

A polynomial bound on the mixing time of a Markov chain for sampling regular directed graphs

Catherine Greenhill

School of Mathematics and Statistics
The University of New South Wales
Sydney NSW 2052, Australia

`csg@unsw.edu.au`

Submitted: May 11, 2011; Accepted: Dec 4, 2011; Published: Dec 19, 2011

Mathematics Subject Classification: 05C20, 60J10

Abstract

The switch chain is a well-known Markov chain for sampling directed graphs with a given degree sequence. While not ergodic in general, we show that it is ergodic for regular degree sequences. We then prove that the switch chain is rapidly mixing for regular directed graphs of degree d , where d is any positive integer-valued function of the number of vertices. We bound the mixing time by bounding the eigenvalues of the chain. A new result is presented and applied to bound the smallest (most negative) eigenvalue. This result is a modification of a lemma by Diaconis and Stroock [9], and by using it we avoid working with a lazy chain. A multicommodity flow argument is used to bound the second-largest eigenvalue of the chain. This argument is based on the analysis of a related Markov chain for undirected regular graphs by Cooper, Dyer and Greenhill [6], but with significant extension required.

1 Introduction

Directed graphs are natural combinatorial objects which are used to model systems in many areas including biology (for example [5, 16]), the social sciences (for example [25, 26]) and computer science (for example [12, 18]). In this paper we consider the problem of sampling directed graphs with a given degree sequence.

For graph-theoretic terminology not introduced here, see [1]. A directed graph (digraph) $G = (V, A)$ consists of a vertex set $V = V(G)$ and an arc set

$$A = A(G) \subseteq \{(v, w) \in V \times V \mid v \neq w\}.$$

Note that digraphs as defined here are simple, which means that they contain no loops and no multiple arcs.

The arc (v, w) is drawn as an arrow from v to w . We refer to v as the *tail* and w as the *head* of the arc. For a vertex v , the *out-degree* $d^+(v)$ of v is the number of arcs with tail v . Similarly, the *in-degree* $d^-(v)$ of v is the number of arcs with head v . For a positive integer d , if $d^+(v) = d^-(v) = d$ for all vertices $v \in V$ then we say that the digraph G is *d-regular* (or *d-in, d-out*).

Let $d = d(n) \geq 1$ be a sequence of positive integers, and let $\Omega_{n,d}$ be the set of all simple d -regular digraphs on the vertex set $[n] = \{1, \dots, n\}$. The *configuration model* of Bollobás [4] (adapted for directed graphs) gives an expected polynomial-time uniform sampling algorithm for $\Omega_{n,d}$ when $d = O(\sqrt{\log n})$.

There is a one-to-one correspondence between $\Omega_{n,d}$ and the set of all d -regular bipartite graphs on $\{1, 2, \dots, n\} \cup \{n+1, n+2, \dots, 2n\}$ with no edges in common with the perfect matching $\{\{j, n+j\} : j = 1, \dots, n\}$. The probability that a d -regular bipartite graph on the given vertex bipartition has no edges in common with this perfect matching is asymptotic to e^{-d} whenever $d = o(n^{1/3})$, by [20, Theorem 4.6]. This probability is polynomially small when $d = O(\log n)$. McKay and Wormald's algorithm [21] for sampling d -regular graphs runs in expected polynomial time for $d = O(n^{1/3})$, and hence gives rise to an expected polynomial-time algorithm for uniformly sampling elements of $\Omega_{n,d}$ when $d = O(\log n)$.

The set of all 1-regular digraphs is in one-to-one correspondence with the set of all derangements of $[n]$, and here the configuration model corresponds to repeatedly sampling uniform permutations of n until one is obtained without fixed points. The proportion of permutations which are derangements tends to $1/e$, so this algorithm has linear expected running time. Other algorithms for uniformly sampling derangements in linear expected time but an improved constant have been proposed, for example [19].

We know of no expected polynomial-time uniform sampling algorithm for regular digraphs other than those mentioned above. Hence we turn our attention to the problem of obtaining approximately uniform samples from $\Omega_{n,d}$ using a Markov chain. (Some Markov chain definitions are given in Section 1.1; for others, see [23].)

There is a very natural Markov chain for digraphs which has arisen in many contexts, which we will call the *switch chain*. A transition of the switch chain is performed by randomly choosing two distinct arcs and exchanging their heads, if the two arcs are non-incident and if the resulting digraph does not contain any multiple arcs. See Figure 1 for a precise description of the transition procedure of the chain. A transition of the switch chain is called *switching along an alternating rectangle* by Rao, Jana and Bandyopadhyay [24]; we will simply call it a *switch*. Similar transformations were used by Ryser [27] to study 0-1 matrices. Besag and Clifford [3] defined a related chain for sampling 0-1 matrices with given row and column sums, while Diaconis and Sturmfels [10] used a similar chain to sample contingency tables.

Rao, Jana and Bandyopadhyay [24] showed that the switch chain is not irreducible for general degree sequences. (However, they mention that degree sequences for which the switch chain is not irreducible are "rather rare".) For completeness, we prove in Lemma 2.2 that the switch chain is irreducible for regular digraphs. This also follows from the existence of the multicommodity flow defined in Sections 3, 4.

The switch chain is aperiodic for $d \geq 1$, as we prove in Lemma 2.1.

In their empirical study of methods for generating directed graphs with given degree sequences, Milo et al. [22] wrote that the switch chain “works well but, as with many Markov chain methods, suffers because in general we have no measure of how long we need to wait for it to mix properly”. Our main result, Theorem 1.1, partially answers this point by providing the first rigorous polynomial bound on the mixing time of the switch chain, in the special case of regular digraphs.

