Balanced Matroids

Tomás Feder and *Milena Mihail*
Bell Communications Research
445 South Street, Morristown NJ 07960

tomas@bellcore.com mihail@bellcore.com

Abstract

We introduce the notion of "balance", and say that a matroid is balanced if the matroid and all its minors satisfy the property that, for a randomly chosen basis, the presence of an element can only make any other element less likely.

We establish strong expansion properties for the bases-exchange graph of balanced matroids; consequently, the set of bases of a balanced matroid can be sampled and approximately counted using rapidly mixing Markov chains. Thus, the general problem of approximately counting bases (known to be $\#P$-complete) is reduced to that of showing balance. Specific classes for which balance is known to hold include graphic and regular matroids.

1 Introduction and Preliminaries

Let S be a finite ground-set, and let $\mathcal{B}$ be a collection of subsets of S. Following standard terminology, the set $\mathcal{B}$ is said to form the collection of *bases* of a *matroid* $\mathcal{M}(S, \mathcal{B})$ if and only if (i) all sets in $\mathcal{B}$ have the same cardinality (called the *rank* of the matroid), and (ii) for any pair of bases B_1 and B_2 the following exchange property holds: for all $e \in B_1$ there exists an $f \in B_2$ such that $B_1 \setminus \{e\} \cup \{f\}$ is in $\mathcal{B}$. A main example is that of graphic matroids, whose ground-set is the set of edges of a given graph and whose bases are the spanning trees of the graph. Another example is that of vectorial matroids, whose ground-set is a set of vectors over some field and whose bases are maximum cardinality linearly independent subsets of the set of vectors.

The *bases-exchange graph* of a matroid $\mathcal{M}$, henceforth denoted by $G(\mathcal{M})$, was introduced by Edmonds [9] as the graph whose vertex-set is the collection of bases of the matroid, with two bases B_1 and B_2 connected by an edge if and only if B_2 can be obtained from B_1 by the fundamental operation of removing and adding one ground-set element: $B_2 = B_1 \setminus \{e\} \cup \{f\}$. The bases-exchange graph has been studied before in various combinatorial contexts [11, 19]; here we focus on expansion related properties of $G(\mathcal{M})$.

If B is a basis chosen uniformly at random from $\mathcal{B}$, and e is an element of S, let e denote the event $e \in B$ indicating that e is in the chosen basis. The matroid $\mathcal{M}(S, \mathcal{B})$ is said to satisfy the *negative correlation* property if the inequality

$$\Pr[ef] \leq \Pr[e] \Pr[f]$$

holds for all pairs of distinct elements e, f in S. Realize that negative correlation is equivalent to $\Pr[e|f] \leq \Pr[e]$, thus expressing the intuitive fact that the presence of an element f can only make another element e less likely. The negative correlation property has been studied before in [5, 22]:

there are essentially only two matroids known to violate it [21], in that every known counterexample contains one of them as a minor. One of these two matroids is S_8 [22]; in fact, it is known that every binary (i.e. vectorial over $GF[2]$) matroid not containing S_8 as a minor is balanced [21]. This class of binary matroids contains all the graphic matroids and more generally the regular matroids. (A *minor* of a matroid is obtained by repeatedly performing the operation of choosing an element $e \in S$ and then selecting either those bases that contain e or those that do not. In the case of graphic matroids, this operation corresponds to contraction or deletion of selected edges from the graph.)

Here we strengthen the notion of negative correlation to that of "balance": We say that a matroid $\mathcal{M}(S, \mathcal{B})$ is *balanced* if all its minors, including $\mathcal{M}$ itself, satisfy the negative correlation property.

We first establish (Theorems 3.1 and 3.4) that the bases-exchange graph $G(\mathcal{M})$ of any balanced matroid $\mathcal{M}$ has *cutset expansion* 1, i.e., for any bipartition of the vertices of $G(\mathcal{M})$, the number of edges incident to both partition classes is at least as large as the size of the smaller partition class. Such strong expansion properties had been conjectured by Mihail and Vazirani [18], in fact for significantly wider classes of graphs. The core of our proof here is to show that the negative correlation property for two elements e and f implies negative correlation between an element e and an arbitrary monotone property m on $S\backslash\{e\}$,

$$\Pr[em] \leq \Pr[e]\Pr[m] \quad ,$$

which implies an enforcement of vertex neighborhood analogous to bipartite expansion for certain subgraphs on $G(\mathcal{M})$, which finally implies expansion for $G(\mathcal{M})$. It should be pointed out that all steps in the reasoning that appears here are entirely new. Previous arguments for expansion of combinatorial graphs involved using elements of the state-space to encode paths and bound path congestion (introduced by Jerrum and Sinclair [13] to approximate permanents), the discrete version of isoperimetries from differential geometry (used by Dyer, Frieze, and Kannan and others [8, 15] to approximate the volumes of convex bodies), and coupling arguments (used by Aldous [2] and Broder [4] to sample spanning trees).

The algorithmic significance of the expansion of the bases-exchange graph derives from a sequence of well known ideas, see for example [1, 3, 8, 13, 14, 23], as follows:

• Consider the *natural random walk* X_t, $t=0,1,\ldots$ on $G(\mathcal{M})$: If X_t is the state (basis) at time t, X_{t+1} is determined by choosing e from X_t and f from S uniformly at random and, if $X' = X_t\backslash\{e\}\cup\{f\} \in \mathcal{B}$ then $X_{t+1}=X'$, otherwise $X_{t+1}=X_t$. It is easy to see that the Markov chain X_t can be simulated efficiently (given, say, an independence oracle), and that it converges to the uniform distribution over $\mathcal{B}$ (it is symmetric). Most importantly, the expansion of $G(\mathcal{M})$ suggests that X_t has large *conductance*, and hence possesses the *rapid mixing property*; as is well known the rapid mixing property amounts, roughly, to X_t approaching its stationary distribution arbitrarily closely for $t = \text{poly-log}(|\mathcal{B}|)$. Therefore, the natural random walk on $G(\mathcal{M})$ can be used as an efficient almost uniform sampling scheme for the set of bases of any balanced matroid.

• In turn, efficient almost uniform sampling for self-reducible combinatorial populations yields efficient Monte-Carlo approximation schemes for their size, and, in view of the fact of self-reducibility for balanced matroids, we obtain an efficient randomized algorithm to approximately count the number of bases of any balanced matroid.

• Therefore, our results reduce the question of efficiently sampling and approximately counting the set of bases of a matroid to that of establishing balance. In general, the problem of exact counting the number of bases of a matroid is $\#P$-complete [24], and thus it would be of special interest to obtain efficient approximations. In particular, such fundamental problems as computing the reliability of a network can be formulated as counting the number of bases of suitably chosen matroids (truncations of the dual of the graphic matroid).

We further show that balance can be used to define random paths between any pair of bases on the bases-exchange graph, such that the length of these paths and the path congestion through any specific edge can be bounded (Section 4). It should be pointed out that the technique used here to bound path congestion is new, in that it is derived by means of establishing the existence of certain *fractional matchings*, which in turn control the flow of paths through any edge of the bases-exchange graph. The only previously known method to bound path congestion was Jerrum and Sinclair's argument of using the state-space to encode paths [13]; the arguments used here are of completely different type. Furthermore, by extending a technique of Diaconis and Strook [6] and Mihail [16] we show that analysis that uses paths rather than conductance can yield a sharper convergence rate for the natural random walk as well as for a modified version of the natural random walk. The algorithmic significance of this improvement is as follows:

• For arbitrary balanced matroids, where n is the rank, m is the cardinality of the ground-set, and ϵ is a bound on total variation distance (here, as usual, we use total variation distance to measure convergence rate), we obtain a sampling scheme based on the natural random walk with $O((n\log m+\log\epsilon^{-1})mn^2)$ convergence rate (Theorem 5.1). In particular, this result applies to graphic and regular matroids and significantly improves upon previously known bounds. For graphic matroids, for the natural random walk, Broder's coupling and conductance arguments yielded, roughly, $O((n\log m+\log\epsilon^{-1})m^2n^4)$ convergence rate [4]. For regular matroids, Dyer and Frieze's geometric arguments yielded $O((n\log m+\log\epsilon^{-1})m^4n^4)$ convergence rate [7].

• We further analyze a modification of the natural random walk and show $O((n\log m+\log\epsilon^{-1})n^3)$ convergence rate (Theorem 5.2). This bound is the first known bound for matroids that remains polynomial even for an exponentially large ground-set.

