximum of its multiplicities in α and β .*

The following theorem summarizes some of the results from [13], obtained using the above properties. It asserts that chip-firing games have a certain “Church-Rosser” property.

Theorem 6.3 *For a given directed graph G and initial distribution s of chips, either every chip-firing game can be continued indefinitely, or every game terminates after the same number of moves with the same terminal position. The number of times a given node is fired is the same in every terminating game. If a game is infinite, then every node gets fired infinitely often.*

In the case of undirected graphs, Tardos [48] proved a strong converse of the last assertion:

Lemma 6.4 *If a chip-firing game on an undirected graph is finite, then there is a node that is never fired.*

This assertion is analogous to Theorem 4.5 for stopping rules; however, it does not remain true for general digraphs. It was shown in [14] that it remains true for eulerian digraphs, and that it can be extended to digraphs in a different way (see Lemma 6.6 below).

Given a graph, we may ask: what is the minimum number of chips that allows an infinite game? What is the maximum number of chips that allows a finite game? In [13] it was shown that for an undirected graph with n nodes and m edges, more than $2m - n$ chips guarantees that the game is infinite; fewer than m chips guarantee that the game is finite; for every number N of chips with $m \leq N \leq 2m - n$, there are initial positions that lead to an infinite game and initial positions that lead to a finite game.

For directed graphs, the second question can still be answered trivially: if G is a directed graph with n nodes and m edges, and we have $N > m - n$ chips, then the game is infinite (there is always a node that can be fired, by the pigeonhole principle), and $N \leq n - m$ chips can be placed so that the game terminates in 0 steps.

It is not known how to determine the minimum number of chips allowing an infinite game on a general digraph. This is not just a function of the number of nodes and edges. For eulerian digraphs, it was mentioned in a

remark added in proof to [14] that the minimum number of chips that can start an infinite game is the *edge-feedback number*, i.e., the minimum number of edges whose removal destroys all directed cycles. Moreover, the feedback number is always a lower bound on the number of chips in an infinite game.

Chip conservation. A useful tool in the study of chip-firing games is the following “chip conservation equation” from [13] (cf. Lemma 4.1). Let s be the initial and t , the final configuration of a finite game, and let x_i denote the number of times the node i is fired. Let a_{ij} be the number of edges from node i to node j . Then

$$\sum_j a_{ji}x_j - d_i^+x_i = t_i - s_i. \quad (13)$$

Let $L \in \mathbb{R}^{V \times V}$ be the matrix defined by

$$L_{ij} = \begin{cases} a_{ij}, & \text{if } i \neq j, \\ -d_i^+, & \text{if } i = j. \end{cases}.$$

We call L the *Laplacian* of the digraph G . Note that $L\mathbf{1} = 0$, so L is singular. It is also well known that for strongly connected digraphs, the co-rank of L is 1. Let $v = (v_i)$ denote the solution of $L^\top v = 0$, scaled so that the v_i are coprime integers. From the Frobenius-Perron theory it follows that we may assume that $v > 0$. If G is an digraph (in particular, if G is an undirected graph), then $v = \mathbf{1}$. The quantities $|v| := \sum_i v_i$ and $\|v\| := \sum_i d_i v_i$ play an important role in chip-firing.

The Laplacian is also related to the transition probability matrix of the random walk:

$$M = D^{-1}L + I,$$

where D is the diagonal matrix with the outdegrees in the diagonal. It follows that the stationary probabilities are given by

$$\pi_i = \frac{d_i^+ v_i}{\|v\|}.$$

It follows by (3) that the maximum hitting time can be estimated as follows:

$$H(G) \leq \|v\| \sum_i \frac{1}{v_i} \leq n\|v\|. \quad (14)$$

In terms of the Laplacian, equation (13) can be written as

$$L^\top x = t - s.$$

Period length. As a first application of this identity, we discuss periodic games. More exactly, consider a period, i.e., a game that starts and ends with the same configuration. Let x be its score vector; then

$$L^\top x = 0,$$

whence it follows that $x = tv$ for some positive integer t . It is easy to see that $x = v$ can be realized: just place a very large number of chips on each node, and fire each node i v_i times in any order. The conservation equation implies that we return to the starting configuration.

