Chapter 5

Matroids

Let E be a finite ground set and $\mathcal{B} \subseteq 2^E$ a collection of subsets of E. We say that $\mathcal{B}$ forms the collection of *bases* of a *matroid* $M = (E, \mathcal{B})$ if the following two conditions hold:

1. All bases (sets in $\mathcal{B}$) have the same size, namely the *rank* of M.

2. For every pair of bases $X, Y \in \mathcal{B}$ and every element $e \in X$, there exists an element $f \in Y$ such that $X \cup \{f\} \setminus \{e\} \in \mathcal{B}$.

The above axioms for a matroid capture the notion of linear independence. Thus if $S = \{u_0, \dots, u_{m-1}\}$ is a set of n-vectors over a field K, then the maximal linearly independent subsets of S form the bases of a matroid with ground set S. The bases in this instance have size equal to the dimension of the vector space spanned by S, and they clearly satisfy the second or "exchange" axiom. A matroid that arises in this way is *vectorial*, and is said to be *representable over* K. A matroid that is representable over every field is called *regular*. Several other equivalent axiomatisations of matroid are possible, each shedding different light on the notion of linear independence; the above choice turns out to be the most appropriate for our needs. For other possible axiomatisations, and more on matroid theory generally, consult Oxley [**?**] or Welsh [**?**].

The advantage of the abstract viewpoint provided by matroid theory is that it allows us to perceive and exploit formal linear independence in a variety of combinatorial situations. Most importantly, the spanning trees in an unlabelled graph $G = (V, E)$ form the bases of a matroid, the *cycle matroid of* G, with ground set E. A matroid that arises as the cycle matroid of some graph is called *graphic*. The *co-cycle matroid of* G again has ground set E but the bases are now complements (in E) of spanning trees. The relationship of the cycle and co-cycle matroids of G is a special case of a general one of *duality*. All graphic matroids are regular, but the converse does not hold: the co-graphic matroid of a non-planar graph is regular but not graphic. A rather trivial class of matroids we shall encounter are the "uniform matroids." The *uniform matroid* $U_{r,m}$ of rank r on a ground set E of size m has as its bases all subsets of E or size r.

Two absolutely central operations on matroids are contraction and deletion. If $e \in E$ is an element of the ground set of M then the matroid $M \setminus e$ obtained by *deleting* e has ground set $E_{-e} = E \setminus \{e\}$ and bases $\mathcal{B}(M \setminus e) = \{X \subseteq E_{-e} : X \in \mathcal{B}(M)\}$; the matroid M/e obtained by *contracting* e has ground set E_{-e} and bases $\mathcal{B}(M/e) = \{X \subseteq E_{-e} : X \cup \{e\} \in \mathcal{B}(M)\}$. Any matroid obtained from M by a series of contractions and deletions is a *minor* of M.

The matroid axioms given above suggest a very natural walk on the set of bases of a matroid M. The *bases-exchange graph* $G(M)$ of a matroid M has vertex set $\mathcal{B}(M)$ and edge set

$$\big\{\{X,Y\} : X, Y \in \mathcal{B} \text{ and } |X \oplus Y| = 2\big\},$$

where $\oplus$ denotes symmetric difference. Note that the edges of the bases-exchange graph $G(M)$ correspond to the transformations guaranteed by the exchange axiom. Indeed, it is straightforward to check, using the exchange axiom, that the graph $G(M)$ is always connected. By simulating a random walk on $G(M)$ it is possible, in principle, to sample a base (almost) u.a.r. from $\mathcal{B}(M)$. Although it has been conjectured that the random walk on $G(M)$ is rapidly mixing for all matroids M, the conjecture has never been proved and the circumstantial evidence in its favour seems slight. Nevertheless there is an interesting class of matroids, the "balanced" matroids for which rapid mixing has been established. The definition of balanced matroid is due to Feder and Mihail [**?**], as is the proof of rapid mixing. We follow their treatment quite closely.

5.1 Balanced matroids

For this section we usually drop explicit reference to the matroid M, and simply write $\mathcal{B}$ and E in place of $\mathcal{B}(M)$ and E. Suppose a base $X \in \mathcal{B}$ is chosen u.a.r. If $e \in E$, we let e stand (with a slight abuse of notation) for the event $e \in X$, and $\bar{e}$ for the event $e \notin X$. Furthermore, we denote conjunction of events by juxtaposition: thus $e\bar{f}$ denotes the event $e \in X \wedge f \notin X$, etc. The matroid M is said to possess the *negative correlation property* if the inequality $\mathbf{Pr}(ef) \le \mathbf{Pr}(e)\mathbf{Pr}(f)$ holds for all pairs of distinct elements $e, f \in E$. Another way of expressing negative correlation is by writing $\mathbf{Pr}(e \mid f) \le \mathbf{Pr}(e)$; in other words the knowledge that f is present in X makes the presence of e less likely.[1] Further, the matroid M is said to be *balanced* if all minors of M (including M itself) possess the negative correlation property. We shall see in §5.1.2 that regular matroids are always balanced. But there are balanced matroids that are not regular: it is easy to check that all uniform matroids satisfy the negative correlation property and that the

[1]We assume here that $\mathbf{Pr}(f) > 0$; an element f such that $\mathbf{Pr}(f) = 0$ is said to be a *loop*.

class of uniform matroids is closed under contraction and deletion; on the other hand, $U_{2,m}$ is not regular when $m \geq 4$. (Refer to Oxley [?, Theorem 13.1.1].)

5.1.1 Efficiently sampling bases of balanced matroids

It is convenient in this section to work with a combinatorial version of conductance rather than conductance itself. The *cutset expansion* of a graph G is the minimum, over all subsets $\mathcal{S} \subset V(G)$ with $0 < |\mathcal{S}| \leq |V(G)|/2$ of the ratio $|\operatorname{cut}(\mathcal{S})|/|\mathcal{S}|$, where $\operatorname{cut}(\mathcal{S}) \subseteq E(G)$ denotes the set of edges with one endpoint in $\mathcal{S}$ and one in the complement of S. The main result of this section is a lower bound on cutset expansion of the bases-exchange graph.

Theorem 5.1.1 *The cutset expansion of the bases-exchange graph $G(M)$ of any balanced matroid M is at least* 1.