• For regular matroids we need $O(mn)$ time to implement each step of the modified random walk (details of this implementations are tedious but straightforward and are left for the full paper), which results in $O((n\log m+\log\epsilon^{-1})mn^4)$ running time. There is an alternative sampling scheme based on Kirkhoff's tree matrix theorem [25] (involving evaluation of determinants and arithmetric modulo primes) with running time $O(mn^4\log m)$. However, even though our scheme is worse by a factor of n, it is very simple conceptually and much easier to implement (futhermore our analysis on the running time is not necessarily tight).

• For graphic matroids we need $O(\sqrt{m})$ time to implement each step of the modified random walk, along the lines of Frederickson [10] which results in $O((n\log m+\log\epsilon^{-1})\sqrt{m}n^3)$ running time. For the natural random walk, each step can be implemented in $O(\log n)$ time [10]. There is an alternative sampling scheme due to Aldous [2] and Broder [4] with $E(C)$ running time, where C is the cover time of the underlying graph. That scheme thus runs in expected time $O(nm)$ and is clearly faster than the one presented here. However, when we introduce a large number of parallel edges in the graph the expectation of the cover time blows up analogously and [2, 4] fail.

• For the latter case of graphs with a large number of parallel edges we obtain a sampling scheme with $O((n\log m'+\log\epsilon^{-1})\sqrt{m'}n^3)$ running time, where m' is the number of edges when parallel edges are counted only once; this follows from Theorem 5.3. (More generally, the introduction of parallel elements for a balanced matroid can significantly increase the convergence rate of the natural random walk, but leaves the bounds for the modified random walk unchanged.) The alternative scheme based on Kirkhoff's tree matrix theorem is both less efficient, $O(m'n^4 \log m)$, and more complicated to implement. The ability to introduce a large number of parallel edges without any increse on the running time makes it possible, for example, to output a spanning tree in a graph where each edge has an associated failure probability (such probabilities can be represented by introducing a proportional number of parallel edges), so that each spanning tree is output with probability proportional to that of its presence in a random configuration of the graph.

We give a combinatorial proof of balance for regular matroids, exhibiting a structure that can be used to significantly reduce the convergence rates for certain regular matroids. Previous proofs were of algebraic nature; e.g. see [5] for the special case of graphic matroids.

We pose the question of whether the balance property is preserved when the truncations of a balanced matroid are considered. If this property holds in the case of graphic matroids or their duals, then the number of forests or, respectively, the number of spanning subgraphs of given number of edges can be efficiently approximated (duality preserves balance). This would lead to an approximation scheme for the $(2,1)$ point of the Tutte plane [12, 26], and for network reliability [24], respectively. Another important family of matroids for which the balance property remains open is that of transversals.

The remainder of the paper is organized as follows. Section 2 gives a new combinatorial proof of the fact that all regular matroids, and in particular graphic matroids, are balanced. Section 3 exploits balance to satisfy a certain ratios condition and establish expansion. Section 4 uses the existence of certain fractional matchings derived from the ratios condition to construct paths in the bases-exchange graph; these paths are then used in Section 5 to bound the convergence rates of two sampling schemes for balanced matroids. Section 6 introduces a decomposition property that follows in certain cases from the construction in Section 2 and implies then tighter bounds on convergence rates.

2 Achieving Balance

Suppose that $\mathcal{F}$ is a minor-closed family of matroids that satisfy the negative correlation property. Then all matroids in $\mathcal{F}$ are balanced as well. The negative correlation property has been established for several classes of matroids. A main example is that of graphic matroids, whose bases are the edge sets of spanning trees in a given graph. More generally, all regular matroids are negatively correlated, and the fact that their minors are also regular implies that all regular matroids are balanced. We give here a new combinatorial proof of the fact that regular matroids are balanced.

We refer to the circuits and cocircuits of a matroid as cycles and cuts, by analogy with the case of graphic matroids. A cycle is thus a minimal dependent set, while a cut is a minimal set whose complement is not spanning. It is known that all regular matroids, as *orientable binary matroids*, satisfy the following property [25]. It is possible to assign (1) values $C(g)$ for each cycle C and each element g, so that $C(g) = \pm 1$ if $g \in C$ and $C(g) = 0$ otherwise; and (2) values $D(g)$ for each cut D and each element g, so that $D(g) = \pm 1$ if $g \in D$ and $D(g) = 0$ otherwise; most importantly,

$$\sum_g C(g)D(g) = 0$$

for all cycles C and cuts D. Intuitively, in the case of graphic matroids, assign a conventional direction to all edges in the underlying graph, then traverse each cycle C in one of the two possible directions, setting $C(g) = \pm 1$ for each edge traversed depending on whether the edge is traversed in the conventional direction; the edges of each cut D separate a set A of vertices in the graph from its complement $\overline{A}$, so traverse all edges in the cut from A to $\overline{A}$, setting $D(g) = \pm 1$ depending on whether the edge is traversed in the conventional direction. Since a cycle will traverse a cut the same number of times from A to $\overline{A}$ as from $\overline{A}$ to A, it follows that the sum over all g of $C(g)D(g)$ is zero.

We refer to the independent sets whose cardinality is one less than the cardinality of bases as *near-bases*, since they can all be obtained from bases by removing one element. Every near-basis N defines a unique cut D_N contained in the complement of N. We refer to the spanning sets whose cardinality is one more than the cardinality of bases as *unicycles*; all such spanning sets can be obtained from bases by adding one element. Every unicycle U defines a unique cycle C_U contained in U. A useful property relating near-bases and unicycles is the following. If $U = N \cup \{e, f\}$, then

$$D_N(e)D_N(f) = -C_U(e)C_U(f).$$

To see this, note that the only elements that could be both in D_N and in C_U are e and f, so that by the sum property we have $C_U(e)D_N(e) + C_U(f)D_N(f) = 0$; the above equality follows then if all four quantities involved are zero, while if one of them is non-zero, say $D_N(e) \neq 0$, then $N \cup \{e\}$ is a basis, so $U \backslash \{f\}$ is a basis and $C_U(f) \neq 0$, giving $D_N(f)/D_N(e) = -C_U(e)/C_U(f)$.

We can now define an important quantity. If $e \neq f$, we let

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

The equality of both expressions follows from the fact that non-zero terms in the sum arise only with pairs N, U such that $U = N \cup \{e, f\}$. For the theorem below, it will be convenient to select a specific element e and require that $D(e) = -C(e) = 1$ for all cycles C and cuts D containing it. This condition can easily be enforced by changing the signs of all elements of chosen cycles or cuts, without violating the conditions on $C(g), D(g)$, and without affecting the value of Δ_{ef}. With this condition, the above equations become simply

$$\sum_{g \neq e} C(g)D(g) = 1 \text{ provided } e \in C \cap D,$$

$$D_N(f) = C_U(f) \text{ for } U = N \cup \{e, f\}, \text{ and}$$

$$\Delta_{ef} = \sum_{N: e \in D_N} D_N(f) = \sum_{U: e \in C_U} C_U(f).$$

Intuitively, for graphic matroids, the quantity Δ_{ef} measures whether cycles containing e, f arising from unicycles tend to traverse e and f in the same or in opposite directions, or equivalently, whether cuts containing e, f arising from near-bases tend to be traversed by e and f in the same or in opposite directions. For graphic matroids, we can show that the quantity Δ_{ef} coincides with a quantity from the theory of electrical networks: If a unit resistance is assigned to each edge in the graph, and a current of $|\mathcal{B}|$ enters and leaves at the endpoints of e, then the potential drop between the endpoints of f is Δ_{ef}. The result below was shown for graphs with Δ_{ef} defined as a potential drop in [5]. We use indices on $\mathcal{B}$ to indicate subsets of $\mathcal{B}$ satisfying certain conditions concerning the presence or absence of certain elements in the chosen bases.

Theorem 2.1 *The bases of a regular matroid satisfy* $|\mathcal{B}| \cdot |\mathcal{B}_{ef}| = |\mathcal{B}_e| \cdot |\mathcal{B}_f| - \Delta_{ef}^2$.

Proof. From pairs $(B, B') \in \mathcal{B}_{\overline{e}} \times \mathcal{B}_{ef}$, we obtain pairs $(B'', B''') \in \mathcal{B}_e \times \mathcal{B}_{\overline{e}f}$, by means of an exchange involving e and an element $g \neq e$. More specifically, we let $B'' = B \cup \{e\} \backslash \{g\}$, and

$B''' = B'\backslash\{e\} \cup \{g\}$, and assign a *weight* to this exchange equal to $C_{B\cup\{e\}}(g)D_{B'\backslash\{e\}}(g)$. If this weight is non-zero, then the resulting B'' and B''' are indeed bases in $\mathcal{B}_e$ and $\mathcal{B}_{\overline{e}f}$ respectively, and a non-zero weight cannot occur here for $g = f$ since $D_{B'\backslash\{e\}}(f) = 0$.