Theorem 1.1. *Let $\Omega_{n,d}$ be the set of all d -regular digraphs on the vertex set $\{1, \dots, n\}$, where $d = d(n)$ is any integer-valued function which satisfies $1 \leq d(n) \leq n - 1$ for all $n \geq 4$. Let $\tau(\epsilon)$ be the mixing time of the Markov chain \mathcal{M} with state space $\Omega_{n,d}$ and transition procedure given by Figure 1, for $d \geq 1$. Then*

$$\tau(\epsilon) \leq 50 d^{25} n^9 (dn \log(dn) + \log(\epsilon^{-1})).$$

Our proof of this result has two parts. To avoid using a lazy chain (which stays where it is at each step, with probability at least $\frac{1}{2}$) we prove and apply a new result which can be used to bound the smallest eigenvalue of an ergodic reversible Markov chain. This new bound is based on Diaconis and Stroock [9, Proposition 2] and inspired by Sinclair [28, Theorem 5]. To bound the second-largest eigenvalue of the chain we adapt the multicommodity flow analysis given in [6] for the undirected case. While some parts of the proof are very similar to [6], significant extra technical difficulties arise in the directed setting. We expect that the bound on the mixing time given in Theorem 1.1 is far from tight, but proving a substantially tighter bound seems beyond the reach of known proof techniques.

The *flip chain* is a Markov chain which performs a restricted set of switches, designed to ensure that the underlying digraph never becomes disconnected. The flip chain for undirected graphs was described in [17], and proposed as a self-randomizing mechanism for peer-to-peer networks. The mixing time of the flip chain for regular undirected graphs was analysed in [7, 11], building on the multicommodity flow analysis of the switch chain [6]. We expect that Theorem 1.1 can be used to show that the flip chain for digraphs is rapidly mixing for regular degree sequences. This result would be of interest since many protocols for communications networks (such as peer-to-peer networks) use directed communications (see for example [12, 18]).

The structure of the rest of paper is as follows. The necessary Markov chain definitions are given in the next subsection, together with the new result (Lemma 1.3) for bounding the smallest eigenvalue of an ergodic, reversible Markov chain. In Section 2 we define the switch chain \mathcal{M} and prove that it is ergodic on $\Omega_{n,d}$ for $d \geq 1$. A bound on the smallest eigenvalue of the chain is given in Lemma 2.4, and a bound on the second-largest eigenvalue is stated in Proposition 2.5. To conclude Section 2, we show how Theorem 1.1 follows from Proposition 2.5, and give an overview of the main steps of the multicommodity flow argument which is used to prove Proposition 2.5. This argument is presented in

Sections 3– 5. Finally, a worked example is given in Section 6 which illustrates several features of the multicommodity flow construction.

Before we begin our analysis, we mention some recent related work. In many practical situations, almost uniformly random samples are required in order to estimate the average value of some observable of a system. Kim et al. [14] describe an alternative approach to this problem in the case of sampling directed graphs with given in-degrees and out-degrees. Let \mathbf{d}^+ , \mathbf{d}^- be two vectors of nonnegative integers with a common sum. Denote by $\Omega_{n,\mathbf{d}^+,\mathbf{d}^-}$ the set of all digraphs on the vertex set $[n]$ with in-degree sequence \mathbf{d}^+ and out-degree sequence \mathbf{d}^- (and assume that this set is nonempty). Kim et al. describe an algorithm which runs in time $O(n^3)$ and produces a random element of $\Omega_{n,\mathbf{d}^+,\mathbf{d}^-}$, drawn from a specific non-uniform distribution. The samples output by the algorithm are statistically independent, and the algorithm can calculate the weight of each digraph that it produces. They then explain how combining their algorithm with biased sampling allows the average value of any function on $\Omega_{n,\mathbf{d}^+,\mathbf{d}^-}$ to be approximated. However, they do not analyse the running time of the biased sampling procedure, which could be very inefficient when the output distribution is far from uniform. (Indeed, in [14, Section 4.1] they assume that the number of samples in the biased sampling is some positive integer multiple of $|\Omega_{n,\mathbf{d}^+,\mathbf{d}^-}|$, which is usually exponentially large.)

We complete this section with a final remark. Milo et al. [22] wrote of the switch chain for directed graphs that “Theoretical bounds on the mixing time exist only for specific near-regular degree sequences”, citing Kannan, Tetali and Vempala [13]. However, this is not correct, as we now explain. Two Markov chains are considered in [13]. The first is an analogue of the switch chain for undirected graphs. A bound on the mixing time is given in [13] for near-regular bipartite undirected graphs, but no conclusion can be drawn from this for directed graphs. The second chain analysed in [13] is a Markov chain for tournaments with a given score sequence. (A *tournament* is a digraph obtained by giving an orientation to each edge in an (undirected) complete graph. Its *score sequence* is the sequence of out-degrees.) Each transition of the Markov chain reverses the arcs of a directed 3-cycle, so it is quite different from the switch chain. Furthermore, tournaments are very special kinds of digraphs. We know of no rigorous polynomial bound on the mixing time of the switch chain for digraphs, other than Theorem 1.1.

Acknowledgements. I am grateful to Brendan McKay for his suggestion that it seemed unnecessary to make the switch chain lazy, which led to the approach taken here. I am also grateful to the anonymous referee for their helpful comments, which improved both the content and the structure of this paper.

1.1 Markov chain definitions and a new bound on the smallest eigenvalue

Let \mathcal{M} be an ergodic, time-reversible Markov chain on the finite state space Ω with transition matrix P and stationary distribution π . The *total variation distance* between

two probability distributions σ, σ' on Ω is given by

$$d_{\text{TV}}(\sigma, \sigma') = \frac{1}{2} \sum_{x \in \Omega} |\sigma(x) - \sigma'(x)|.$$

The *mixing time* $\tau(\varepsilon)$ is defined by

$$\tau(\varepsilon) = \max_{x \in \Omega} \min \{T \geq 0 \mid d_{\text{TV}}(P_x^t, \pi) \leq \varepsilon \text{ for all } t \geq T\},$$

where P_x^t is the distribution of the state X_t of the Markov chain after t steps from the initial state $X_0 = x$. Let $\pi^* = \min\{\pi(x) \mid x \in \Omega\}$ be the minimum stationary probability.