Suppose we implement the random walk on the bases-exchange graph $G(M)$ in the following natural way. The current state (base) is X.

Step 1 With probability $\frac{1}{2}$ set $Y = X$.

Step 2 Otherwise, choose e u.a.r. from $E \setminus X$.

Step 3 Choose $f \in E$ u.a.r. from the elements of the ground set satisfying $Y = X \cup \{e\} \setminus \{f\} \in \mathcal{B}$.

The new state is Y. Call this the *bases-exchange walk*. Note that the transition matrix implicitly described by the above implementation is symmetric. Since we have already observed that $G(M)$ is connected, we see that the bases-exchange walk converges to a stationary distribution that is uniform over states. Furthermore, the non-zero transition probabilities (corresponding to edges of $G(M)$) are all at least $1/2mr$. Thus, according to Theorem 5.1.1, the conductance of the random walk is bounded below by $1/2mr$, and by Theorem 2.2.1 we obtain:

Corollary 5.1.1 *The mixing time of the bases-exchange walk on any balanced matroid of rank r on a ground set of size m is at most* $4m^2r^2(r \ln m + \ln \epsilon^{-1})$.

We approach the proof of Theorem 5.1.1 via a couple of lemmas. If $E' \subseteq E$, then a *increasing property*over E' is a property of subsets of E' that is closed under the superset relation; equivalently, it is a property that may be expressed as a monotone Boolean formula in the indicator variables of the elements in E'. A *decreasing property* is defined analogously.

Lemma 5.1.1 *Let M be a balanced matroid and let $e \in E$.*

(a) *If μ is an increasing property over E_{-e}, then $\mathbf{Pr}(\mu \mid e) \leq \mathbf{Pr}(\mu \mid \bar{e})$.*

(b) *If μ is a decreasing property over E_{-e}, then $\mathbf{Pr}(\mu \mid e) \geq \mathbf{Pr}(\mu \mid \bar{e})$.*

Proof We prove (a), (b) follows by consideration of $\bar{\mu}$. The proof is by induction on the size of the ground set. We may assume that $\mathbf{Pr}(\mu e) > 0$, otherwise the result is immediate. Conditional probabilities with respect to e and μe are thus well defined, and we may re-express our goal as $\mathbf{Pr}(\mu \mid e) \leq \mathbf{Pr}(\mu)$. If the rank of M is 1 then either (i) $\emptyset \in \mu$ and $\mathbf{Pr}(\mu) = 1$ or (ii) $\emptyset \notin \mu$ and $\mathbf{Pr}(\mu e) = 0$. Thus we may assume that the rank r of M is at least 2.

From the identity

$$\mathbf{E}(|X \setminus e| \mid \mu e) = \sum_{f \neq e} \mathbf{Pr}(f \mid \mu e) = r - 1 = \mathbf{E}(|X \setminus e| \mid e) = \sum_{f \neq e} \mathbf{Pr}(f \mid e),$$

and the assumption that $r \geq 2$, we deduce the existence of an element f satisfying $\mathbf{Pr}(f \mid \mu e) \geq \mathbf{Pr}(f \mid e) > 0$, and hence

$$\mathbf{Pr}(\mu \mid ef) \geq \mathbf{Pr}(\mu \mid e); \tag{5.1}$$

note that the conditional probability on the left is well defined. Two further inequalities that hold between conditional probabilities are

$$\mathbf{Pr}(f \mid e) \leq \mathbf{Pr}(f) \tag{5.2}$$

and

$$\mathbf{Pr}(\mu \mid ef) \leq \mathbf{Pr}(\mu \mid f); \tag{5.3}$$

the former comes simply from the negative correlation property, and the latter from applying the inductive hypothesis to the matroid M/f and the property derived from μ by forcing f to 1.

At this point we dispense with the degenerate case $\mathbf{Pr}(\bar{f} \mid e) = 0$. It follows from (5.2) that $\mathbf{Pr}(f) = 1$, and then from (5.3) that $\mathbf{Pr}(\mu \mid e) \leq \mathbf{Pr}(\mu)$, as desired. So we may now assume $\mathbf{Pr}(\bar{f} \mid e) > 0$ and hence that probabilities conditional on the event $e\bar{f}$ are well defined. In particular,

$$\mathbf{Pr}(\mu \mid e\bar{f}) \leq \mathbf{Pr}(\mu \mid \bar{f}), \tag{5.4}$$

as can be seen by applying the inductive hypothesis to the matroid $M \setminus f$ and the property derived from μ by forcing f to 0. Further, inequality (5.1) may be re-expressed as

$$\mathbf{Pr}(\mu \mid ef) \geq \mathbf{Pr}(\mu \mid e\bar{f}). \tag{5.5}$$

The inductive step is now achieved through a chain of inequalities based on (5.2)–(5.5):

$$\begin{aligned}
\mathbf{Pr}(\mu \mid e) &= \mathbf{Pr}(\mu \mid ef)\mathbf{Pr}(f \mid e) + \mathbf{Pr}(\mu \mid e\bar{f})\mathbf{Pr}(\bar{f} \mid e) \\
&= \mathbf{Pr}(\mu \mid ef)\mathbf{Pr}(f \mid e) + \mathbf{Pr}(\mu \mid e\bar{f})(1 - \mathbf{Pr}(f \mid e)) \\
&= \big[\mathbf{Pr}(\mu \mid ef) - \mathbf{Pr}(\mu \mid e\bar{f})\big]\mathbf{Pr}(f \mid e) + \mathbf{Pr}(\mu \mid e\bar{f}) \\
&\le \big[\mathbf{Pr}(\mu \mid ef) - \mathbf{Pr}(\mu \mid e\bar{f})\big]\mathbf{Pr}(f) + \mathbf{Pr}(\mu \mid e\bar{f}) \qquad (5.6) \\
&= \mathbf{Pr}(\mu \mid ef)\mathbf{Pr}(f) + \mathbf{Pr}(\mu \mid e\bar{f})\mathbf{Pr}(\bar{f}) \\
&\le \mathbf{Pr}(\mu \mid f)\mathbf{Pr}(f) + \mathbf{Pr}(\mu \mid \bar{f})\mathbf{Pr}(\bar{f}) \qquad (5.7) \\
&= \mathbf{Pr}(\mu),
\end{aligned}$$

where inequality (5.6) uses (5.2) and (5.5), and inequality (5.7) uses (5.3) and (5.4). □