A key property of the vector v is the following:

Lemma 6.5 *Let $\alpha \in \mathcal{L}_s$. Let α' be obtained by deleting the first v_i occurrences of node i in α (if i occurs fewer than v_i times, we delete all of its occurrences). Then α' is the record of a game from the same initial position.*

This lemma (which is easy to prove by counting chips) has a number of consequences. First, 6.2 implies that the deleted elements can be added to α' , and so we get a game that is a rearrangement of α but starts with α' . From this it is easy to derive that *if a configuration starts a periodic game, it also starts one with period score v* . We also get an extension of Lemma 6.4:

Lemma 6.6 *In every terminating game, there is a node i that is fired fewer than v_i times.*

From these considerations, one obtains:

Proposition 6.7 *The minimum length of a period of any game on the graph G is $|v|$, and the number of chips moved during a minimal period is $\|v\|$.*

Game length. Deviating from earlier papers, we measure the length of a game by the number of chip-motions (so the firing of a node of outdegree d contributes d to the length). This is of course an upper bound on the number of firings, and is never more than a factor of m larger.

Tardos [48] proved that on an undirected graph, every terminating game ends in a polynomial number of steps. We sketch a new proof based on the conservation equation. Consider a game that terminates, and let z_i be the number of times node i is fired. Then we have

$$\sum_j a_{ji} z_j - d_i z_i = t_i - s_i.$$

Here, by termination, $s_i < d_i$. We can rewrite this equation as

$$\sum_j \frac{a_{ij}}{d_j} \cdot \frac{d_j z_j}{m} - \frac{d_i z_i}{m} = \frac{t_i + d_i - s_i}{m} - \frac{d_i}{m}.$$

Thus the numbers $d_i z_i / m$ are the exit frequencies of a stopping rule from the distribution defined by $\tau_i = (t_i + d_i - s_i) / m$ to π . By Lemma 6.4, the minimum

of these exit frequencies is 0, and hence the stopping rule is optimal. This implies that

$$\sum_i \frac{d_i z_i}{m} = H(\tau, \pi) \leq T_{\text{mix}}.$$

Hence we get:

Theorem 6.8 *The number of chips moved during a terminating game on an undirected graph G is at most mT_{mix} .*

Eriksson [28] showed that on a directed graph (even on a graph with all but one edges undirected) a terminating game can be exponentially long. It was proved in [14] that the maximum length of a terminating game can exceed the period length by a polynomial factor only. It was conjectured that a converse inequality, bounding the period length by a polynomial multiple of the maximum game length, also holds. It turns out that this conjecture is true, and in fact it follows quite simply using the conservation equation. Results on random walks discussed above also yield an improvement in the first direction.

Theorem 6.9 *Let M denote the maximum number of chip-motions in a terminating finite game. Then*

$$\|v\| - m \leq M \leq nm\|v\|.$$

Sketch of proof. 1. Consider a terminating game, and let p and q be the beginning and terminating configurations of it. Obviously, $|p| = |q| < m$. Let u be the score vector of the game. By Lemma 6.6, there is a node i such that $u_i < v_i$. Hence we can write $u = tv + w$, where $0 \leq t < 1$ and $\min w_i = 0$. By the conservation equation, we have

$$L^\top u = q - p,$$

and hence we also have

$$L^\top w = q - p.$$

Let N denote the number of chips, then $N \leq m - n$ since the game terminates. We get that the numbers $x_i = d_j^+ w_j / N$ are the exit frequencies of an optimum stopping rule from $(1/N)p$ to $(1/N)q$. This implies that

$$\sum_i \frac{d_j^+ w_j}{N} = H\left(\frac{1}{N}p, \frac{1}{N}q\right) \leq H(G).$$

Thus by (14), the number of chips moved is

$$\sum_i d_i^+ u_i = t \sum_i d_i^+ v_i + \sum_i d_i^+ w_i < \|v\| + NH(G) < nm\|v\|.$$

2. To prove the other inequality (and thereby verify a conjecture from [14]), place $d_i^- - 1$ chips on node i . We claim that every game from this starting position is finite; in fact, we claim that no node can be fired v_i times. Assume that this is false, and consider the first step when a node i is fired the v_i -th time. Let y be the score vector up to and including this step, and q , the configuration after this step. Then the conservation equation says:

$$\sum_j a_{ji} y_j - d_i^+ y_i = q_i - (d_i^- - 1).$$

But here the left hand side is

$$\sum_j a_{ji} y_j - d_i^+ y_i \leq \sum_j a_{ji} (v_j - 1) + d_i^+ v_i = - \sum_j a_{ji} = -d_i^-,$$

which is a contradiction since $q_i \geq 0$.

Thus every game started from this position terminates. But with $m - n$ chips on board, the only terminating configuration is having $d_i^+ - 1$ chips on node i . Moreover, substitution into the chip conservation equation shows that in order to get from $d_i^- - 1$ chips to $d_i^+ - 1$ chips on each node, we have to fire every node i exactly $v_i - 1$ times. Hence there is always a terminating game of length

$$\sum_i d_i (v_i - 1) = \|v\| - m.$$

□

We have seen three relations between the three diffusion parameters we considered: the maximum hitting time $H(G)$, the period length $\|v\|$, and the maximum game length M . The last two are equal up to a polynomial factor, while the first is at most this large.