The transition matrix P has real eigenvalues

$$1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{N-1} \geq -1,$$

where $N = |\Omega|$, and the Markov chain is aperiodic if and only if $\lambda_{N-1} > -1$. Let

$$\lambda_* = \max\{\lambda_1, |\lambda_{N-1}|\} \tag{1}$$

be the second-largest eigenvalue in absolute value. The following result follows from Sinclair [28, Proposition 1], which is based on a result of Diaconis and Stroock [9, Proposition 3].

Lemma 1.2. ([28, Proposition 1]) *The mixing time of the Markov chain \mathcal{M} satisfies*

$$\tau(\varepsilon) \leq (1 - \lambda_*)^{-1} (\log(1/\pi^*) + \log(\varepsilon^{-1})).$$

It has become common practice when applying this bound to first make the Markov chain lazy (that is, replace the transition matrix P by $(I + P)/2$). This ensures that all eigenvalues of the chain are nonnegative, so that $\lambda_* = \lambda_1$ and only $(1 - \lambda_1)^{-1}$ needs to be bounded. However, we prefer not to use introduce unnecessary laziness and seek an alternative approach.

Diaconis and Stroock proved a result [9, Proposition 2] which provides an upper bound on $(1 + \lambda_{N-1})^{-1}$, where λ_{N-1} is the smallest eigenvalue of a Markov chain as in (1). In Lemma 1.3 below, we give a new method for bounding on λ_{N-1} . The new bound is obtained by modifying [9, Proposition 2] in the same way that Sinclair modified [9, Proposition 1] to produce [28, Theorem 5]. The modification results in a bound which is more local in character and seems easier to apply than [9, Proposition 2]. (See also the discussion in [28, Section 2].)

To state the new bound we need some notation. Write \mathcal{G} for the underlying graph of the Markov chain \mathcal{M} , where $\mathcal{G} = (\Omega, \Gamma)$ and each edge $e \in \Gamma$ corresponds to a transition of \mathcal{M} . That is, $e = \{x, y\}$ is an edge of \mathcal{G} if and only if $P(x, y) > 0$. Define $Q(e) = Q(x, y) = \pi(x)P(x, y)$ for the edge $e = \{x, y\}$. (If $P(x, x) > 0$ then the edge $\{x, x\}$ is called a *self-loop* at x .)

For each $x \in \Omega$, fix a particular cycle from x to x in \mathcal{G} with an odd number of edges, and denote it by σ_x . (Such a cycle exists for each x , since the Markov chain is aperiodic.) Note that σ_x may be a 1-cycle, which is a walk along a self-loop edge at x . Write $|\sigma_x|$ to denote the length of the cycle σ_x , which is a positive odd number. Let $\Sigma = \{\sigma_x : x \in \Omega\}$ be the set of these odd cycles, and define the parameter

$$\eta = \eta(\Sigma) = \max_{e \in \Gamma} \frac{1}{Q(e)} \sum_{x \in \Omega, e \in \sigma_x} |\sigma_x| \pi(x).$$

Lemma 1.3. *Suppose that \mathcal{M} is a reversible, ergodic Markov chain with state space Ω . Let $N = |\Omega|$ and let the eigenvalues of \mathcal{M} be given by (1). Then*

$$(1 + \lambda_{N-1})^{-1} \leq \frac{\eta}{2}.$$

Proof. The proof is very similar to the proof of [9, Proposition 2], but using a different application of the Cauchy-Schwarz inequality, as in the proof of [28, Theorem 5]. Assign an arbitrary orientation to each cycle σ_x and let $e = (e^-, e^+)$ under this orientation. Also define $\ell(e)$ to be the distance from x to e^- along the oriented cycle σ_x . Then for any function $\psi : \Omega \rightarrow \mathbb{R}$ we have

$$\psi(x) = \frac{1}{2} \sum_{e \in \sigma_x} (-1)^{\ell(e)} (\psi(e^+) + \psi(e^-))$$

for all $x \in \Omega$. Given $\psi, \varphi : \Omega \rightarrow \mathbb{R}$, let

$$\langle \psi, \varphi \rangle_\pi = \sum_{x \in \Omega} \psi(x) \varphi(x) \pi(x), \quad \mathbf{E}_\pi(\psi) = \sum_{x \in \Omega} \psi(x) \pi(x).$$

Then for any nonzero function $\psi : \Omega \rightarrow \mathbb{R}$ we have

$$\begin{aligned} \mathbf{E}_\pi(\psi^2) &= \sum_{x \in \Omega} \psi(x)^2 \pi(x) = \sum_{x \in \Omega} \pi(x) \left(\frac{1}{2} \sum_{e \in \sigma_x} (-1)^{\ell(e)} (\psi(e^+) + \psi(e^-)) \right)^2 \\ &\leq \frac{1}{4} \sum_x \pi(x) |\sigma_x| \sum_{e \in \sigma_x} (\psi(e^+) + \psi(e^-))^2, \end{aligned}$$

using the Cauchy-Schwarz inequality. Exchanging the order of summation (and now orienting each edge $e \in \Gamma$ arbitrarily) gives

$$\begin{aligned} \mathbf{E}_\pi(\psi^2) &= \frac{1}{4} \sum_{e \in \Gamma} (\psi(e^+) + \psi(e^-))^2 \sum_{x \in \Omega, e \in \sigma_x} |\sigma_x| \pi(x) \\ &\leq \frac{\eta}{4} \sum_{e \in \Gamma} (\psi(e^+) + \psi(e^-))^2 Q(e) \\ &= \frac{\eta}{2} (\mathbf{E}_\pi(\psi^2) + \langle \psi, P\psi \rangle_\pi). \end{aligned}$$

Divide through by $\mathbf{E}_\pi(\psi^2)$ to obtain

$$1 \leq \frac{\eta}{2} \left(1 + \frac{\langle \psi, P\psi \rangle_\pi}{\mathbf{E}_\pi(\psi^2)} \right).$$

Now set ψ equal to any eigenfunction ψ_{N-1} corresponding to λ_{N-1} . After rearranging this completes the proof, since

$$\langle \psi_{N-1}, P\psi_{N-1} \rangle_\pi = \lambda_{N-1} \langle \psi_{N-1}, \psi_{N-1} \rangle_\pi = \lambda_{N-1} \mathbf{E}_\pi(\psi_{N-1}^2).$$

□

This leads to an analogue of [28, Corollary 6]. We also prove a bound for a special case which often arises.