Given $e \in E$, the set of bases $\mathcal{B}$ may be partitioned as $\mathcal{B} = \mathcal{B}_e \cup \mathcal{B}_{\bar{e}}$, where $\mathcal{B}_e = \{X \in \mathcal{B} : e \in X\}$ and $\mathcal{B}_{\bar{e}} = \{X \in \mathcal{B} : e \notin X\}$; observe that $\mathcal{B}_e$ and $\mathcal{B}_{\bar{e}}$ are isomorphic to $\mathcal{B}(M/e)$ and $\mathcal{B}(M \setminus e)$, respectively. For $\mathcal{A} \subseteq \mathcal{B}_e$ (respectively, $\mathcal{A} \subseteq \mathcal{B}_{\bar{e}}$), let $\Gamma_e(\mathcal{A})$ denote the set of all vertices in $\mathcal{B}_{\bar{e}}$ (respectively, $\mathcal{B}_e$) that are adjacent to some vertex in $\mathcal{A}$. The bipartite subgraph of the bases-exchange graph induced by the bipartition $\mathcal{B} = \mathcal{B}_e \cup \mathcal{B}_{\bar{e}}$ satisfies a natural expansion property. For $\mathcal{S} \subseteq \mathcal{B}$ we let $\mathcal{S}_e = \{X \in \mathcal{S} : e \in X\}$ and $\mathcal{S}_{\bar{e}} = \mathcal{S} \setminus \mathcal{S}_e$.

Lemma 5.1.2 *Suppose M is a balanced matroid, $e \in E$, and that the partition $\mathcal{B} = \mathcal{B}_e \cup \mathcal{B}_{\bar{e}}$ is non-trivial. Then for all $\mathcal{S} \subseteq \mathcal{B}$,*

$$\frac{|\Gamma_e(\mathcal{S}_e)|}{|\mathcal{B}_{\bar{e}}|} \ge \frac{|\mathcal{S}_e|}{|\mathcal{B}_e|}, \quad \textit{and}$$
$$\frac{|\Gamma_e(\mathcal{S}_{\bar{e}})|}{|\mathcal{B}_e|} \ge \frac{|\mathcal{S}_{\bar{e}}|}{|\mathcal{B}_{\bar{e}}|}.$$

Proof $\mu_1 = \{Y \subseteq E_{-e} : \exists X \in \mathcal{S}_e \text{ s.t. } Y \supseteq X \setminus \{e\}\}$ is an increasing property. The collection of all bases in $\mathcal{B}_e$ satisfying μ_1 is precisely $\mathcal{S}_e$, while the collection of all bases in $\mathcal{B}_{\bar{e}}$ satisfying μ_1 is precisely $\Gamma_e(\mathcal{S}_e)$. Hence the first part of the lemma is equivalent to the inequality $\mathbf{Pr}(\mu_1 \mid \bar{e}) \ge \mathbf{Pr}(\mu_1 \mid e)$, which follows from Lemma 5.1.1. Similarly, $\mu_2 = \{Y \subseteq E_{-e} : \exists X \in \mathcal{S}_{\bar{e}} \text{ s.t. } Y \subseteq X \cup \{e\}\}$ is a decreasing property. The set of all bases in $\mathcal{B}_{\bar{e}}$ satisfying μ_2 is precisely $\mathcal{S}_{\bar{e}}$, while the set of all bases in $\mathcal{B}_e$ satisfying μ_2 is precisely $\Gamma_e(\mathcal{S}_{\bar{e}})$. Hence the second part of the lemma is equivalent to the inequality $\mathbf{Pr}(\mu_2 \mid e) \ge \mathbf{Pr}(\mu_2 \mid \bar{e})$, which again follows from Lemma 5.1.1. □

We now have the tools needed to bound the cutset expansion of the bases-exchange graph.

Proof of Theorem 5.1.1 We proceed by induction on the rank of M. Let $\mathcal{S} \subset \mathcal{B}$ be a collection of bases, with $|\mathcal{S}| \le |\mathcal{B}|/2$, defining a cut in the bases-exchange graph of M. Let $\mathcal{S}_e = \mathcal{S} \cap \mathcal{B}_e$ and $\mathcal{S}_{\bar{e}} = \mathcal{S} \cap \mathcal{B}_{\bar{e}}$, and define α and β by $|\mathcal{S}_e| = \alpha|\mathcal{B}_e|$ and

$|\mathcal{S}_{\bar{e}}| = \beta|\mathcal{B}_{\bar{e}}|$. The edges forming the cut are of three kinds: (i) those whose endpoints are both within $\mathcal{B}_e$, (ii) those whose endpoints are both within $\mathcal{B}_{\bar{e}}$, and (iii) those which span $\mathcal{B}_e$ and $\mathcal{B}_{\bar{e}}$. By the induction hypothesis, the numbers of edges of kinds (i) and (ii) are at least $\min\{\alpha, 1-\alpha\}|\mathcal{B}_e|$ and $\min\{\beta, 1-\beta\}|\mathcal{B}_{\bar{e}}|$, respectively. To lower bound the number of edges of kind (iii), assume first that $\alpha \geq \beta$. By Lemma 5.1.2, there are at least $\alpha|\mathcal{B}_{\bar{e}}|$ bases in $\mathcal{B}_{\bar{e}}$ adjacent to some base in $\mathcal{S}_e$; of these, at least $(\alpha-\beta)|\mathcal{B}_{\bar{e}}|$ must lie outside $\mathcal{S}_{\bar{e}}$. Thus there are at least $(\alpha-\beta)|\mathcal{B}_{\bar{e}}|$ edges of type (iii). This argument can equally well be applied in the opposite direction, starting at the set $\mathcal{B}_{\bar{e}} \setminus \mathcal{S}_{\bar{e}}$, yielding a second lower bound of $((1-\beta)-(1-\alpha))|\mathcal{B}_e| = (\alpha-\beta)|\mathcal{B}_e|$. Thus the number of edges of kind (iii) is at least $(\alpha-\beta)\max\{|\mathcal{B}_e|, |\mathcal{B}_{\bar{e}}|\}$. Since the case $\alpha < \beta$ is entirely symmetric, we obtain, summing the contributions from edges of kinds (i)–(iii):

$$|\operatorname{cut}(\mathcal{S})| \geq \min\{\alpha, 1-\alpha\}|\mathcal{B}_e| + \min\{\beta, 1-\beta\}|\mathcal{B}_{\bar{e}}| + |\alpha-\beta|\max\{|\mathcal{B}_e|, |\mathcal{B}_{\bar{e}}|\}. \tag{5.8}$$