The above weight can also be expressed as $D_{B''\backslash\{e\}}(g)C_{B'''\cup\{e\}}(g)$. From the point of view of a pair $(B'', B''') \in \mathcal{B}_e \times \mathcal{B}_{\overline{e}f}$, if this weight is non-zero for some $g \neq e, f$, then the reverse exchange $B = B''\backslash\{e\} \cup \{g\}$ and $B' = B''' \cup \{e\}\backslash\{g\}$ gives back a pair $(B, B') \in \mathcal{B}_{\overline{e}} \times \mathcal{B}_{ef}$. Therefore

$$
\begin{aligned}
|\mathcal{B}_e| \cdot |\mathcal{B}_{\overline{e}f}| &= \sum_{(B'',B''')\in\mathcal{B}_e\times\mathcal{B}_{\overline{e}f}} \Big(\sum_{g\neq e} D_{B''\backslash\{e\}}(g)C_{B'''\cup\{e\}}(g)\Big) \\
&= \sum_{(B'',B''')\in\mathcal{B}_e\times\mathcal{B}_{\overline{e}f}} \Big(\sum_{g\neq e,f} D_{B''\backslash\{e\}}(g)C_{B'''\cup\{e\}}(g)\Big) + \sum_{(B'',B''')\in\mathcal{B}_e\times\mathcal{B}_{\overline{e}}} D_{B''\backslash\{e\}}(f)C_{B'''\cup\{e\}}(f) \\
&= \sum_{(B,B')\in\mathcal{B}_{\overline{e}}\times\mathcal{B}_{ef}} \Big(\sum_{g\neq e} C_{B\cup\{e\}}(g)D_{B'\backslash\{e\}}(g)\Big) + \Big(\sum_{B''\in\mathcal{B}_e} D_{B''\backslash\{e\}}(f)\Big)\Big(\sum_{B'''\in\mathcal{B}_{\overline{e}}} C_{B'''\cup\{e\}}(f)\Big) \\
&= |\mathcal{B}_{\overline{e}}| \cdot |\mathcal{B}_{ef}| + \Delta_{ef}^2,
\end{aligned}
$$

and the theorem follows. □

The theorem immediately implies $\Pr[ef] \leq \Pr[e]\Pr[f]$, so negative correlation and hence balance follow for regular matroids.

3 From Balance to Ratios, Expansion, and Fractional Matchings

In this section we derive expansion for balanced matroids. The proof is essentially inductive on the rank and, in an intermediate step, relies on the fact that balance implies a certain enforcement of ratios analogous to bipartite expansion for the bipartite subgraph of $G(\mathcal{M})$, $G_e(\mathcal{M}) = G(\mathcal{B}_e, \mathcal{B}_{\overline{e}}, E)$, where edges that correspond to exchanges not involving a specific element e are ommitted. In particular, let Γ_e denote vertex neighborhood in $G_e(\mathcal{M})$, and say that the bases-exchange graph of a matroid $\mathcal{M}(\mathcal{B}, S)$ *enforces ratios* if for every element $e \in S$ the following holds:

$$\forall \mathcal{A} \subseteq \mathcal{B}_e \qquad \frac{|\Gamma_e(\mathcal{A})|}{|\mathcal{B}_{\overline{e}}|} \geq \frac{|\mathcal{A}|}{|\mathcal{B}_e|} \tag{1}$$

$$\forall \mathcal{A} \subseteq \mathcal{B}_{\overline{e}} \qquad \frac{|\Gamma_e(\mathcal{A})|}{|\mathcal{B}_e|} \geq \frac{|\mathcal{A}|}{|\mathcal{B}_{\overline{e}}|} \quad . \tag{2}$$

Theorem 3.1 *For every balanced matroid $\mathcal{M}(\mathcal{B}, S)$, the bases exchange graph of $\mathcal{M}$ enforces ratios.*

Proof. Let $\mathcal{A} \subseteq \mathcal{B}_e$ and let $m_{\mathcal{A}} = \bigvee_{B\in\mathcal{A}} \bigwedge_{e_i\in B, e_i\neq e} e_i$. Notice that the set of bases in $\mathcal{B}_e$ satisfying $m_{\mathcal{A}}$ is precisely the set $\mathcal{A}$, while the set of bases in $\mathcal{B}_{\overline{e}}$ satisfying $m_{\mathcal{A}}$ is precisely the set $\Gamma_e(\mathcal{A})$. Hence (1) is equivalent to

$$\Pr[m_{\mathcal{A}}|\overline{e}] \geq \Pr[m_{\mathcal{A}}|e] \quad . \tag{3}$$

Analogously, for $\mathcal{A} \subseteq \mathcal{B}_{\overline{e}}$, let $\overline{m_{\mathcal{A}}} = \bigvee_{B\in\mathcal{A}} \bigwedge_{e_i\notin B, e_i\neq e} \overline{e_i}$, and notice that the set of bases in $\mathcal{B}_{\overline{e}}$ satisfying $\overline{m_{\mathcal{A}}}$ is the set $\mathcal{A}$, and the set of bases in $\mathcal{B}_e$ satisfying $\overline{m_{\mathcal{A}}}$ is the set $\Gamma_e(\mathcal{A})$. Hence (2) is equivalent to

$$\Pr[\overline{m_{\mathcal{A}}}|e] \geq \Pr[\overline{m_{\mathcal{A}}}|\overline{e}] \quad . \tag{4}$$

In turn, (3) and (4) follow from the following lemma:

Lemma 3.2 *For any balanced matroid $\mathcal{M}(\mathcal{B}, S)$ on ground set S, and for any monotone property m over the variables in $S \setminus \{e\}$, the following holds:*

$$\Pr[me] \leq \Pr[m]\Pr[e] \quad .$$

Proof. We show equivalently that $\Pr[m|e] \leq \Pr[m]$. The reasoning is inductive on the size of the ground-set. The case where $\mathcal{M}$ has rank $n=1$ is easy to verify (and this is also the only point where the monotonicity of m is used). For the inductive step note that

$$\begin{aligned} \Pr[m|e] &= \Pr[f|e]\Pr[m|fe] + \Pr[\overline{f}|e]\Pr[m|\overline{f}e], \text{ and} \\ \Pr[m] &= \Pr[f]\Pr[m|f] + \Pr[\overline{f}]\Pr[m|\overline{f}]. \end{aligned}$$

Note further that (i) $\Pr[f|e] \leq \Pr[f]$ from the fact that $\mathcal{M}$ is balanced, (ii) $\Pr[m|fe] \leq \Pr[m|f]$ by applying the inductive hypothesis on $\mathcal{B}_f$ for the property $m|_{f=1}$, and (iii) $\Pr[m|\overline{f}e] \leq \Pr[m|\overline{f}]$ by applying the inductive hypothesis on $\mathcal{B}_{\overline{f}}$ for the monotone property $m|_{f=0}$. If, in addition, it was the case that (iv) $\Pr[m|fe] \geq \Pr[m|\overline{f}e]$, then the lemma would follow by averaging principles: $\Pr[m|e] \leq \Pr[f]\Pr[m|fe] + \Pr[\overline{f}]\Pr[m|\overline{f}e]$ by (i) and (iv), $\leq \Pr[f]\Pr[m|f] + \Pr[\overline{f}]\Pr[m|\overline{f}] = \Pr[m]$ by (ii) and (iii).

We argue that there exists always some element f such that (iv) holds. In particular, note that

$$\sum_{f \neq e} \Pr[f|me] = n - 1 = \sum_{f \neq e} \Pr[f|e]$$

Hence for some f, $\Pr[f|me] \geq \Pr[f|e]$ (with $\Pr[f|e] > 0$), which is equivalent to $\Pr[m|fe] \geq \Pr[m|e]$, which is finally equivalent to $\Pr[m|fe] \geq \Pr[m|\overline{f}e]$. This completes the proof of Lemma 3.2, and Theorem 3.1. □□
Remark: The proof of the lemma can be extended to show that if m_1 and m_2 are two monotone properties over disjoint sets of variables from the ground-set, then $\Pr[m_1 m_2] \leq \Pr[m_1]\Pr[m_2]$.

Say that a bipartite graph $G(U, V, E)$ has a *fractional matching* if there exists an assignment of nonnegative weights to the edges in E such that for each vertex in $u \in U$, the sum of the weights of edges incident on u is $|V|$, and for each vertex in $v \in V$, the sum of the weights of edges incident on v is $|U|$.