Corollary 1.4. *Under the same conditions as Lemma 1.3 we have*

$$(1 + \lambda_{N-1})^{-1} \leq \frac{\eta'(\Sigma) \ell(\Sigma)}{2}$$

where

$$\eta'(\Sigma) = \max_{e \in \Gamma} \frac{1}{Q(e)} \sum_{x \in \Omega, e \in \sigma_x} \pi(x), \quad \ell(\Sigma) = \max_{x \in \Omega} |\sigma_x|.$$

In particular, if $\ell(\Sigma) = 1$ then

$$(1 + \lambda_{N-1})^{-1} \leq \frac{1}{2} \max_{x \in \Omega} P(x, x)^{-1}$$

(where P denotes the transition matrix of the Markov chain).

Proof. The first statement follows immediately from Theorem 1.3. Now suppose that $\ell(\Sigma) = 1$. Then each $\sigma_x \in \Sigma$ is a self-loop. If $e = (y, y) \in \Gamma$ is a self-loop at y then e is contained in exactly one element of Σ , namely σ_y . In this case

$$\frac{1}{Q(e)} \sum_{x, e \in \sigma_x} \pi(x) = \frac{\pi(y)}{Q(y, y)} = P(y, y)^{-1}.$$

If $e \in \Gamma$ is not a self-loop then e is not contained in any element of Σ , and in this case

$$\frac{1}{Q(e)} \sum_{x, e \in \sigma_x} \pi(x) = 0.$$

Therefore $\eta'(\Sigma) = \max_{x \in \Omega} P(x, x)^{-1}$ and the second statement follows from the first. □

We will use the *multicommodity flow* method of Sinclair [28] to bound the second eigenvalue λ_1 . A *flow* in \mathcal{G} is a function $f : \mathcal{P} \rightarrow [0, \infty)$ which satisfies

$$\sum_{p \in \mathcal{P}_{xy}} f(p) = \pi(x)\pi(y) \quad \text{for all } x, y \in \Omega, x \neq y, \tag{2}$$

where \mathcal{P}_{xy} is the set of all simple directed paths from x to y in \mathcal{G} and $\mathcal{P} = \cup_{x \neq y} \mathcal{P}_{xy}$. Extend f to a function on oriented edges by setting

$$f(e) = \sum_{p \ni e} f(p),$$

so that $f(e)$ is the total flow routed through e . Let $\ell(f)$ be the length of the longest path with $f(p) > 0$, and let

$$\rho(e) = f(e)/Q(e)$$

be the *load* of the edge e . The *maximum load* of the flow is

$$\rho(f) = \max_e \rho(e).$$

Sinclair [28, Corollary 6'] proves the following.

Lemma 1.5. ([28, Corollary 6']) *For any reversible ergodic Markov chain \mathcal{M} and any flow f , the second eigenvalue λ_1 satisfies*

$$(1 - \lambda_1)^{-1} \leq \rho(f)\ell(f).$$

2 The switch chain

Let $d : \{4, 5, \dots\} \rightarrow \mathbb{N}$ be any function such that $1 \leq d(n) \leq n-1$ for all $n \geq 4$, and denote by $\Omega_{n,d(n)}$ the set of all $d(n)$ -regular simple digraphs with vertex set $[n] = \{1, 2, \dots, n\}$. We usually hide the dependence of d on n and just write d rather than $d(n)$; similarly we write $\Omega_{n,d}$ for $\Omega_{n,d(n)}$.

We will study the Markov chain \mathcal{M} described in Figure 1, which we call the *switch chain*. From a given state, an unordered pair of two distinct arcs are chosen uniformly at random. Then the two chosen arcs exchange heads, unless the chosen arcs are incident or exchanging their heads would create a repeated arc. Note that two arcs are non-incident if and only if the set of endvertices of the two arcs contains exactly four vertices. We will write $[ijk\ell]$ as shorthand notation for the switch that replaces the arcs $(i, j), (k, \ell)$ with the arcs $(i, \ell), (k, j)$, as in Figure 1.

The transition matrix P of the Markov chain satisfies $P(X, Y) = P(Y, X) = 1/\binom{dn}{2}$ if X and Y differ by just a switch, with all other non-diagonal entries equal to zero. Therefore P is symmetric, so the stationary distribution of the Markov chain is uniform over $\Omega_{n,d}$.

It is not difficult to see that the switch chain is aperiodic, but for completeness we give a brief proof.