To complete the proof we must show that $|\operatorname{cut}(\mathcal{S})|$ is always at least $\alpha|\mathcal{B}_e| + \beta|\mathcal{B}_{\bar{e}}| = |\mathcal{S}|$, whenever $|\mathcal{S}| \leq |\mathcal{B}|/2$. Note that this last condition may be expressed as

$$(\tfrac{1}{2}-\alpha)|\mathcal{B}_e| + (\tfrac{1}{2}-\beta)|\mathcal{B}_{\bar{e}}| \geq 0. \tag{5.9}$$

If $\alpha, \beta \leq \frac{1}{2}$, the required lower bound on $|\operatorname{cut}(\mathcal{S})|$ follows immediately from (5.8). We therefore just need to treat the cases when one of α or β is greater than $\frac{1}{2}$. To simplify the working, we'll exploit the symmetry of (5.8) and assume, without loss of generality, that

$$|\mathcal{B}_e| \geq |\mathcal{B}_{\bar{e}}|. \tag{5.10}$$

Suppose first that $\alpha > \frac{1}{2}$. Then inequalities (5.9) and (5.10) entail $\beta < 1-\alpha < \frac{1}{2}$, and inequality (5.8) simplifies to

$$|\operatorname{cut}(\mathcal{S})| \geq (1-\alpha)|\mathcal{B}_e| + \beta|\mathcal{B}_{\bar{e}}| + (\alpha-\beta)|\mathcal{B}_e|.$$

Hence,

$$|\operatorname{cut}(\mathcal{S})| \geq (1-\beta)|\mathcal{B}_e| + \beta|\mathcal{B}_{\bar{e}}| \geq \alpha|\mathcal{B}_e| + \beta|\mathcal{B}_{\bar{e}}| = |\mathcal{S}|,$$

as required.

Finally, suppose that $\beta > \frac{1}{2}$. Then necessarily $\alpha < \frac{1}{2}$ and inequality (5.8) simplifies to

$$\begin{aligned}|\operatorname{cut}(\mathcal{S})| \geq \alpha|\mathcal{B}_e| + (1-\beta)|\mathcal{B}_{\bar{e}}| + (\beta-\alpha)|\mathcal{B}_e| &\geq \beta|\mathcal{B}_e| + (1-\beta)|\mathcal{B}_{\bar{e}}| \\ &\geq (\alpha+\beta-\tfrac{1}{2})|\mathcal{B}_e| + \tfrac{1}{2}|\mathcal{B}_{\bar{e}}| \geq \alpha|\mathcal{B}_e| + \beta|\mathcal{B}_{\bar{e}}| = |\mathcal{S}|.\end{aligned}$$

This completes the inductive step. □

5.1.2 Regular matroids are balanced

A natural question now presents itself: how big is the class of balanced matroids? Recall that a regular matroid is one that is representable over every field. In this section we prove that all regular matroids are balanced. More precisely, we prove the equivalent result that all "orientable" matroids are balanced. The class of orientable matroids is known to be the same as the class of regular matroids [**?**, Corollary 13.4.6].[2]

In order to define the property of being orientable, we need some further matroid terminology. A *cycle* $C \subset E$ in a matroid $M = (E, \mathcal{B})$ is a minimal (under set inclusion) subset of elements that cannot be extended to a base. A *cut* is a minimal set of elements whose complement cannot be extended to a base. Note that in the case of the cycle matroid of a graph, in which the bases are spanning trees, these terms are consistent with the usual graph theoretic ones. Let $\mathcal{C} \subseteq 2^E$ denote the set of all cycles in M and $\mathcal{D} \subseteq 2^E$ the set of all cuts. We say that M is *orientable* if functions $\gamma : \mathcal{C} \times E \to \{-1, 0, +1\}$ and $\delta : \mathcal{D} \times E \to \{-1, 0, +1\}$ exist which satisfy the following three conditions, for all $C \in \mathcal{C}$ and $D \in \mathcal{D}$:

$$\begin{aligned} \gamma(C, g) &\neq 0 \text{ iff } g \in C, \\ \delta(D, g) &\neq 0 \text{ iff } g \in D, \text{ and} \\ \sum_{g \in E} \gamma(C, g)\delta(D, g) &= 0. \end{aligned} \tag{5.11}$$

We work in this section towards the following result.

Theorem 5.1.2 *Orientable (and hence regular) matroids are balanced.*

A *near base* of M is a set $N \subseteq E$ that can be augmented to a base by the addition of a single element from the ground set. A *unicycle* of M is a set $U \subseteq E$ that can be reduced to a base by the removal of a single element. A near base N defines a unique cut D_N consisting of all elements of the ground set whose addition to N results in a base. A unicycle U defines a unique cycle C_U consisting of all elements which whose removal from U results in a base. Let e, f be distinct elements of the ground set E. We claim that

$$\gamma(C_U, e)\gamma(C_U, f) + \delta(D_N, e)\delta(D_N, f) = 0, \tag{5.12}$$

for all near-bases N and unicycles U that are related by $U = N \cup \{e, f\}$. To see this, note that the equation (5.11) simplifies in this situation to

$$\gamma(C_U, e)\delta(D_N, e) + \gamma(C_U, f)\delta(D_N, f) = 0, \tag{5.13}$$

[2] When consulting this corollary, it is important to realise that Oxley applies the term "signable" to the class of matroids Feder and Mihail call "orientable," preferring to apply the latter term to a different and larger class. We follow Feder and Mihail's terminology.

since all terms in the sum are zero except from those obtained by setting $g = e$ and $g = f$. Now it may be that all four quantities in (5.13) are zero, in which case we are done. Otherwise, some quantity, say $\delta(D_N, e)$, is non-zero, in which case $D_N \cup \{e\} = C_U \setminus \{f\}$ is a base and $\gamma(C_U, f)$ is non-zero also. Multiplying (5.13) through by $\gamma(C_U, f)\delta(D_N, e)$ yields

$$\gamma(C_U, e)\gamma(C_U, f)\delta(D_N, e)^2 + \gamma(C_U, f)^2\delta(D_N, e)\delta(D_N, f) = 0,$$

which simplifies to equation (5.12) as required, since the square factors are both one.