The hitting time can be much smaller than the other two quantities. Consider a 2-connected undirected graph G and orient one edge (leave the rest two-way). Then one can argue that the hitting time remains polynomial; on the other hand, the example of Eriksson mentioned above is of this type, and here the game length and period length are exponentially large.

Algorithmic issues. Results mentioned above were used in [14] to give an algorithm for checking whether a given position on an undirected graph can be transformed to another given position by a sequence of firings. The running time of the algorithm is polynomial in the period length $\|v\|$, so in the case of undirected graphs, it is polynomial in m . The idea is to use Lemma 6.6 in a manner similar to the proof of Theorem 6.9 to show that if there is a sequence of chip-firings then there is one of length polynomial in $\|v\|$, and in fact the firing frequencies z_i can be calculated by simple arithmetic. Then one can show that any game with the additional restriction that no node i is

fired more than z_i times, must terminate in the prescribed target position, or else the target position is not reachable.

Unfortunately, no truly polynomial algorithm is known to decide the reachability question. It is also not known how to decide in polynomial time whether a given initial position starts a finite or infinite game.

These questions are quite interesting because chip-firing on a digraph may be considered as a “totally asynchronous” distributed protocol (by Theorem 6.3). The comparison of the class of functions computable by such a protocol with the class P seems both interesting and difficult.

Avalanches. Let each node of a digraph represent a site where snow is accumulating. One special node s is considered the “outside world”. Once the amount of snow on a site (other than s) surpasses a given threshold, the site can “break”, sending one unit of snow to each of its out-neighbors. This may result in overloading some of the children of the node, and then these nodes break etc. If the digraph is strongly connected (which we assume for simplicity) then after a finite number of steps, no node will have too much snow (except s , which cannot break), and the avalanche terminates.

To maintain the dynamics of the model, snow is supposed to fall on the nodes. There are various ways to model this; simplest of these is to assume that each node i gets a given a_i amount of snow in unit time. We add snow until some node reaches the breaking threshold and starts an avalanche again (which happens so fast that no new snow falls during the avalanche).

The breaking threshold can be chosen, after some easy reductions, to be the outdegree of the node; then the avalanche is just a chip-firing game (where s is not allowed to be fired). But we can also include snow-fall in this model: we connect s to node i by a_i edges. Then a snowfall just corresponds to firing node s . We assume that there is enough snow in s (all those oceans, snow-caps etc) so that it can always be fired.

Thus a sequence of avalanche–snowfall–avalanche–snowfall–... is just an infinite chip-firing game on the graph, with the additional restriction that the special node s is only fired if no other node can be fired. We call such a restricted chip-firing game an *avalanche game*. When an avalanche starts, it consists of a sequence of firings which may happen in many ways, but the length of the avalanche, the number of times a node is fired, as well as the ending position are uniquely determined. The ending position of an avalanche is called *stable*.

Consider a periodic avalanche game. The (stable) position immediately before snowfall is called a *recurrent* position. A snowfall followed by an avalanche leads to another recurrent position, and this defines a permutation of recurrent positions. Each cycle in this permutation corresponds to a periodic avalanche game. The score vector of this game is an integer multiple of the period vector v . It follows by an argument almost identical to the

second half of the proof of Theorem 6.9 that in fact we get the period vector. Hence the number of recurrent positions in the cycle is v_s . It follows that the average length of an avalanche is

$$\frac{1}{v_s} \sum_{i \neq s} v_i,$$

independently of the cycle.

The conservation equation is very useful in the study of recurrent configurations. Gabrielov ([31]) introduces the lattice

$$\mathcal{L} = \{L^\top u : u \in \mathbb{Z}^V, u_s = 0\}$$

The conservation equation implies that if a position p can be obtained from a position q by a sequence of firings of nodes different from s , then $p - q \in \mathcal{L}$. It is not difficult to prove that if two positions p and q satisfy $p - q \in \mathcal{L}$ and p_i, q_i are large enough for all $i \neq s$, then there is a position that can be reached from each of p and q . Hence the stable positions in which the avalanches starting from p and q end are the same. It is also easy to see that this position is recurrent. These considerations imply that every translated copy $u + \mathcal{L}$ of the lattice \mathcal{L} (with $u \in \mathbb{Z}^V$) contains a unique recurrent position. Thus the number of recurrent positions is the same as the number of cosets of \mathcal{L} in \mathbb{Z}^V , which is $\det(\mathcal{L})$. Hence an easy computation gives the following interesting theorem of Dhar ([20]):

Theorem 6.10 *The number of recurrent positions is $\det(L')$, where L' is the matrix obtained from L by deleting the row and column corresponding to the node s .*

The reader may recognize that this determinant is just the number of spanning arborescences of G rooted at s , by the ‘‘Matrix-Tree Theorem’’ of Tutte. This relation is explained and exploited in [30], [31].