Corollary 3.3 *For a balanced matroid $\mathcal{M}(\mathcal{B}, S)$, every element $e \in S$ is such that the bipartite graph $G_e(\mathcal{M}) = G(\mathcal{B}_e, \mathcal{B}_{\overline{e}}, E)$, obtained from the bases exchange graph by ommitting exchanges not involving e, admits a fractional matching.*

Proof. Consider the bipartite graph G^* obtained from G_e by making $|\mathcal{B}_{\overline{e}}|$ copies of each basis in $\mathcal{B}_e$, making $|\mathcal{B}_e|$ copies of each basis in $\mathcal{B}_{\overline{e}}$, and including edges between all copies of each pair of adjacent bases in G_e. The fact that the basis exchange graph of $\mathcal{M}$ enforces ratios implies now that G^* satisfies Hall's condition, and hence G^* has a perfect matching P. Finally, by identifying the copies of each basis and assigning to each edge of E a weight equal to the number of edges of P that correspond to it, the conditions for a fractional matching are satisfied. □

Theorem 3.4 *For every balanced matroid $\mathcal{M}(\mathcal{B}, S)$, the bases-exchange graph of $\mathcal{M}$ has cutset expansion 1.*

The proof shows that enforcement of ratios implies expansion (Mihail and Sudan [17]) and is given in Appendix 1. An alternative proof involving fractional matchings and the paths constructed in Section 4 is also given.

Corollary 3.5 *For any balanced matroid $\mathcal{M}(S,\mathcal{B})$, the natural random walk X_t, $t = 0, 1, \ldots$, on the bases exchange graph $G(\mathcal{M})$ is rapidly mixing: the conductance is $\Phi \geq 1/2mn$, and hence the total variation distance $d(t)$ can be bounded by ϵ in time $t = \Omega(n \log m + \log \epsilon^{-1})m^2n^2$.*

Proof. The total variation distance $d(t)$ has been bounded from above by $|\mathcal{B}|\left(1-\Phi^2/2\right)^t$ in [23], where the conductance Φ in the case of symmetric Markov chains is the product of the transition probabilities with the cutset expansion. Hence, from Theorem 3.4 and the definition of the random walk, Φ can be bounded by by $1/2mn$. □

4 From Fractional Matchings to Path Congestion

In order to obtain tighter bounds on convergence rates, we shall use the fractional matchings defined in Corollary 3.3 of the previous section to construct paths joining all pairs of bases in the bases-exchange graph of a balanced matroid, while keeping the path lengths and the number of paths through each basis small.

We thus wish to construct $|\mathcal{B}|^2$ paths in the bases-exchange graph of a balanced matroid $\mathcal{M}(S,\mathcal{B})$, one path for each choice of an origin B and a destination B'. The construction begins with the choice a random permutation $e_1, \ldots, e_m$ of the m elements of S. Consider a specific destination $B' = B_{\dot{e}_1 \ldots \dot{e}_m}$, where each $\dot{e}_i$ is either e_i or $\overline{e_i}$, depending on whether it is present or absent in B'. Note that the paths destined to B' are initially uniformly arranged over the set of all bases $\mathcal{B}$, in that each basis B has exactly one path destined to B'. We shall ensure that after i steps, the paths destined to B' are uniformly arranged over $\mathcal{B}_i = \mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i}$, the set of bases that agree with B' in the first i elements, in the sense that the expected number of paths at each such basis after i steps is the same for all of them. This condition holds initially, when $i = 0$ and $\mathcal{B}_i = \mathcal{B}$. Assume inductively that after i steps, the paths destined to B' are uniformly arranged over $\mathcal{B}_i = \mathcal{B}_{i+1} \cup \mathcal{B}'_{i+1}$, where $\mathcal{B}_{i+1} = \mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i \dot{e}_{i+1}}$ and $\mathcal{B}'_{i+1} = \mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i \overline{\dot{e}_{i+1}}}$. The paths at $\mathcal{B}_i$ that are already at $\mathcal{B}_{i+1}$ are left there for the $(i+1)$th step. The paths at $\mathcal{B}_i$ that are currently at $\mathcal{B}'_{i+1}$ are sent to $\mathcal{B}_{i+1}$ by means of a fractional matching from $\mathcal{B}'_{i+1}$ to $\mathcal{B}_{i+1}$. Such a fractional matching exists by Corollary 3.3 applied to the minor $\mathcal{M}_{\dot{e}_1 \ldots \dot{e}_i}$; the probabilities given to each possible exchange are chosen proportionally to the associated weights in the fractional matching, thus ensuring that after $(i+1)$ steps, the paths destined to $\mathcal{B}'$ are uniformly arranged over $\mathcal{B}_{i+1}$. After m steps, all paths are at $\mathcal{B}_m = \{B'\}$, as desired.

Lemma 4.1 *The expected number of paths leaving a basis $\hat{B}$ in $\mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i}$ at step i is $|\mathcal{B}| \Pr[\overline{\dot{e}_i}|\dot{e}_1 \ldots \dot{e}_{i-1}]$.*

Proof. The paths leaving a basis $\hat{B}$ in $\mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i}$ at step i are those destined to a B' in $\mathcal{B}_{\dot{e}_1 \ldots \overline{\dot{e}_i}}$. For each such B', before step i, the $|\mathcal{B}|$ paths destined to it are uniformly arranged over $\mathcal{B}_{\dot{e}_1 \ldots \dot{e}_{i-1}}$, so at each of these bases there are $|\mathcal{B}|/|\mathcal{B}_{\dot{e}_1 \ldots \dot{e}_{i-1}}|$ paths destined to B', in expectation. Mutliplying this quantity by the number $|\mathcal{B}_{\dot{e}_1 \ldots \overline{\dot{e}_i}}|$ of possible destinations B' gives the quantity $|\mathcal{B}| \Pr[\overline{\dot{e}_i}|\dot{e}_1 \ldots \dot{e}_{i-1}]$ as required. □

The same bound holds for paths entering $\hat{B}$. Note that a particular bases exchange involving two element e_i and e_j with $i < j$ can only be used during step i, and that the expected number of paths using such an exchange at step i is at most $|\mathcal{B}| \min(\Pr[e_i|\dot{e}_1 \ldots \dot{e}_{i-1}], \Pr[\overline{e_i}|\dot{e}_1 \ldots \dot{e}_{i-1}])$.

Corollary 4.2 *The expected number of paths through any edge $(\hat{B}, \hat{B}')$ is at most $|\mathcal{B}|/2$.*

In order to bound the expected length of paths and the expected number of paths through a basis, we give the following lemma.

Lemma 4.3 *Given a basis $\hat{B}$, consider a sequence of bases and elements $B_0, e_1, B_1, e_2, B_2, \ldots, e_m, B_m = \hat{B}$, where B_i agrees with $\hat{B}$ in $e_1, \ldots, e_i$ and each e_i is chosen uniformly from the elements other than $e_1, \ldots, e_{i-1}$. Then the expected number of B_i's that differ from $\hat{B}$ in e_i is at most $2n$.*

Proof. If e_i is chosen so that B_{i-1} and $\hat{B}$ differ in it, then with probability 1/2 the element e_i is present in $\hat{B}$, since the symmetric difference of two bases contains the same number of elements from each. Since $\hat{B}$ has only n elements, we expect such a choice to happen at most $2n$ times. ☐

Corollary 4.4 *The expected length of a path from any B to any B' is at most $2n$.*

Proof. Follows from Lemma 4.3, where $B = B_0, B_1, \ldots, \mathcal{B}_m = B' = \hat{B}$ is the path from B to B' and $e_1, \ldots, e_m$ is the random order in which the edges were chosen to be fixed. ☐

Say that a path goes through a basis if this basis is on the path but is not its destination.

Corollary 4.5 *The expected number of paths through a basis $\hat{B} = B_{\dot{e}_1 \ldots \dot{e}_m}$ is at most $2n|\mathcal{B}|$.*

Proof. By Lemma 4.1, this expectation is $|\mathcal{B}| E\left[\sum_{i=1}^m \Pr[\overline{\dot{e}}_i | \dot{e}_1, \ldots, \dot{e}_{i-1}]\right]$, the expectation being over all permutations in which the e_i's are chosen. Now Lemma 4.3, where B_0 is chosen uniformly from $\mathcal{B}$ and B_i is chosen uniformly from $\mathcal{B}_{\dot{e}_1 \ldots \dot{e}_i}$, implies that $E\left[\sum_{i=1}^m \Pr[\overline{\dot{e}}_i | e_1, \ldots, e_{i-1}]\right] \leq 2n$. ☐

5 From Path Congestion to Rapid Mixing

In this section we show how the bounds on path congestion for the bases-exchange graph $G(\mathcal{M})$ of a balanced matroid $\mathcal{M}$ yield bounds on the convergence rate of the natural random walk on $G(\mathcal{M})$; the bounds derived here are significantly better than those derived in Corollary 3.5 in terms of conductance. For technical reasons we shall add self loops with probability 1/2 to the natural random walk, so that the walk X_t is as follows: Toss a fair coin. If "heads" then $X_{t+1} = X_t$. If "tails" then choose e from X_t and f from S uniformly at random and, if $X' = X_t \setminus \{e\} \cup \{f\} \in \mathcal{B}$ then $X_{t+1} = X'$ else $X_{t+1} = X_t$.