For distinct elements $e, f \in E$, define

$$\Delta_{ef} = \sum_N \delta(D_N, e)\delta(D_N, f) = -\sum_U \gamma(C_U, e)\gamma(C_U, f),$$

where the sums are over all near bases N and unicycles U. The equality of the two expressions above is a consequence of (5.12), and the bijection between non-zero terms in the two sums that is given by $N \mapsto N \cup \{e, f\} = U$. Select a distinguished element $e \in E$ and force $\gamma(C, e) = -1$ and $\delta(D, e) = 1$ for all cycles $C \ni e$ and cuts $D \ni e$. This can be done by flipping signs around cycles and cuts, without compromising the condition (5.11) for orientability, nor changing the value of Δ_{ef}. With this convention we have

$$\sum_{g \neq e} \gamma(C, g)\delta(D, g) = 1, \quad \text{provided } C \ni e \text{ and } D \ni e; \tag{5.14}$$

$$\gamma(C_U, f) = \delta(D_N, f), \quad \text{provided } U = N \cup \{e, f\}; \tag{5.15}$$

and

$$\Delta_{ef} = \sum_{U: e \in C_U} \gamma(C_U, f) = \sum_{N: e \in D_N} \delta(D_N, f), \tag{5.16}$$

where C, D, U and N denote, respectively, arbitrary cycles, cuts, unicycles and near bases satisfying the stated conditions. An intuitive reading of Δ_{ef} is as a measure of whether cycles containing e, f arising from unicycles tends to traverse e and f in the same or opposite directions; similarly for cuts arising from near bases.

We extend earlier notation in an obvious way, so that $\mathcal{B}_{ef}$ is the set of bases of M containing both e and f, and $\mathcal{B}_{\bar{e}f}$ is the set of bases excluding e but including f, etc.

Theorem 5.1.3 *The bases $\mathcal{B} = \mathcal{B}(M)$ of an oriented matroid M satisfy $|\mathcal{B}| \cdot |\mathcal{B}_{ef}| = |\mathcal{B}_e| \cdot |\mathcal{B}_f| - \Delta_{ef}^2$.*

Proof We consider a pair of bases $(X, Y) \in \mathcal{B}_{\bar{e}} \times \mathcal{B}_{ef}$ to be adjacent to a pair $(X', Y') \in \mathcal{B}_e \times \mathcal{B}_{\bar{e}f}$ if (X', Y') can be obtained by an exchange involving e and a second element $g \neq e$:

$$X' = X \cup \{e\} \setminus \{g\} \tag{5.17}$$

$$Y' = Y \cup \{g\} \setminus \{e\}. \tag{5.18}$$

With each adjacent pair we associate a weight

$$\gamma(C_{X\cup\{e\}}, g)\delta(D_{Y\setminus\{e\}}, g). \tag{5.19}$$

Given a pair $(X, Y) \in \mathcal{B}_{\bar{e}} \times \mathcal{B}_{ef}$, the condition that an exchange involving g leads to a valid pair of bases (X', Y') via (5.17) and (5.18) is precisely that the weight (5.19) is non-zero. Note that whenever this occurs, $(X', Y') \in \mathcal{B}_e \times \mathcal{B}_{\bar{e}f}$. Thus

$$\begin{aligned}
|\mathcal{B}_{\bar{e}}| \cdot |\mathcal{B}_{ef}| &= \sum_{(X,Y)\in\mathcal{B}_{\bar{e}}\times\mathcal{B}_{ef}} \left[\sum_{g\neq e} \gamma(C_{X\cup\{e\}}, g)\delta(D_{Y\setminus\{e\}}, g)\right] \\
&= W, \qquad (5.20)
\end{aligned}$$

where W is the total weight of adjacent pairs. Here we have used equation (5.14).

Now we perform a similar calculation, but in the other direction, starting at pairs $(X', Y') \in \mathcal{B}_e \times \mathcal{B}_{\bar{e}f}$. We apply a weight

$$\delta(D_{X'\setminus\{e\}}, g)\gamma(C_{Y'\cup\{e\}}, g) \tag{5.21}$$

to each adjacent pair, which is consistent, by (5.15), with the weight (5.19) applied earlier. Again, starting at (X', Y'), the condition that the pair (X, Y) obtained by inverting the exchange given in (5.17) and (5.18) is that the weight (5.21) in non-zero. But now, even if the weight is non-zero, there is a possibility that the new pair of bases (X, Y) will not be a member of $\mathcal{B}_{\bar{e}} \times \mathcal{B}_{ef}$; this will happen precisely when $g = f$. Thus

$$\begin{aligned}
|\mathcal{B}_e| \cdot |\mathcal{B}_{\bar{e}f}| &= \sum_{(X',Y')\in\mathcal{B}_e\times\mathcal{B}_{\bar{e}f}} \left[\sum_{g\neq e} \delta(D_{X'\setminus\{e\}}, g)\gamma(C_{Y'\cup\{e\}}, g)\right] \qquad (5.22) \\
&= \sum_{(X',Y')\in\mathcal{B}_e\times\mathcal{B}_{\bar{e}f}} \left[\sum_{g\neq e,f} \delta(D_{X'\setminus\{e\}}, g)\gamma(C_{Y'\cup\{e\}}, g)\right] \\
&\qquad + \sum_{(X',Y')\in\mathcal{B}_e\times\mathcal{B}_{\bar{e}f}} \delta(D_{X'\setminus\{e\}}, f)\gamma(C_{Y'\cup\{e\}}, f) \\
&= W + \sum_{(X',Y')\in\mathcal{B}_e\times\mathcal{B}_{\bar{e}}} \delta(D_{X'\setminus\{e\}}, f)\gamma(C_{Y'\cup\{e\}}, f) \qquad (5.23) \\
&= W + \sum_{X'\in\mathcal{B}_e} \delta(D_{X'\setminus\{e\}}, f) \sum_{Y'\in\mathcal{B}_{\bar{e}}} \gamma(C_{Y'\cup\{e\}}, f) \\
&= W + \Delta_{ef}^2. \qquad (5.24)
\end{aligned}$$

Here, step (5.22) is by (5.14); step (5.23) uses the observation that terms are non-zero only when $f \in Y'$; and (5.24) is from the definition (5.16) of Δ_{ef}.