There are many characterizations of recurrent positions. For example, a position p is recurrent if and only if there exists a position q with $p_i \leq q_i$ for each node $i \neq s$ such that the avalanche starting from q ends with p .

Speer [46] gives a characterization that gives a way to test for recurrence. To describe this, we introduce a version of the period vector. We say that a vector $v \in \mathbb{Z}_+^V$ is *reducing*, if $v_s = 0$ and starting with N chips on each node (where N is a large integer), and firing each node i v_i times, we obtain a position with at most N chips on each node $i \neq s$. It is easy to see that a reducing vector must satisfy $v_i > 0$ for $i \neq s$. So we may fire every node once right away. This may produce a position with more than N chips on some node; this node must be fired at least twice during the game, so we may

as well fire it right away, etc. This way we construct a “canonical” reducing vector \hat{v} such that $\hat{v} \leq v$ for every reducing vector v .

Now one can prove that following analogue of Lemma 6.5:

Lemma 6.11 *Let $\alpha \in \mathcal{L}_p$, and assume that α does not contain s . Let α' be obtained by deleting the first \hat{v}_i occurrences of node i from α (if i occurs fewer than v_i times, we delete all of its occurrences). Then $\alpha' \in \mathcal{L}_p$.*

Corollary 6.12 *A stable configuration p is recurrent if and only if the avalanche starting from $p - L^\top \hat{v}$ ends with p .*

References

- [1] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász and C.W. Rackoff, Random walks, universal travelling sequences, and the complexity of maze problems, *Proc. 20th Ann. Symp. on Foundations of Computer Science* (1979), 218–223.
- [2] W. Aiello, B. Awerbuch, B. Maggs and S. Rao, Approximate load balancing on dynamic and asynchronous networks, *Proc. 25th ACM Symp. of Theory of Computing* (1993), 632–634.
- [3] N. Alon, Eigenvalues and expanders, *Combinatorica* **6** (1986), 83–96.
- [4] R.J. Anderson, L. Lovász, P.W. Shor, J. Spencer, É. Tardos and S. Winograd, Disks, balls, and walls: analysis of a combinatorial game, *Amer. Math. Monthly* **96** (1989), 481–493.
- [5] D.J. Aldous, *Reversible Markov Chains and Random Walks on Graphs* (book), to appear.
- [6] D.J. Aldous, Some inequalities for reversible Markov chains, *J. London Math. Soc.* **25** (1982), 564–576.
- [7] D.J. Aldous, On simulating a Markov chain stationary distribution when transition probabilities are unknown, preprint (1993).
- [8] S. Asmussen, *Applied Probability and Queues*, Wiley, New York 1987.
- [9] S. Asmussen, P. W. Glynn and H. Thorisson, Stationary detection in the initial transient problem, *ACM Transactions on Modeling and Computer Simulation* **2** (1992), 130–157.
- [10] J.R. Baxter and R.V. Chacon, Stopping times for recurrent Markov processes, *Illinois J. Math.* **20** (1976), 467–475.