We further consider the modification of the natural random walk which uses both bases and near-bases; for this modified random walk the convergence rate becomes even faster. The modified random walk Y_t is as follows: Toss a fair coin. If "heads" then $Y_{t+1} = Y_t$. If "tails" then (i) if Y_t is a basis then choose e from Y_t uniformly at random and $Y_{t+1} = Y_t \setminus \{e\}$; (ii) if Y_t is a near-basis and $D_{Y_t} \subseteq S$ is the cut of Y_t, then choose f from D_{Y_t} uniformly at random and $Y_{t+1} = Y_t \cup \{f\}$. It is easy to verify that when Y_t is stationary, the probability of a basis is $1/2|\mathcal{B}|$, and the probability of a near-basis with cut D is $|D|/2n|\mathcal{B}|$; furthermore the probability that Y_t is on a basis is exactly 1/2. Hence Y_t is also appropriate to use as a sampling scheme for $\mathcal{B}$.

We proceed to bound convergence rate in terms of path congestion. In particular, for a Markov chain Z_t, $t = 0, 1, \ldots$ on state space $\mathcal{S}$ and transition matrix $P' = \frac{1}{2}(P+I)$, let $\vec{\pi}$ denote the stationary distribution of Z_t. Let $w_{ij} = \pi_i p_{ij}$ measure the *ergodic flow* between states i and j, and suppose that the ergodic flow w_{ij} is the same between any two states for which $p_{ij} \neq 0$, so let w denote this ergodic flow. Assuming that random paths have been defined between any pair A and B of states of the Markov chain (so that these paths use edges (i, j) such that $p_{ij} \neq 0$), let l denote an upper bound on the expected length of the path from A to B. Finally, among all paths that have been defined between pairs of states choose one path at random according to the distribution that assigns probability $\pi_A \pi_B$ to the path from state A to state B. Let L be an upper bound on the probability that the chosen path contains a specific edge (i,j). In Appendix 2 we show:

$$d(t) \leq \frac{1}{2}\left(|\mathcal{S}|\frac{\pi_{\max}}{\pi_{\min}}\right)^{\frac{1}{2}}\left(1-\frac{w}{2Ll}\right)^{\frac{t}{2}} \tag{5}$$

For the natural random walk X_t on any balanced matroid $\mathcal{M}(S,\mathcal{B})$ we have $w=1/|\mathcal{B}|mn$, $l=2n$ by Corollary 4.4, and $L=1/2|\mathcal{B}|$ by Corollary 4.2 and the definition of L. Hence

Theorem 5.1 *For the natural random walk X_t on any balanced matroid $\mathcal{M}(S,\mathcal{B})$ the total variation distance $d(t)$ can be bounded by ϵ in time $t=\Omega(n\log m+\log\epsilon^{-1})n^2m$.*

For the modified random walk, paths between bases are defined as for the bases-exchange graph in Section 4, except that near-bases are used to simulate the edges of the bases-exchange graph in the obvious way. The expected number of such paths through an edge can be bounded by the expected number of paths through a vertex of the bases-exchange graph; the latter expectation is by Corollary 4.5 $2n|\mathcal{B}|$. Paths whose endpoints include near-bases are defined by first choosing a neighboring basis uniformly at random, and then choosing a random path between bases, thus reducing path congestion due to paths with arbitrary endpoints to path congestion due to paths whose endpoints are only bases. From the above observations and the definition of L it is not hard to verify that for the modified random walk $L=2n/|\mathcal{B}|$. Furthermore, $w=1/2n|\mathcal{B}|$ and $l=2n$. Thus

Theorem 5.2 *For the modified random walk Y_t on any balanced matroid $\mathcal{M}(S,\mathcal{B})$ the total variation distance $d(t)$ can be bounded by ϵ in time $t=\Omega(n\log m+\log\epsilon^{-1})n^3$.*

Consider now the case where each element may have a very large number of elements parallel to it. These parallel elements may cause a large increase in the state space size $|\mathcal{S}|$. To avoid this increase, collapse all parallel elements, and view the number of parallel elements as a weight on the resulting element after collapsing; the probability of any given basis is now proportional to the product of the weights of elements in it, multiplied by n for bases, and by the sum of the weights of elements in the associated cut for near-bases. This transformation reduces the size of the state space, but may increase the factor $\pi_{\max}/\pi_{\min}$ due to the weight differences. To reduce this quantity, we choose as a starting point for the random walk a basis of maximum weight, so that $\pi_{\min}$ can be replaced by $\pi_{\max-\text{basis}}$ and $\pi_{\max}/\pi_{\max-\text{basis}}$ is at most m/n, where m is the number of elements after parallel elements have been collapsed. Therefore for the modified random walk, unlike the natural random walk, the number of parallel elements does not affect the running time. A basis of maximum weight needed to start the random walk can be found by greedily selecting elements of largest weight while preserving independence.

Theorem 5.3 *The bounds for the modified random walk on any balance matroid remain the same after an arbitrary number of parallel elements are added while maintaining balance.*

6 Better Path Congestion, Decomposition, Symmetry, Series Parallel

In the graph consisting of bases and near-bases for a balance matroid, we have constructed paths between each pair of bases, and bounded the load on edges using a bound of about $n|\mathcal{B}|$ on the load of basis vertices. Since each basis has n incident edges, one for each element of the basis, one might expect the load on each edge to be bounded by $|\mathcal{B}|$. If a bound of $\alpha|\mathcal{B}|$ holds for some α, the n^3 factor of convergence rate for the modified random walk can be reduced to αn^2. We shall now examine conditions that can yield such bounds.

Consider the bipartite graph with edges joining bases and near-bases. Given an element e, one can find a fractional matching between all bases, each with weight $|\mathcal{B}_e|$, and all near-bases $\mathcal{N}_e$ such that e belongs to their associated cut, each with weight $|\mathcal{B}|$. To see this, note that the bases in $\mathcal{B}_e$ can be matched to the near-bases in $\mathcal{N}_e$, simply by removing e from each of them. The weight assigned to this edge is $|\mathcal{B}_e|$, satisfying the weight constraint for each basis in $\mathcal{B}_e$, and leaving at each near-basis in $\mathcal{N}_e$ a weight of $|\mathcal{B}_{\overline{e}}|$ to be satisfied using the edges joining $\mathcal{B}_{\overline{e}}$ to $\mathcal{N}_e$. If we view each near-basis in $\mathcal{N}_e$ as representing a corresponding basis in $\mathcal{B}_e$, then this bipartite graph is simply the bases-exchange graph associated with e, and we can infer that a fractional matching giving weight $|\mathcal{B}_e|$ to each basis in $\mathcal{B}_{\overline{e}}$ and weight $|\mathcal{B}_{\overline{e}}|$ to each near-basis in $\mathcal{N}_e$ does exist.

Given a near-basis N in $\mathcal{N}_e$ and an element $f \neq e$ that belongs to the cut of N, denote by w_{efN} the weight assigned to the edge joining N to the basis $N \cup \{f\}$ by the above fractional matching for e. For convenience, if $U = N \cup \{e, f\}$, we write $w_{efU} = w_{efN}$. For all other choices of f,N,U, we set $w_{efN} = 0$, $w_{efU} = 0$. The definition of fractional matchings requires $w_{efN} \geq 0$, and in addition, the two conservation equations

$$\sum_f w_{efN} = |\mathcal{B}_{\overline{e}}| \text{ for given } e \text{ and } N \text{ with } e \in D_N,$$

$$\sum_f w_{efU} = |\mathcal{B}_e| \text{ for given } e \text{ and } U \text{ with } e \in C_U.$$

We have thus constructed m fractional matchings from $\mathcal{B}$ to $\mathcal{N}$, one for each choice of e. Suppose that we now add these m fractional matchings, by adding, for each basis, each near-basis, and each edge, its weight from each of the fractional matchings. The resulting weight for each basis is $\sum_e |\mathcal{B}_e| = n|\mathcal{B}|$, and for each near-basis N it is $\sum_{e \in D_N} |\mathcal{B}| = |D_N| \cdot |\mathcal{B}|$. In both cases, the weight is equal to the degree of the basis or near-basis times $|\mathcal{B}|$. Note that there is a very simple fractional matching satisfying these weight constraints, namely the *uniform fractional matching* that assigns weight $|\mathcal{B}|$ to each edge joining a basis and a near-basis. We say that a balanced matroid has the *exact decomposition property* if the uniform fractional matching can indeed be obtained as the sum of m fractional matchings, one for each element e. In terms of weights, this means that the weights assigned to a particular edge joining N and $N \cup \{f\}$ must add to $|\mathcal{B}|$, and since the weight of this edge for the fractional matching with $e = f$ is $|\mathcal{B}_f|$, we must have

$$\sum_e w_{efN} = |\mathcal{B}_{\overline{f}}| \text{ for given } f \text{ and } N \text{ with } f \in D_N.$$

Note that this property must hold if *symmetry* holds, namely, if $w_{efN} = w_{feN}$ for all e, f, N. If symmetry holds, then exact decomposition follows from the first conservation equation. One can also consider the weaker *α-decomposition* property, which only requires

$$\sum_e w_{efN} \leq \alpha|\mathcal{B}|$$

for some α. Note that this property always holds for $\alpha = n$, since the sum weight of the basis $N \cup \{f\}$ is $n|\mathcal{B}|$. The proofs of the following lemma and theorem are given in Appendix 3.