Comparing 5.20 and 5.24 we have

$$|\mathcal{B}_e| \cdot |\mathcal{B}_{\bar{e}f}| = |\mathcal{B}_{\bar{e}}| \cdot |\mathcal{B}_{ef}| + \Delta_{ef}^2,$$

and the result now follows by adding $|\mathcal{B}_e| \cdot |\mathcal{B}_{ef}|$ to both sides. □

Proof of Theorem 5.1.2 According to Theorem 5.1.3, all orientable matroids satisfy the negative correlation property. Moreover, it is easily checked that the class of orientable matroids is closed under contraction and deletion. □

Remark 1: (Flesh this out.) Number of bases of a regular matroid may be computed exactly in time ? by matrix-tree theorem + Gaussian elimination. This gives alternative polynomial-time sampling procedure. However, as we have seen, the class of balanced matroids is strictly larger than the class of regular matroids.

Remark 2: (Flesh this out.) There exist non-balanced matroids. Let M be a matroid of rank r on ground set E. For any $0 < r' < r$,

$$\mathcal{B}' = \{X' : |X'| = r' \wedge \exists X \in \mathcal{B}(M).\, X' \subset X\}$$

is the collection of bases of a matroid M' on ground set E, the *truncation* of M to rank r'. The truncation of a graphic matroid may fail to be balanced. Consider the graph G with vertex set

$$\{u, v, y, z, 0, 1, 2, 3, 4\}$$

and edge set

$$\big\{\{u, v\}, \{y, z\}\big\} \cup \big\{\{u, i\} : 0 \le i \le 4\big\} \cup \big\{\{v, i\} : 0 \le i \le 4\big\}.$$

Let e denote the edge $\{u, v\}$ and f the edge $\{y, z\}$. Let $\mathcal{F}^6$ denote the set of forests in G with six edges, $\mathcal{F}^6_{ef}$ the number of such forests including edges e and f, etc. Then $\mathcal{F}^6_{ef} = 80$, $\mathcal{F}^6_{e\bar{f}} = 32$, $\mathcal{F}^6_{\bar{e}f} = 80$ and $\mathcal{F}^6_{\bar{e}\bar{f}} = 192$. Thus

$$\mathbf{Pr}(e \mid f) = 5/17 > 7/24 = \mathbf{Pr}(e),$$

contradicting negative correlation.

5.2 Graphic matroids in particular

Since graphic matroids are balanced, the bases-exchange walk may be used to sample, efficiently and almost u.a.r., spanning trees in an undirected graph. However there are a number of other procedures for sampling bases in this special case, some of them providing *exactly* uniform samples. Perhaps the most efficient proposal is the "cycle-popping" technique of Wilson. We describe this now in the somewhat more general setting of sampling a directed tree in a directed graph.

5.2.1 Cycle popping: the general setting

Let $G = (V, A, r)$ be a directed graph with vertices V, arcs A and a distinguished root r. A *directed tree with root* r in G is a subgraph (V, T) in which there is a unique path from each vertex $v \in V$ to the root r. Note that a tree (we drop the qualifier "directed" at this point) has $n-1$ arcs, and every vertex other than r has outdegree 1 (the root has outdegree 0). Thus another way of viewing a tree is as a function $f : V \setminus \{r\} \to V$ that is cycle free: that is, $f^i(v) = v$ entails $i = 0$, for all v, i such that $f^i(v)$ is defined. One way to sample a tree is to select u.a.r. a function $f : V \setminus \{r\} \to V$ and accept if it is cycle-free. However the rejection probability will in general be high, as can be seen by considering the $n \times n$ grid: there are $O(n^2)$ disjoint 4-cycles in the grid, and each of them will, independently with probability $1/256$, lead to a cycle in f. Thus the probability that f is cycle-free is exponentially small in the number of vertices. The idea behind the cycle-popping strategy is to remove cycles and re-randomise f on the affected vertices.

We first describe cycle-popping in a setting that is convenient for proof, but not for implementation. For each vertex $u \in V \setminus \{r\}$, we postulate a sequence $(S_u^0, S_u^1, S_u^2, \dots) \in \Gamma(u)^\omega$ of r.v's, where $\Gamma(u) = \{v : (u, v) \in A\}$ is the set of neighbours of u. Each r.v. S_u^i is distributed uniformly over $\Gamma(u)$, and is independent of all the other r.v's. We call the indices i "colours." At any instant there is a visible colour $c(u)$ at vertex u; initially, $c(u) = 0$ for all $u \in V \setminus \{r\}$. As time progresses, higher colours become visible, corresponding to r.v's further along the lists being revealed.

Consider the following procedure, guided by the r.v's (S_u^i). Let the currently visible colours be $c : V \setminus \{r\} \to \mathbf{N}$, and consider the function $f : V \setminus \{r\} \to V$ given by $f(u) = S_u^{c(u)}$ for all $u \in V \setminus \{r\}$. The digraph $D_f = (V, \{(v, f(v)) : v \in V \setminus \{r\}\}$ has the following structure. The week component containing r is a directed tree with root r. Every other weak component consists of a single cycle plus disjoint directed trees rooted at a vertex of the cycle. If f is cycle-free we are done. Otherwise, select an ℓ-cycle $C = (u, f(u), f^2(u), \dots, f^{\ell-1}(u))$ and "pop" it; that is, increment $c(v)$ for all vertices v on the cycle C, revealing a fresh set of colours/r.v's. This process, if iterated, might continue indefinitely, but if it terminates, f will define a tree in G with root r. We shall argue that the process terminates with probability 1, and that the tree produced is exactly uniform.

The cycle-popping process is nondeterministic, since a number of cycles may be available for popping at any instant. The key observation is that the order of popping does not matter: if the process terminates then it always terminates with the colour labelling c (and hence the same cycle-free function f). Consider any configuration of the process, uniquely determined by the colour assignment c. A number of cycles $C_0, \dots, C_{s-1}$ may be available for popping. Assume $s \geq 2$ and that C_j and C_k are distinct cycles. Necessarily, C_j and C_k are disjoint, so that if we decide to pop C_j first we can then pop C_k and be in exactly the same configuration as if we had popped C_k first and then C_j. Thus the process has the "diamond property" and hence is Church-Rosser:

either the process continues indefinitely, or terminates at a well defined configuration independent of the order in which cycles are popped. This is a result of Newman [**?**], see also Sperschneider and Antoniou [**?**].