- [11] D. Bayer and P. Diaconis, Trailing the dovetail shuffle to its lair, *Ann. Appl. Probab.* **2** (1992), 294–313.
- [12] P. Bak, C. Tang and K. Wiesenfeld, Self-organized criticality, *Physical Revue A* **38** (1988), 364–374.
- [13] A. Björner, L. Lovász and P. Shor, Chip-firing games on graphs, *Europ. J. Comb.* **12** (1991), 283–291.
- [14] A. Björner and L. Lovász, Chip-firing games on directed graphs, *J. Algebraic Combinatorics* **1** (1992), 305–328.
- [15] G. Brightwell and P. Winkler, Maximum hitting time for random walks on graphs, *J. Random Structures and Algorithms* **1** (1990), 263–276.
- [16] R.V. Chacon and D.S. Ornstein, A general ergodic theorem, *Illinois J. Math.* **4** (1960), 153–160.
- [17] F.R.K. Chung and S.-T. Yau, Eigenvalues of graphs and Sobolev inequalities, to appear.
- [18] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky and P. Tiwari, The electrical resistance of a graph captures its commute and cover times, *Proc. 21st ACM Annual ACM Symposium on the Theory of Computing* (1989), 574–586.
- [19] D. Coppersmith, P. Tetali and P. Winkler, Collisions among Random Walks on a Graph, *SIAM J. on Discrete Mathematics* **6** No. 3 (1993), 363–374.
- [20] D. Dhar, Self-organized critical state of sandpile automaton models, *Physical Revue Letters* **64** (1990), 1613–1616.
- [21] P. Diaconis, *Group Representations in Probability and Statistics*, Inst. of Math. Statistics, Hayward, CA (1988).
- [22] P. Diaconis and D. Stroock, Geometric bounds for eigenvalues of Markov chains, *Annals of Appl. Prob.* **1** (1991), 36–62.
- [23] J. Dodziuk and W.S. Kendall, Combinatorial Laplacians and isoperimetric inequality, in: *From Local Times to Global Geometry, Control and Physics*, (ed. K. D. Ellworthy), Pitman Res. Notes in Math. Series **150** (1986), 68–74.
- [24] P.G. Doyle and J.L. Snell, *Random Walks and Electric Networks*, Mathematical Assoc. of America, Washington, DC 1984.

- [25] M. Dyer, A. Frieze and R. Kannan, A random polynomial time algorithm for estimating volumes of convex bodies, *Proc. 21st Annual ACM Symposium on the Theory of Computing* (1989), 375–381.
- [26] A. Engel, The probabilistic abacus, *Educ. Stud. in Math.* **6** (1975), 1–22.
- [27] A. Engel, Why does the probabilistic abacus work? *Educ. Stud. in Math.* **7** (1976), 59–69.
- [28] K. Eriksson, No polynomial bound for the chip-firing game on directed graphs, *Proc. Amer. Math. Soc.* **112** (1991), 1203–1205.
- [29] J.A. Fill, Eigenvalue bounds on convergence to stationary for nonreversible Markov chains, with an application to the exclusion process, *Ann. of Appl. Prob.* **1** (1991), 62–87.
- [30] A. Gabriellov, Avalanches, sandpiles, and Tutte decomposition for directed graphs, preprint (1993).
- [31] A. Gabriellov, Asymmetric abelian avalanches and sandpiles, preprint (1993).
- [32] B. Ghosh and S. Muthukrishnan, Dynamic load balancing by random matchings, *J. Comp. Sys. Sci.*, to appear.
- [33] M.R. Jerrum and A. Sinclair, Approximating the permanent, *SIAM J. Comput.* **18** (1989), 1149–1178.
- [34] B. Korte, L. Lovász and R. Schrader, *Greedoids*, Springer, 1991.
- [35] L. Lovász, Random walks on graphs: a survey, in: *Combinatorics, Paul Erdős is Eighty*, (eds. D. Miklós, V.T.Sós and T. Szőnyi) J. Bolyai Math. Soc., Vol. II, to appear.
- [36] L. Lovász and M. Simonovits, Random walks in a convex body and an improved volume algorithm, *Random Structures and Alg.* **4** (1993), 359–412.
- [37] L. Lovász and P. Winkler, A note on the last new vertex visited by a random walk, *J. Graph Theory* **17** (1993), 593–596.
- [38] L. Lovász and P. Winkler, Exact mixing in an unknown Markov chain, preprint (1994).
- [39] L. Lovász and P. Winkler, Fast mixing in a Markov chain, preprint (1994).

- [40] L. Lovász and P. Winkler, The forget time of a Markov chain, preprint (1994).
- [41] M. Mihail, Conductance and convergence of Markov chains: a combinatorial treatment of expanders, *Proc. 30th Ann. Symp. on Foundations of Computer Science* (1989), 526–531.
- [42] J.W. Pitman, Occupation measures for Markov chains, *Adv. Appl. Prob.* **9** (1977), 69-86.
- [43] W. Reisig, *Petri Nets: An Introduction*, Springer-Verlag, New York 1985.
- [44] A. Skorokhod, *Studies in the Theory of Random Processes*, orig. pub. Addison-Wesley (1965), 2nd ed. Dover, New York 1982.
- [45] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability and Computing* **1** (1992), 351–370.
- [46] E. R. Speer, Asymmetric Abelian sandpile models, *J. Stat. Phys.* **71** (1993), 61–74.
- [47] J. Spencer, Balancing vectors in the max norm, *Combinatorica* **6** (1986), 55–66.
- [48] G. Tardos, Polynomial bound for a chip firing game on graphs *SIAM J. Disc. Math* **1** (1988), 397–398.