Lemma 6.1 *If a balanced matroid and all its minors satisfy the α-decomposition property, then its path congestion is at most $(\alpha + 1)|\mathcal{B}|$ and the n^3 factor of the convergence rate for the modified random walk reduces to αn^2.*

The main obstacle in obtaining bounds for the values of α achievable for α-decompositions is the fact that the fractional matchings are in general not explicitly known. We now give a case where such an explicit construction can be obtained.

Theorem 6.2 *If a regular matroid satisfies $D(e)D(f)\Delta_{ef} \geq 0$ for all cuts D and elements e, f, then symmetry and exact decomposition hold. The regular matroids such that exact decomposisition (as well as symmetry) holds for the matroid and all its minors are precisely the graphic matroids for which each biconnected component is either K_4 or a series-parallel graph.*

We do not know whether an α constant exist such that the α-decomposition property holds for all balanced matroids.

7 Conclusions and Open Problems

We have introduced a new class of matroids, the balanced matroids. This class contains some important matroids, in particular the graphic matroids, whose bases are the spanning trees of a given graph, and more generally all regular matroids. We have then shown that on balanced matroids, one can efficiently perform almost uniform sampling of bases; this in turn leads to efficient approximate counting of bases in view of self-reducibility for balanced matroids.

The results on balanced matroids extend known results for graphic and for regular matroids and improve the known bounds on the rate of convergence for the natural random walk on bases. A new modified random walk on bases and near-bases shows that the rate of convergence for sampling schemes can be made to depend essentially only on the size of basis, and only logarithmically on the size of the underlying ground set, while remaining unaffected by the presence of parallel elements.

The basic technique involves a new construction of random paths in the bases-exchange graph, based on the existence of certain fractional matchings that can be inferred from the negative correlation property, and introduces a notion of decomposition that can lead to better bounds when the quality of a decomposition can be measured by examining the underlying fractional matchings.

It remains open to determine whether other well-known classes of matroids are balanced. A class of matroids closely related to matchings consists of the transversals, and their minor-closed extension, the gammoids. It is of particular interest to determine whether the truncations of balanced matroids remain balanced, especially for graphic matroids and their duals, in connection with the network reliability problem. For classes of matroids that are not always balanced, such as the full range of binary matroids, the techniques introduced on balanced matroids may provide a fruitful approach to sampling.

Acknowledgements

We are grateful to Paul Seymour for his insightful comments; in particular, Paul pointed out the negative correlation property. We also wish to thank Alistair Sinclair, Umesh Vazirani, and Peter Winkler for several valuable discussions.

References

[1] D. Aldous, "On the Markov chain simulation method for uniform combinatorial distributions and simulated annealing", *Probability in Eng. and Inf. Sci.* 1, 1987, pp 33-46.

[2] D. Aldous, "The random walk construction for spanning trees and uniform labelled trees," *SIAM J. of Disc. Math.* 3, 1990, pp450-465.

[3] A.Z. Broder, "How hard is it to marry at random? (On the approximation of the permanent)," *STOC* 1986, pp 50-58.

[4] A.Z. Broder, "Generating random spanning trees," *FOCS* 1989, pp 442-447.

[5] R.L. Brooks, C.A.B. Smith, A.H. Stone, and W.T. Tutte, "The dissection of rectangles into squares," *Duke Math. J.* 7, 1940, 312-340.

[6] P. Diaconis and D.Strock, "Geometric bounds for eigenvalues of Markov chains," *Annals of Applied Probability*, Vol. 1, 36-61, 1991.

[7] M. Dyer and A. Frieze, "Random walks on unimodular zonotopes," 1991, preprint.

[8] M. Dyer, A. Frieze, and R. Kannan, "A random polynomial time algorithm for estimating volumes of convex bodies,"*STOC* 1989, pp 375-381.

[9] J. Edmonds, "Submodular functions, matroids and certain polyhedra," *Combinatorial structures and their applications*, Proceedings Calgary International Conference, 1969

[10] G.N. Frederickson, "Ambivalent data structures for dynamic 2-edge connectivity and k smallest spanning trees," *FOCS* 1991, pp 632-641.

[11] C. Holtzmann, and F. Harary, On the tree graph of a matroid, *SIAM J. Applied Math.* 25, pp 187-193.

[12] F. Jaeger, D.L. Vertigan, and D.J.A. Welsh, "On the comutational complexity of Jones and Tutte polynomials," *Math. Proc. Camb. Phil. Soc.* 1990, 108, pp 35-53.

[13] M.R. Jerrum and A. Sinclair, "Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved," *STOC* 1988, pp 235-243.

[14] M.R. Jerrum, L.G. Valiant, and V.V. Vazirani, "Random generation of combinatorial structures from a uniform distribution," *Theoretical Computer Science* 43, 1986, pp 169-188.

[15] L. Lovász and M. Simonovis, "The mixing rate of M arkov chains, An isoperimetric inequality, and computing the volume," *FOCS* 1990, pp 346-354.

[16] M. Mihail, "Conductance and convergence of Markov chains (A combinatorial Treatment of expanders)," *FOCS* 1989, pp 526-531.

[17] M. Mihail and M. Sudan, "Connectivity properties of matroids", TR-89, U.C. Berkeley.

[18] M. Mihail, and U. Vazirani, "On the magnification of 0-1 polytopes," TR-3-89, Aiken Comp. Lab., Harvard

[19] D. Naddef and W. Pulleyblank, "Hamiltonicity in 0-1 polyhedra," *Journal of Combinatorial Theory*, Series B 37, 1984, pp 41-52.

[20] A. Sinclair, "Improved bounds for mixing rates of Markov chains and multicomodity flows," *Proc. of Latin* 1992, to appear.

[21] P. Seymour, personal communication.

[22] P. Seymour and D.J.A. Welsh, "Combinatorial applications of an inequality from statistical mechanics," *Math. Proc. Camb. Phil. Soc.* 1975, 77, pp 485-495.

[23] A. Sinclair and M.R. Jerrum, "Approximate counting, uniform generation and rapidly mixing markov chains," *Information and Computing*, to appear.

[24] L. G. Valiant, "The complexity of enumeration and reliability problems," *SIAM Journal on Computing* 8, 1979, pp 410-421.

[25] D. Welsh, Matroid Theory, *Academic Press* , 1976.

[26] D.J.A. Welsh, "On the Tutte map," Lecture Notes, DIMACS May '91/Dagstuhl June 91.

Appendix 1: Expansion Proofs

We give here the proof of Theorem 3.4, establishing cutset expansion 1 for balanced matroids.
Proof. (Mihail and Sudan [17]). Let $\mathcal{A} \subseteq \mathcal{B}$ such that $|\mathcal{A}| \leq |\mathcal{B}|/2$. We wish to show that the cutset of $\mathcal{A}$, $C(\mathcal{A})$, has cardinality $|C(\mathcal{A})| \geq |\mathcal{A}|$. The argument proceeds inductively. Let $\mathcal{A}_e = \mathcal{A} \cap \mathcal{B}_e$ and $\mathcal{A}_{\overline{e}} = \mathcal{A} \cap \mathcal{B}_{\overline{e}}$. Let C_e be the set of edges incident to $\mathcal{A}_e$ and $\mathcal{B}_e \setminus \mathcal{A}_e$. Let $C_{\overline{e}}$ be the set of edges incident to $\mathcal{A}_{\overline{e}}$ and $\mathcal{B}_{\overline{e}} \setminus \mathcal{A}_{\overline{e}}$. Finally let $C_{e\overline{e}}$ be the set of edges incident to $\mathcal{A}_e$ and $\mathcal{B}_{\overline{e}} \setminus \mathcal{A}_{\overline{e}}$. Obviously $|C(\mathcal{A})| \geq |C_e| + |C_{\overline{e}}| + |C_{e\overline{e}}|$.