In fact, more is true. Label each transition of the cycle-popping process by the *coloured* cycle—i.e., the sequence of vertices u on the cycle together with their corresponding colours $c(u)$—whose popping generates that transition. In each diamond, the same two coloured cycles are involved in the two paths through the diamond. Thus in any sequence of transitions leading to the unique terminating configuration (assuming it exists) exactly the same set of coloured cycles are popped, only the order of popping varies. Thus we can think of the cycle-popping process (assuming it terminates) as defining an underlying tree T, rooted at r, on which are superimposed a partially ordered set $\mathcal{C}$ of coloured cycles. Conditioned on the set $\mathcal{C}$ of cycles, the tree T is uniform. Thus, conditioned on termination, the cycle-popping process generates a rooted tree u.a.r. We collect these discoveries in the following theorem.

Theorem 5.2.1 *The order in which cycles are popped in the cycle-popping process is of no consequence: for a given collection of r.v's (S_u^i) the process either always continues indefinitely, or always terminates at the same configuration (colouring) c. Conditioned on termination, the tree defined by the colouring c is distributed uniformly.*

In fact, termination occurs with probability 1, but this is easier to appreciate once we move to an alternative, more implementation-friendly version of the process.

5.2.2 Cycle popping: the implementation

The process of the previous section is straightforwardly implementable, provided we view the r.v's (S_u^i) as being revealed to us on demand. We know from Theorem 5.2.1 that the order in which the cycles are popped is of no consequence. A particularly elegant way of performing the computations is to perform a random walk on G, popping cycles as soon as they are discovered. (Refer to Figure 5.1.) We shall refer to this particular implementation of cycle-popping as the *cycle-erased random walk*, even though it is more commonly called the *loop-erased random walk* in the literature. Note that by storing the current function f as an array *Tree*[ ], the effect of popping a cycle is achieved automatically through overwriting an array element.

Assume that G is strongly connected and contains at least one odd cycle, so that the random walk on G has a well defined stationary distribution π.

Perhaps we should relax this.

Theorem 5.2.2 *The procedure call* TreeSample(G) *halts with probability 1, returning a uniform random tree in G, rooted at r. The expected running time of* TreeSample(G) *is proportional to $\sum_u \pi(u)C_{u,r}$, i.e., the expected commute time between r and a π-random vertex.*

TreeSample(G, r)
begin
  InTree$[u] \leftarrow$ **false**, for all $u \in V \setminus \{r\}$;
  InTree$[r] \leftarrow$ **true**;
  for all $s \in V$:
    $u \leftarrow s$;
    while not *InTree*$[u]$:
      Select $v \in \Gamma(u)$, u.a.r.;
      Tree$[u] \leftarrow v$;
      $u \leftarrow v$;
    $u \leftarrow s$;
    while not *InTree*$[u]$:
      InTree$[u] \leftarrow$ **true**;
      $u \leftarrow$ *Tree*$[u]$;
  return *Tree*
end

Figure 5.1: An implementation of the cycle-popping strategy.

Proof The running time of TreeSample is proportional to the number of steps in the simulated random walk. For each $u \neq r$ we estimate the number of steps that are taken *from* u; the total number of steps will be the sum of these. The key observation is that the number of steps from u is one greater that the number of coloured cycles containing u that are popped. So the number of steps from u is dependent on u and the r.v's (S_u^i), but *not* on the order in which the starting points are considered by TreeSample, i.e., the order in which vertices s are taken in the outer loop. Thus we may assume without loss of generality that $s = u$ is the first vertex to be selected. The expected number of steps from u is the expected number of visits to u (including the visit at time 0) made by a random walk started at u before hitting r. The latter quantity is $\pi(u)C_{u,r}$, see Lemma 2.5.2. The result now follows from Theorem 5.2.1. □

The commute time between any pair of vertices is bounded by the twice the cover time. Thus the expected running time of TreeSample(G) when applied to an undirected graph G (i.e., all directed edges occur in antiparallel pairs) is $O(nm)$, where m is the number of edges, see Lemma 2.5.4.

5.3 Independent sets in matroids: forests

The *independent sets* $\mathcal{I}(M)$ of a matroid $M = (E, \mathcal{B})$ are the subsets of the ground set E that may be extended to a base; thus, $\mathcal{I}(M) = \{I \subseteq E : \exists X \in \mathcal{B}(M).\, X \supseteq I\}$. In the case of the cycle matroid of a graph G, the independent sets are the forests in G.

There is a natural random walk on the independent sets of a matroid. Suppose the current independent set is I. Select an element $e \in E$ of the ground set u.a.r., and let $I' = I \oplus \{e\}$; if $I' \in \mathcal{I}(M)$ then move to I', otherwise remain at I. (If desired, exchange moves akin to those employed in the bases-exchange walk may be added.) At first sight, performing a random walk on all independent sets rather than just the maximal independent sets (i.e., bases) appears to allow more freedom, only increasing the potential for rapid mixing. However, this initial impression is misleading, and it is not known whether the natural random walk on independent sets is rapidly mixing, even in the special case of graphic matroids. Exponential mixing time is consistent with our present knowledge.

There is another way to connect forests with matroids. The set of all k-edge forests in a graph G can be viewed as the set of bases of a matroid: a truncation of the cycle matroid of G. Unfortunately, we saw at the end of §5.1.2 that the truncation of a graphic matroid is not necessarily balanced, so we cannot employ the machinery so far established.

However, there is a special situation where we do know how to sample forests in a graph G, and that is when G is sufficiently dense. For $\alpha > 0$ we say that a graph G is *α-dense* if every vertex in G has degree at least αn. The main result of this section is that there is a polynomial-time uniform sampler for forests in α-dense graphs. (The degree of the polynomial governing the runtime of the sampler grows unboundedly as $\alpha \to 0$.) The idea, due to Annan [?], is to reduce the current problem to the already solved problem of sampling spanning trees.