Lemma 2.1. *The switch chain on $\Omega_{n,d}$ is aperiodic for $n \geq 4$ and $1 \leq d \leq n-1$.*

<p>From $G \in \Omega_{n,d}$ do choose an unordered pair of two distinct arcs $\{(i, j), (k, \ell)\}$, u.a.r., if $\{i, j, k, \ell\} = 4$ and $\{(i, \ell), (k, j)\} \cap A(G) = \emptyset$ then delete the arcs $(i, j), (k, \ell)$ and add the arcs $(i, \ell), (k, j)$, else do nothing; end if; end.</p>

Figure 1: The Markov chain on $\Omega_{n,d}$

Proof. Fix $G \in \Omega_{n,d}$ and choose an arc $(\alpha, \beta) \in A(G)$. Since $d \geq 1$, there exists an arc $(\gamma, \alpha) \in A(G)$. These two arcs are distinct but incident, and if they are the arcs chosen in the transition procedure then the switch will be rejected and the chain will remain at G . Hence $P(G, G) \geq 1/\binom{dn}{2} > 0$. So there is a self-loop at every state of $\Omega_{n,d}$, which proves that the chain is aperiodic. \square

Rao, Jana and Bandyopadhyay [24] showed that the switch chain is not always irreducible on the set of all digraphs with a given degree sequence. Characterisations of degree sequences for which the chain is irreducible were given in [2, 15]. We will now prove that when $n \geq 4$ and $d \geq 1$ the set $\Omega_{n,d}$ is connected under switches; that is, that the switch chain is irreducible on $\Omega_{n,d}$. (This was already known when $d = 1$, see Diaconis, Graham and Holmes [8, Remark 2].)

We will use results from LaMar [15]. For a set U of vertices in a digraph G , define the sets $\mathcal{W}^{(i,j)} = \mathcal{W}^{(i,j)}(U, G)$ for $(i, j) \in \mathbb{Z}_2^2$, as follows:

$$\begin{aligned} \mathcal{W}^{(0,0)} &= \{x \in V(G) - U : (x, u) \notin A(G), (u, x) \notin A(G) \text{ for all } u \in U\}, \\ \mathcal{W}^{(0,1)} &= \{x \in V(G) - U : (x, u) \notin A(G), (u, x) \in A(G) \text{ for all } u \in U\}, \\ \mathcal{W}^{(1,0)} &= \{x \in V(G) - U : (x, u) \in A(G), (u, x) \notin A(G) \text{ for all } u \in U\}, \\ \mathcal{W}^{(1,1)} &= \{x \in V(G) - U : (x, u) \in A(G), (u, x) \in A(G) \text{ for all } u \in U\}. \end{aligned}$$

(In [15] these sets are called $\mathcal{C}^0, \mathcal{C}^+, \mathcal{C}^-, \mathcal{C}^\pm$, respectively.)

Lemma 2.2. *The space $\Omega_{n,d}$ is connected under switches when $n \geq 4$ and $1 \leq d \leq n - 1$.*

Proof. For a contradiction, assume that $\Omega_{n,d}$ is not connected under switches. Then by LaMar [15, Theorems 3.3 and 3.4], for any digraph $G \in \Omega_{n,d}$ there is a set of vertices $\{v_0, v_1, v_2\}$ such that the induced digraph $G[\{v_0, v_1, v_2\}]$ is a directed 3-cycle and, writing $\mathcal{W}^{(i,j)} = \mathcal{W}^{(i,j)}(\{v_0, v_1, v_2\}, G)$ for all $(i, j) \in \mathbb{Z}_2^2$,

- (i) all vertices in $V(G)$ other than $\{v_0, v_1, v_2\}$ belong to $\bigcup_{(i,j) \in \mathbb{Z}_2^2} \mathcal{W}^{(i,j)}$,
- (ii) no arcs from $\mathcal{W}^{(0,0)} \cup \mathcal{W}^{(0,1)}$ to $\mathcal{W}^{(0,0)} \cup \mathcal{W}^{(1,0)}$ are present,

(iii) all (non-loop) arcs from $\mathcal{W}^{(1,0)} \cup \mathcal{W}^{(1,1)}$ to $\mathcal{W}^{(0,1)} \cup \mathcal{W}^{(1,1)}$ are present.

Let $n^{(i,j)} = |\mathcal{W}^{(i,j)}|$ for $(i, j) \in \mathbb{Z}_2^2$. Considering the in-degree and out-degree of v_0 gives, using (i),

$$d = n^{(1,0)} + n^{(1,1)} + 1 = n^{(0,1)} + n^{(1,1)} + 1$$

(and in particular, $n^{(0,1)} = n^{(1,0)}$). However, by (ii), the in-degree of any element of $\mathcal{W}^{(1,0)}$ is at most $n^{(1,0)} + n^{(1,1)} - 1 = d - 2$, the out-degree of any element of $\mathcal{W}^{(0,1)}$ is at most $n^{(0,1)} + n^{(1,1)} - 1 = d - 2$ and the out-degree of any element of $\mathcal{W}^{(0,0)}$ is at most $n^{(0,1)} + n^{(1,1)} = d - 1$. This contradicts the assumption that $G \in \Omega_{n,d}$ unless

$$\mathcal{W}^{(0,0)} \cup \mathcal{W}^{(0,1)} \cup \mathcal{W}^{(1,0)} = \emptyset.$$

But then $\mathcal{W}^{(1,1)} = V - \{v_0, v_1, v_2\}$, and this set is nonempty as $n \geq 4$. By (iii), the in-degree of any element of $\mathcal{W}^{(1,1)}$ is $n - 1$. Since the in-degree of v_0 is $n - 2$, we obtain a contradiction. \square

Suppose that $G \in \Omega_{n,d}$ contains a directed 3-cycle on the vertices v_0, v_1, v_2 . Consider the sets $\mathcal{W}^{(i,j)} = \mathcal{W}^{(i,j)}(\{v_0, v_1, v_2\}, G)$ where $(i, j) \in \mathbb{Z}_2^2$. If $x \in V(G) - \{v_0, v_1, v_2\}$ does not belong to $\bigcup_{(i,j) \in \mathbb{Z}_2^2} \mathcal{W}^{(i,j)}$ then we say that x is a *useful neighbour* for the given 3-cycle. (Note that x must be an in-neighbour or an out-neighbour of at least one vertex on the 3-cycle, since $x \notin \mathcal{W}^{(0,0)}$.) Similarly, (x, y) is called a *useful arc* for the given 3-cycle if $x \neq y$ and one of the following conditions holds:

(U1) $(x, y) \in A(G)$, with $x \in \mathcal{W}^{(0,0)} \cup \mathcal{W}^{(0,1)}$ and $y \in \mathcal{W}^{(0,0)} \cup \mathcal{W}^{(1,0)}$;

(U2) $(x, y) \notin A(G)$, with $x \in \mathcal{W}^{(1,0)} \cup \mathcal{W}^{(1,1)}$ and $y \in \mathcal{W}^{(0,1)} \cup \mathcal{W}^{(1,1)}$.