If $|\mathcal{A}_e| \leq |\mathcal{B}_e|/2$ and $|\mathcal{A}_{\overline{e}}| \leq |\mathcal{B}_{\overline{e}}|/2$ then $|C_e| \geq |\mathcal{A}_e|$ and $|C_{\overline{e}}| \geq |\mathcal{A}_{\overline{e}}|$ by the inductive hypothesis, and hence $|C(\mathcal{A})| \geq |\mathcal{A}|$.

If $|\mathcal{A}_e| \geq |\mathcal{B}_e|/2$ and $|\mathcal{A}_{\overline{e}}| \leq |\mathcal{B}_{\overline{e}}|/2$ then the inductive hypothesis gives (i)$|C_e| \geq |\mathcal{B}_e| - |\mathcal{A}_e|$ and (ii)$|C_{\overline{e}}| \geq |\mathcal{A}_{\overline{e}}|$. By Theorem 3.1 we have $|\Gamma_e(\mathcal{A}_e)|/|\mathcal{B}_{\overline{e}}| \geq |\mathcal{A}_e|/|\mathcal{B}_e| \geq 1/2$, hence some $\mathcal{H} \subseteq \Gamma_e(\mathcal{A}_e)$ satisfies $\mathcal{H} \cap \mathcal{A}_{\overline{e}} = \emptyset$ and $|\mathcal{H}| = |\mathcal{B}_{\overline{e}}|/2 - |\mathcal{A}_{\overline{e}}|$. Let $C_{e\overline{e},\mathcal{H}}$ be the set of edges incident to $\mathcal{A}_e$ and to $\mathcal{H}$; clearly (iii)$|C_{e\overline{e},\mathcal{H}}| \geq |\mathcal{H}| = |\mathcal{B}_{\overline{e}}|/2 - |\mathcal{A}_{\overline{e}}|$. Now let $\mathcal{H}' = \mathcal{B}_{\overline{e}} \setminus (\mathcal{H} \cup \mathcal{A}_{\overline{e}})$, and by choice of $\mathcal{H}$, $|\mathcal{H}'| = |\mathcal{B}_{\overline{e}}|/2$. Furthermore by Theorem 3.1 $|\Gamma_e(\mathcal{H}')|/|\mathcal{B}_e| \geq |\mathcal{H}'|/|\mathcal{B}_{\overline{e}}| \geq 1/2$, so $|\Gamma_e(\mathcal{H}') \cap \mathcal{A}_e| \geq |\mathcal{A}_e| - |\mathcal{B}_e|/2$. Let $C_{e\overline{e},\mathcal{H}'}$ be the set of edges incident to $\mathcal{A}_e$ and to $\mathcal{H}'$; clearly (iv)$|C_{e\overline{e},\mathcal{H}'}| \geq |\Gamma_e(\mathcal{H}') \cap \mathcal{A}_e| \geq |\mathcal{A}_e| - |\mathcal{B}_e|/2$. Now (i), (ii), (iii), and (iv) yield $|C(\mathcal{A})| \geq |C_e| + |C_{\overline{e}}| + |C_{e\overline{e},\mathcal{H}}| + |C_{e\overline{e},\mathcal{H}'}| \geq |\mathcal{B}|/2 \geq |\mathcal{A}|$.

The final case where $|\mathcal{A}_e| \leq |\mathcal{B}_e|/2$ and $|\mathcal{A}_{\overline{e}}| \geq |\mathcal{B}_{\overline{e}}|/2$ can be treated analogously in the obvious way, and this completes the proof of the theorem. □

An alternative proof involving fractional matchings and the random paths from Section 4 is as follows. For $\mathcal{A} \subseteq \mathcal{B}$, there are $|\mathcal{A}| \cdot |\overline{\mathcal{A}}|$ paths constructed from $\mathcal{A}$ to $\overline{\mathcal{A}}$. Each path leaves $\mathcal{A}$ through some edge in $C(\mathcal{A})$. The expected number of paths leaving through some specific edge is at most $|\mathcal{B}|/2$, so the total number of paths leaving $\mathcal{A}$ is at most $|C(\mathcal{A})| \cdot |\mathcal{B}|/2$, and therefore $|C(\mathcal{A})| \geq 2|\mathcal{A}| \cdot |\overline{\mathcal{A}}|/|\mathcal{B}| \geq \min(|\mathcal{A}|, |\overline{\mathcal{A}}|)$.

Appendix 2: Convergence Bounds

Proof of (5): Let

$$h_i(t) = \Pr[Z_t = i] - \pi_i$$

and we shall use h_i as short for $h_i(t)$. Note that $\sum_i h_i = 0$, and moreover that in stationarity all h_i's are zero. It is useful to think of discrepancy from stationarity in terms of the discrepancies of the h_i's. We wish to bound the total variation distance

$$d(t) = \frac{1}{2}\sum_i |h_i| \leq \frac{1}{2}|\mathcal{S}|^{\frac{1}{2}} \left(\sum_i h_i^2\right)^{\frac{1}{2}} \leq \frac{1}{2}\left(|\mathcal{S}|\pi_{\max}\right)^{\frac{1}{2}} \left(\sum_i \frac{h_i^2}{\pi_i}\right)^{\frac{1}{2}} \quad . \tag{6}$$

Let $d'(t) = \sum_i \frac{h_i^2}{\pi_i}$, and recall that $w_{ij} = \pi_i p_{ij}$ is the ergodic flow between states i and j. Note that $\pi_i = \sum_j w_{ij} = \sum_j w_{ji}$. Suppose that the ergodic flow w_{ij} is the same between any two states for which $p_{ij} \neq 0$, so let w denote this ergodic flow.

Applying the techniques that have appeared in [16] we get:

$$\begin{aligned} d'(t) &= \sum_i \frac{h_i^2}{\pi_i} \\ &= \sum_i \left(\frac{h_i}{\pi_i}\right)^2 \pi_i \end{aligned}$$

$$
\begin{aligned}
&= \sum_i \left(\frac{h_i}{\pi_i}\right)^2 \frac{1}{2}\left(\sum_j w_{ij} + \sum_j w_{ji}\right) \\
&= \frac{1}{2}\sum_i \sum_j w_{ji}\left(\left(\frac{h_j}{\pi_j}\right)^2 + \left(\frac{h_i}{\pi_i}\right)^2\right)
\end{aligned}
$$

Remark: Intuitively, the above equation attributes discrepancy from stationarity to discrepancy of the h_i's at the endpoints of (j, i)'s such that $p_{ji} \neq 0$.

Furthermore,

$$\vec{h}(t+1) = (\vec{h}(t)+\vec{\pi})P' - \vec{\pi}P' = \vec{h}(t)P'$$

Hence

$$
\begin{aligned}
h_i(t+1) &= \frac{1}{2}\left(\sum_j \pi_{ji} h_j + h_i\right) \\
&= \frac{1}{2}\left(\sum_j \frac{w_{ji}}{\pi_j} h_j + \frac{1}{\pi_i}\sum_j w_{ji} h_i\right) \\
&= \frac{1}{2}\sum_j w_{ji}\left(\frac{h_j}{\pi_j} + \frac{h_i}{\pi_i}\right)
\end{aligned}
$$

Remark: Intuitively, the above equality states that each step of the Markov chain results in an averaging of the discrepancies along the endpoints of (j, i)'s such that $p_{ji} \neq 0$; thus the total discrepancy should decrease.

Combining the above we further get

$$
\begin{aligned}
d'(t) - d'(t+1) &= \frac{1}{2}\sum_i \sum_j w_{ji}\left(\left(\frac{h_j}{\pi_j}\right)^2 + \left(\frac{h_i}{\pi_i}\right)^2\right) - \frac{1}{4}\sum_i \frac{1}{\pi_i}\left(\sum_j w_{ji}\left(\frac{h_j}{\pi_j} + \frac{h_i}{\pi_i}\right)\right)^2 \\
&\geq \frac{1}{4}\sum_i \sum_j 2w_{ji}\left(\left(\frac{h_j}{\pi_j}\right)^2 + \left(\frac{h_i}{\pi_i}\right)^2\right) - \frac{1}{4}\sum_i \sum_j w_{ji}\left(\frac{h_j}{\pi_j} + \frac{h_i}{\pi_i}\right)^2 \\
&= \frac{1}{4}\sum_i \sum_j w_{ij}\left(\frac{h_j}{\pi_j} - \frac{h_i}{\pi_i}\right)^2 \\
&\geq \frac{w}{4}\sum_{ij:p_{ij}\neq 0}\left(\frac{h_j}{\pi_j} - \frac{h_i}{\pi_i}\right)^2 \quad . \qquad (7)
\end{aligned}
$$

Remark: Intuitively, the above bound states that the averaging of the discrepancies along the edges is more effective if it involves (j, i)'s such that $p_{ji} \neq 0$ with significantly different disceparcies at the endpoints of j and i.