Theorem 5.3.1 *Suppose we have a procedure* TreeSample(H) *for sampling, u.a.r., spanning trees in a graph H. There is a polynomial-time algorithm* ForestSample *that, given access to* TreeSample, *takes an α-dense graph G as input and satisfies the following specification:*

- ForestSample *either produces a forest in G or no output; the output distribution, conditioned on there being an output, is uniform over all forests in G.*
- ForestSample *produces an output with probability at least $\frac{1}{2}$.*
- *The number of calls to* TreeSample *is bounded by $n^{4/\alpha}$.*

Let $G = (V, E)$ be an n-vertex graph with vertex set V and edge set E. Denote by G^+ the derived graph with vertex set $V^+ = V \cup \{t\}$ and edge set $E^+ = E \cup \{\{v, t\} : v \in V\}$. Each spanning tree (V^+, T) in G^+ projects to a forest (V, F) in G, where the edge set of the forest is simply $F = T \cap E$. Moreover, every forest in G may be derived from at least one spanning tree in G^+ by projecting in this way. This observation in itself does not provide a reduction from forest sampling to spanning tree sampling, as the number of distinct spanning trees in G^+ projecting to given forest F varies widely as a function of F. At one extreme, the forest consisting of n trivial components arises in just one way; while, at the other, the forest consisting of $n/2$ components of size 2 (i.e., a perfect matching, assuming n is even) arises in $2^{n/2}$ ways. In general, the number of spanning trees corresponding to a specified forest is the product of the sizes of the connected components forming that forest.

ForestSample(G)
begin
    Construct G^+ as described in the text;
    repeat $2n^{4/\alpha}$ times, or until successful:
        $(V^+, T) \leftarrow$ TreeSample(G^+);
        $F \leftarrow T \cap V$
        if (V^+, T) is the canonical tree for forest (V, F) **return** (V, F)
end

Figure 5.2: A procedure for sampling forests in an α-dense graph.

To overcome this problem we nominate a canonical spanning tree in G^+ for each forest F.[3] For example, assume a linear ordering on the vertices of V and deem a spanning tree T canonical if it contains edges from t to the least vertex in every connected component of G. Clearly, this rule results in one canonical tree for each forest. Certainly, then, the output distribution of the procedure presented in Figure 5.2 is uniform over forests in G. What is not immediately clear is that the procedure will ouput *some* forest with probability at least $\frac{1}{2}$. The key fact we need to prove this is the following.

Lemma 5.3.1 *Suppose the n-vertex graph G is α-dense, and let G^+ be the derived graph as defined above. Let $\{v, t\}$ be any of the n edges in G^+ with an endpoint at t. Select T u.a.r. from the set of all spanning trees in G^+. Then the probability that T contains the edge $\{v, t\}$ is at most $2/(\alpha n + 2)$.*

Proof In Annan's proof, this lemma is established using connections between spanning trees and resistances in electrical networks. For us, it is more convenient to appeal to the property of balance, since we have already set up the machinery. Let $e \in E$ be any edge in G. Since graphic matroids are balanced, the probability that edge $\{v, t\}$ is contained in a random spanning tree of G^+ is not decreased by the removal of edge e. Removing in turn all edges $e \in E$ that are not incident at v yields a subgraph of G^+ whose only remaining edges are $(\{t\} \times V) \cup (\{v\} \times \Gamma(v))$, where $\Gamma(v)$ denotes the set of neighbours of v in G. By direct calculation, the number of spanning trees in this

[3]For convenience, in the remainder of the section, we blur the distinction between a forest or tree and the edges that compose it. Since all forests and trees are spanning, this will cause no confusion.

vestigial graph which contain (respectively, do not contain) the edge $\{t, v\}$ is 2^k (respectively, $k2^{k-1}$), where $k = |\Gamma(v)| \geq \alpha n$. Thus the probability that $\{t, v\}$ is contained in a random spanning tree of G^+ is at most $2^k/(k2^{k-1} + 2^k) \leq 2/(\alpha n + 2)$, as claimed. □

Lemma 5.3.1 assures us that a typical random spanning tree in G^+ projects to a forest with few components.

Proof of Theorem 5.3.1 It is only the second of the three claims in the statement of the theorem that remains to be proved. Let T be a spanning tree in G^+ selected u.a.r., and let F be the derived forest in G. According to Lemma 5.3.1, the expected degree of vertex t in T is at most $2/\alpha$. By Markov's inequality, with probability at least $\frac{1}{2}$, the degree of t is no greater than $4/\alpha$. Conditioned on the event that the degree of t is at most $4/\alpha$, the probability that T is canonical for F is at least $n^{-4/\alpha}$. So the probability that *some* forest is output is at least $\frac{1}{2}n^{-4/\alpha}$. The probability that *none* of the $n^{4/\alpha}$ trials produces an output is therefore bounded above by $(1 - \frac{1}{2}n^{-4/\alpha})^{2n^{4/\alpha}} \leq 1/e \leq \frac{1}{2}$. □

As a simple corollary of the above, we can show how to generate and count *trees* (not necessarily spanning) in a dense graph.

Corollary 5.3.1 *Let G be an α-dense graph. There is a good sampler and an FPRAS for the set of* all *trees of G.*

Proof Consider the following algorithm:

(i) Choose forest F at random.

(ii) Accept if F contains one non-trivial tree plus a collection of isolated vertices.

We claim that

$$\mathbf{Pr}(\text{accept in (ii)}) \geq n^{-\lfloor 4/\alpha \rfloor}. \tag{5.25}$$

Let f_k denote the number of forests with k non-trivial trees and $f = \sum_{k=1}^{n} f_k$ and let $b = \lfloor 4/\alpha \rfloor$. Then

(a) $\mathbf{Pr}(\text{accept in (ii)}) \geq \frac{f_1}{f}$.

(b) $f_1 + f_2 + \cdots + f_b \geq f/2$.

(c) $f_{k+1} \leq n f_k$.

Here, (a) is clear, (b) follows from Lemma 5.3.1 and (c) is a consequence of the fact that we can obtain all $k + 1$ tree forests by deleting an edge of a k tree forest. So

$$\begin{aligned} f_1 + f_2 + \cdots + f_b &< (1 + n + n^2 + \cdots + n^{b-1}) f_1 \\ &= \frac{n^b - 1}{n - 1} f_1 \end{aligned}$$

and the result follows. □