The following result will be used later.

Lemma 2.3. *Suppose that $G \in \Omega_{n,d}$ contains a set of three vertices $\{v_0, v_1, v_2\}$ such that the induced digraph $G[\{v_0, v_1, v_2\}]$ is a directed 3-cycle. Then there exists a useful neighbour or a useful arc for this 3-cycle.*

Proof. Suppose that there is no useful neighbour for the 3-cycle. Then condition (i) from the proof of Lemma 2.2 holds. For a contradiction, assume that there is no useful arc (x, y) . Then all (non-loop) arcs from $\mathcal{W}^{(0,0)} \cup \mathcal{W}^{(0,1)}$ to $\mathcal{W}^{(0,0)} \cup \mathcal{W}^{(1,0)}$ are absent in G , and all (non-loop) arcs from $\mathcal{W}^{(1,0)} \cup \mathcal{W}^{(1,1)}$ to $\mathcal{W}^{(0,1)} \cup \mathcal{W}^{(1,1)}$ are present in G . That is, conditions (ii) and (iii) from the proof of Lemma 2.2 also hold. Arguing as in the proof of Lemma 2.2 leads to a contradiction. \square

Now we prove a bound on the smallest eigenvalue of the switch chain.

Lemma 2.4. *Suppose that $n \geq 4$ and $1 \leq d \leq n - 1$, and let $N = |\Omega_{n,d}|$. The smallest eigenvalue λ_{N-1} of the switch chain satisfies*

$$(1 + \lambda_{N-1})^{-1} \leq \frac{1}{4} d^2 n^2.$$

Proof. By Lemma 2.1, there is a self-loop σ_x in Γ at every $x \in \Omega_{n,d}$. Let $\Sigma = \{\sigma_x\}$ be the set of these 1-cycles. Since $\ell(\Sigma) = 1$, the result follows from the second statement of Corollary 1.4 since

$$\max_{x \in \Omega_{n,d}} P(x, x)^{-1} \leq \binom{dn}{2}.$$

□

We also need a bound on the second-largest eigenvalue of the switch chain.

Proposition 2.5. *Suppose that $n \geq 4$ and $1 \leq d \leq n - 1$, and let λ_1 be the second-largest eigenvalue of the switch chain on $\Omega_{n,d}$. Then*

$$(1 - \lambda_1)^{-1} \leq 50d^{25}n^9.$$

The proof of Proposition 2.5 is lengthy and quite technical. We give an outline of the proof below, and full details in Sections 3–5. But first, we show how Theorem 1.1 can be proved from Proposition 2.5.

Proof of Theorem 1.1. If the smallest eigenvalue λ_{N-1} is nonnegative then $\lambda_* = \lambda_1$, and by Proposition 2.5 we have

$$(1 - \lambda_*)^{-1} \leq 50d^{25}n^9. \tag{3}$$

Suppose now that λ_{N-1} is negative. Then $1 - |\lambda_{N-1}| = 1 + \lambda_{N-1}$ and it follows from Lemma 2.4 and Proposition 2.5 that (3) also holds in this case.

Finally, we note that

$$\log |\Omega_{n,d}| \leq dn \log(dn). \tag{4}$$

(This is well-known but for completeness we sketch a proof. Take a bipartite graph on $n + n$ vertices and assign d “half-edges” to each vertex on the side. Arbitrarily match each half-edge on the left to a half-edge on the right. There are at most $(dn)^{dn}$ ways to perform this matching. Finally, orient each edge from left to right and identify the j 'th vertex on each side, giving a digraph on n vertices which may have loops or multiple arcs. As each element of $\Omega_{n,d}$ can be formed from at least one matching in this way, we obtain an upper bound.) Hence, since π is uniform,

$$\log 1/\pi^* = \log |\Omega_{n,d}| \leq dn \log(dn).$$

Substituting (3) and (4) into Lemma 1.2 gives the stated bound on the mixing time, completing the proof. □

Hence it remains to establish Proposition 2.5. We use a multicommodity flow argument to prove this result. Before embarking on the proof, we outline the major steps in the argument. (We note that our proof follows the same general outline as most canonical path or multicommodity flow arguments, where encodings are often used. In particular, our proof builds upon the argument from [6].)

- Given distinct digraphs $G, G' \in \Omega_{n,d}$, we define a finite set $\Psi(G, G')$ of objects, called pairings. For each $\psi \in \Psi(G, G')$ we will define a canonical path $\gamma_\psi(G, G')$ from G to G' , indexed by ψ . Then the flow f is defined on

$$\bigcup_{(G, G')} \{\gamma_\psi(G, G') \mid \psi \in \Psi(G, G')\}$$

by

$$f(\gamma_\psi(G, G')) = \frac{\pi(G)\pi(G')}{|\Psi(G, G')|} = (|\Omega_{n,d}|^2 |\Psi(G, G')|)^{-1},$$

and is set to zero for all other paths. Note that f satisfies (2).