Recall that paths have been defined between any pair A and B of states of the Markov chain (so that these paths use edges (i, j) such that $p_{ij} \neq 0$). Let P_{AB} denote the chosen path from A to B and the event that P_{AB} was chosen, and let $ij \in P_{AB}$ denote the event that edge (i, j) was used in the path from A to B. Let l_{AB} denote the length of the path from A to B, and recall that l denotes an upper bound on the expected length of l_{AB}. Finally, recall that among all paths that have been defined between pairs of states we choose one path at random according to the distribution that assigns probability $\pi_A\pi_B$ to the path from state A to state B. Recall that L is an upper bound on the probability that the chosen path contains a specific edge (i,j). Then the bounding of (7) continues as follows, along the lines of [6] (see also [20], except that we are using expected rather than maximum path length here):

$$\begin{aligned}
d'(t)-d'(t+1) &\geq \frac{w}{4}\sum_{ij:p_{ij}\neq 0}\left(\frac{h_j}{\pi_j}-\frac{h_i}{\pi_i}\right)^2 \frac{\sum_A\sum_B \Pr[ij\in P_{AB}]\pi_A\pi_B}{L} \\
&= \frac{w}{4L}\sum_A\sum_B \pi_A\pi_B \sum_{P_{AB}}\Pr[P_{AB}]\sum_{ij\in P_{AB}}\left(\frac{h_j}{\pi_j}-\frac{h_i}{\pi_i}\right)^2 \\
&\geq \frac{w}{4L}\sum_A\sum_B \pi_A\pi_B \sum_{P_{AB}}\Pr[P_{AB}]\frac{1}{l_{AB}}\left(\frac{h_A}{\pi_A}-\frac{h_B}{\pi_B}\right)^2 \qquad \text{by convexity,} \\
&= \frac{w}{4L}\sum_A\sum_B \pi_A\pi_B \left(\frac{h_A}{\pi_A}-\frac{h_B}{\pi_B}\right)^2 E\left(\frac{1}{l_{AB}}\right) \\
&\geq \frac{w}{4Ll}\sum_A\sum_B \pi_A\pi_B \left(\frac{h_A}{\pi_A}-\frac{h_B}{\pi_B}\right)^2 \qquad \text{by convexity,} \\
&= \frac{w}{2Ll}\sum_A \frac{h_i^2}{\pi_i} = \frac{w}{2Ll}d'(t)
\end{aligned}$$

where the last bound follows by expanding the squares and recalling that $\sum_i h_i = 0$ and $\sum_i \pi_i = 1$.
Remark: Intuitively the above bound can be, roughly, understood as follows: If paths are defined between any two states, then discrepancy can be attributed to discrepancy at the endpoints of the path, and in turn, to discrepancy at the enpoints of each edge along the path (along a breakup that favors short paths). Since there is no path congestion, pairs $(i,j): p_{ji}\neq 0$ with significantly different discrepancies at their endpoints will appear often enough to result in a large drop of the total discrepancy.

Hence

$$d'(t) \leq \left(1-\frac{w}{2Ll}\right)^t d'(0).$$

Finally using all the above and noting that $d'(0)\leq 1/\pi_{\min}$, we get

$$d(t) \leq \frac{1}{2}\left(|\mathcal{S}|\frac{\pi_{\max}}{\pi_{\min}}\right)^{\frac{1}{2}}\left(1-\frac{w}{2Ll}\right)^{\frac{t}{2}}.$$

Appendix 3: Decomposition Bounds

We give the proof of Lemma 6.1, stating that the α-decomposition property implies path congestion of at most $(\alpha+1)|\mathcal{B}|$, reducing the convergence rate accordingly.
Proof. Consider the first step of all paths. Let e be the element chosen for the first step, and consider an edge joining N to $N\cup\{f\}$. If $e=f$, then the load of this edge is $|\mathcal{B}_{\overline{f}}|$. If $e\neq f$, then the load is determined by the fractional matching, and it is precisely w_{efN}. Since e is chosen uniformly at random from all m elements, the expected load is

$$\frac{1}{m}(|\mathcal{B}_{\overline{f}}|+\sum_e w_{efN}) \leq \frac{\alpha+1}{m}|\mathcal{B}|.$$

At the $(i+1)$th step of the paths, if f is one of the i elements $e_1, e_2, \ldots, e_i$ used for the first i steps, then the load for the edge under consideration is zero. This case occurs with probability i/m. Otherwise, the edge joining N to $N \cup \{f\}$ will be involved in the fractional matching for some $e = e_{i+1}$ after the chosen $e_1, e_2, \ldots, e_i$ have been contracted or deleted, depending on whether they occur in N or not. This minor has only $m-i$ elements, so the expected load is now bounded by $\frac{\alpha+1}{m-i}|\mathcal{B}|$ in this case. Since this second case occurs with probability $(m-i)/i$, the expected load for this step is again bounded by $\frac{\alpha+1}{m}|\mathcal{B}|$. Adding this quantity over all m steps of the paths, we obtain a bound of $(\alpha+1)|\mathcal{B}|$ on the path congestion, thus reducing the n^3 factor for the convergence rate to αn^2. $\square$

We now give the proof of Theorem 6.2 showing that the condition $D(e)D(f)\Delta_{ef} \geq 0$ for regular matroids implies symmetry and exact decomposition, and characterizing the regular matroids satisfying exact decomposition as those graphic matroids whose biconnected components are either K_4 or a series-parallel graph.

Proof. We set $w_{efN} = D_N(e)D_N(f)\Delta_{ef}$. These quantities are nonnegative by assumption, and satisfy the conservation laws. If $e \in D_N$, then

$$\begin{aligned}\sum_f w_{efN} &= -\sum_{f \neq e} D_N(e)D_N(f) \sum_U C_U(e)C_U(f) \\ &= -\sum_U D_N(e)C_U(e) \sum_{f \neq e} D_N(f)C_U(f) \\ &= \sum_U (D_N(e)C_U(e))^2 = \sum_{U: e \in C_U} 1 = |\mathcal{B}_{\overline{e}}|\end{aligned}$$

Similarly, if $e \in C_U$, then

$$\begin{aligned}\sum_f w_{efU} &= -\sum_{f \neq e} C_U(e)C_U(f) \sum_N D_N(e)D_N(f) \\ &= -\sum_N C_U(e)D_N(e) \sum_{f \neq e} C_U(f)D_N(f) \\ &= \sum_N (C_U(e)D_N(e))^2 = \sum_{N: e \in D_N} 1 = |\mathcal{B}_e|.\end{aligned}$$

Furthermore, the symmetry property $w_{efN} = w_{feN}$ holds by the definition of Δ_{ef}, and thus exact decomposition holds as well.

There are two dual graphic matroids that do not satisfy symmetry or decomposition, namely K_4^1 and K_4^2, obtained from K_4 by replacing a single edge with either (1) two parallel edges, or (2) two edges in series. If both are excluded as minors, then the duals of K_5 and $K_{3,3}$ are also excluded, so the regular matoid must be graphic [25]; furthermore, a biconnected component containing K_4 as a minor must then in fact be K_4, and a biconnected component has no K_4 as minor if and only if it is a series-parallel graph. In a series-parallel graph, one cannot find two cycles through two given edges e, f that traverse e in the same direction but traverse f in opposite directions; otherwise, the two cycles would give a K_4 minor. It follows that $C(e)C(f)$ cannot be nonzero and have different signs for two different cycles C. Therefore $D_N(e)D_N(f)\Delta_{ef} = -C_U(e)C_U(f)\Delta_{ef} \geq 0$, satisfying the conditions for symmetry and exact decomposition. In a clique K_n, one has $\Delta_{ef} = 0$ if e, f do not share any vertices, because there is an isomorphism of K_n that leaves e fixed but swaps the endpoints of f. If e shares an endpoint with f, then the direction used by a cycle through e, f in traversing e determines the direction used in traversing f, so symmetry and exact decomposition hold as before for series-parallel graphs. They hold in particular for K_4, and since all its minors are series-parallel graphs, the two properties hold for them as well. $\square$