- To define $\gamma_\psi(G, G')$ we work with the symmetric difference $H = G \Delta G'$ of G and G' (with arcs of $G - G'$ coloured blue and arcs of $G' - G$ coloured red). In Sections 3.1 and 3.2 we show how to decompose H into a sequence of arc-disjoint subdigraphs called 1-circuits and 2-circuits, in a canonical way. The canonical path $\gamma_\psi(G, G')$ is formed by processing each of these 1-circuits and 2-circuits in the given order.
- We can process 1-circuits, and certain 2-circuits, in a way which is very similar to the method used in [6] for undirected graphs. The 2-circuits which can be handled in this way are called *normal*. Section 3.3 explains how to process a 1-circuit and Section 4.1 describes how to process a normal 2-circuit.
- The main difficulties in the proof arise from the need to handle 2-circuits which are not normal. We further categorise these as *eccentric 2-circuits* or *triangles*. Sections 4.2 and 4.3 describe how to process these 2-circuits.

By this stage, the multicommodity flow is completely defined. Next we must analyse the flow in order to bound the maximum load of the flow, and hence the second-largest eigenvalue (using Lemma 1.5).

- Let (Z, Z') be a transition along one of the canonical paths $\gamma_\psi(G, G')$, and suppose that this transition is performed while processing the 1-circuit or 2-circuit S . A set of *interesting arcs* for Z with respect to (G, G', ψ) is defined. These are arcs which have been disturbed during the processing of S and not yet returned to their original state, and they will play a key role in our analysis. Lemma 5.1 describes the structure of the digraph formed by the interesting arcs (see also Figure 9).
- Next we identify G, G' and Z with their adjacency matrices and define a matrix L by $L + Z = G + G'$. Then L is an $n \times n$ matrix with entries in $\{-1, 0, 1, 2\}$. We say that L is an *encoding* for Z with respect to (G, G') . Lemma 5.2 shows that given $(Z, Z'), L$ and ψ there are at most four possibilities for (G, G') such that (Z, Z') is a transition on $\gamma_\psi(G, G')$ and L is an encoding for Z with respect to (G, G') . Further information about the structure of L is given in Lemma 5.3.

- Now the notion of encoding is broadened to encompass any $n \times n$ matrix with entries in $\{-1, 0, 1, 2\}$ such that all row sums and column sums equal d . Given $Z \in \Omega_{n,d}$, we say that the encoding L is Z -valid if every entry of $L + Z$ belongs to $\{0, 1, 2\}$ and L, Z, H satisfy the conclusions of Lemma 5.3. (Here H is the digraph defined by all entries of $L + Z$ which equal 1.) Lemma 5.4 proves a useful fact about Z -valid encodings.
- Next we explain how to apply switches to encodings, and prove in Lemma 5.5 (using Lemma 5.4) that any Z -valid encoding can be transformed into an element of $\Omega_{n,d}$ using at most three switches. Counting the number of ways these switches can be performed in reverse leads to an upper bound of the form $\text{poly}(n, d) |\Omega_{n,d}|$ on the number of Z -valid encodings, as proved in Lemma 5.6.
- Combining all this allows us to prove an upper bound on the total flow routed through an arbitrary transition of the Markov chain. This bound, of the form $\text{poly}(n, d) |\Omega_{n,d}|^{-1}$, is proved in Lemma 5.7. With this in hand it is easy to establish a polynomial bound on the maximum load $\rho(f)$ of the flow, and hence to prove Proposition 2.5.

3 Defining the flow

We now define the multicommodity flow which will be used to bound the second largest eigenvalue, and hence the mixing time, of the switch chain for regular directed graphs.

For $G, G' \in \Omega_{n,d}$, let $H = G \Delta G'$ be the symmetric difference of G and G' , together with an arc-colouring which colours all arcs of $G - G'$ blue and all arcs of $G' - G$ red. This arc colouring means that we can think of H as the symmetric difference of the *ordered pair* (G, G') .

For $v \in V$ let θ_v be the blue in-degree of v , which equals the red in-degree of v , and let ϕ_v be the blue out-degree of v , which equals the red out-degree of v . Choose a *pairing* of the red and blue arcs around each vertex as follows: each blue arc with head v is paired with a red arc with head v , and each blue arc with tail v is paired with a red arc with tail v , defining two bijections (one from the set of blue arcs with head v to the set of red arcs with head v , and one from the set of blue arcs with tail v to the set of red arcs with tail v). Denote the set of all such pairings by $\Psi(G, G')$. Then

$$|\Psi(G, G')| = \prod_{v \in V} \theta_v! \phi_v! \tag{5}$$

is the total number of pairings.

Write \mathcal{G} for the underlying graph of the Markov chain \mathcal{M} , where $\mathcal{G} = (\Omega_{n,d}, \Gamma)$ and each edge $e \in \Gamma$ corresponds to a transition of \mathcal{M} . For each pairing in $\Psi(G, G')$ we construct a canonical path from G to G' in \mathcal{G} . Each of these paths will carry $1/|\Psi(G, G')|$ of the total flow from G to G' .

We now introduce some necessary terminology. A *forward circuit* in H is a string $C = w_0w_1 \cdots w_{2k-1}$ over the alphabet V such that the arcs

$$(w_0, w_1), (w_2, w_1), (w_2, w_3), (w_4, w_3), \dots, (w_{2k-2}, w_{2k-1}), (w_0, w_{2k-1}) \quad (6)$$

are all distinct, all belong to A_H and alternate in colour: that is, the arcs in

$$\{(w_{2i}, w_{2i+1}) : i = 0, 1, \dots, k-1\}$$

all have one colour and the arcs in

$$\{(w_{2i+2}, w_{2i+1}) : i = 0, 1, \dots, k-2\} \cup \{(w_0, w_{2k-1})\}$$

all have the other colour.

The *converse* of H is the digraph obtained from H by reversing the direction of every arc (but keeping the colours the same). A *reverse circuit* in H is a string $C = w_0w_1 \cdots w_{2k-1}$ over the alphabet V which forms a forward circuit in the converse of H . That is, the arcs

$$(w_1, w_0), (w_1, w_2), (w_3, w_2), (w_3, w_4), \dots, (w_{2k-1}, w_{2k-2}), (w_{2k-1}, w_0) \quad (7)$$

are all distinct, all belong to A_H and alternate in colour, so that the arcs in

$$\{(w_{2i+1}, w_{2i}) : i = 0, 1, \dots, k-1\}$$

all have one colour and the arcs in

$$\{(w_{2i+1}, w_{2i+2}) : i = 0, 1, \dots, k-2\} \cup \{(w_{2k-1}, w_0)\}$$

all have the other colour. By *circuit* we mean either a forward circuit or a reverse circuit. For a forward or reverse circuit C , denote by $A(C)$ the set of arcs in (6) or (7), respectively. It is important to note that the arcs of a circuit alternate both in colour and orientation at each step. While a circuit may contain both the arcs (x, y) and (y, x) , any three consecutive vertices on the circuit are distinct.

We now define two operations on digraphs. Let ζ denote the operation which takes a digraph to its converse (that is, it reverses every arc in the digraph), and let χ be the operation which takes a digraph to its complement. Writing $[n]^{(2)}$ for the set of all ordered pairs of distinct elements of $[n]$, the complement χG of a digraph G has arc set $[n]^{(2)} - A(G)$. Note that the operations ζ and χ commute and are both involutions.

We can also apply ζ and χ to the (arc-coloured) symmetric difference $H = G \Delta G'$. Here ζH is the result of reversing every arc in H , without changing the colour of any arc. Similarly, χH is the result of exchanging the colour of every arc in H (so that blue becomes red and vice-versa), without changing the orientation of any arc. To see this, note that the set of blue arcs in $H = G \Delta G'$ equals

$$A(G) - A(G') = A(\chi G') - A(\chi G),$$

but this equals the set of red arcs in $(\chi G)\Delta(\chi G')$ (and similarly, the set of red arcs in $G\Delta G'$ equals the set of blue arcs in $(\chi G)\Delta(\chi G')$).

Finally, we generalise these definitions so that they also apply to (arc-coloured) subdigraphs U of H . That is, ζU is the result of reversing every arc in U , without changing the colour of any arc, while χU is the result of exchanging the colour of every arc in U without changing the orientation of any arc.

3.1 Decomposition into circuits

Fix a pairing $\psi \in \Psi(G, G')$. We decompose H into a sequence of circuits depending on ψ , as follows. Let (w_0, w_1) be the lexicographically least arc in H . Choose the arc (w_2, w_1) which is paired with (w_0, w_1) at w_1 . (Note that if (w_0, w_1) is blue then (w_2, w_1) is red, and vice-versa. Furthermore, $w_2 \neq w_0$ since H is a symmetric difference.) Next choose the arc (w_2, w_3) which is paired with (w_2, w_1) at w_2 . (This arc will have the same colour as (w_0, w_1) .) Continue in this fashion. Specifically, for $i \geq 1$, if $w_{2i} \neq w_0$ then let (w_{2i}, w_{2i+1}) be the arc which is paired with (w_{2i}, w_{2i-1}) at w_{2i} and let (w_{2i+2}, w_{2i+1}) be the arc which is paired with (w_{2i}, w_{2i+1}) at w_{2i+1} . The vertices w_i are not necessarily distinct, but the arcs are distinct. The process terminates when (w_0, w_{2k-1}) is paired with (w_0, w_1) at w_0 , giving a forward circuit $C_1 = w_0 w_1 \cdots w_{2k-1}$.

If $A(C_1) = A_H$ then $\mathcal{C} = \{C_1\}$ and we are done. Otherwise, take the lexicographically least arc not in C_1 and generate a new circuit C_2 by the above procedure. Continue generating circuits until

$$A_H = A(C_1) \cup A(C_2) \cup \cdots \cup A(C_s).$$

Then $\mathcal{C} = \{C_1, C_2, \dots, C_s\}$ and the circuits C_1, C_2, \dots, C_s are arc-disjoint. Note that, once the pairing has been chosen, \mathcal{C} is formed without regard to the colouring of H . This property will be needed later.

Using \mathcal{C} , we form a path

$$G = Z_0, Z_1, \dots, Z_M = G'$$

from G to G' in the underlying graph of the Markov chain (that is, to get from Z_a to Z_{a+1} we perform a switch). The path is defined by processing each circuit C_i in turn. Processing a circuit changes its arcs from agreeing with G to agreeing with G' , with no other arcs being permanently altered (though some may be temporarily changed while processing the circuit C_i). The canonical path is defined inductively. If

$$G = Z_0, Z_1, \dots, Z_r$$

is the canonical path obtained by processing the first circuit C_1 , and

$$Z_r, Z_{r+1}, \dots, Z_N = G'$$

is the canonical path from Z_r to G' obtained by processing the circuits (C_2, \dots, C_s) , in order, then the canonical path from G to G' corresponding to ψ is given by the concatenation of these two paths.

Thus it suffices to describe the canonical path corresponding to a particular circuit $C = w_0w_1 \cdots w_{2k-1}$. First we may need to decompose the circuit C further. A *1-circuit* $S = v_0v_1v_2 \cdots v_t$ is a string on the letters $\{w_0, w_1, \dots, w_{2k-1}\}$ such that $v_0 = w_0$ and w_0 appears only once in S . Usually a 1-circuit will be a contiguous substring of C (allowing reversal of direction and/or cyclic wrapping around C), but it may also contain one arc which is not an arc of C .

We now show how to decompose a circuit into a sequence of 1-circuits (and possibly some single switches) which will then be processed in order (as described in Section 3.3) to form the canonical path corresponding to C .

3.2 Decomposition of a circuit

Given a circuit $C = w_0w_1 \cdots w_{2k-1}$, let $v = w_0$ and let $C^{(0)} = C$ (which is the currently unprocessed segment of C). Suppose that the current digraph on the canonical path from G to G' is Z_J . If $w_i \neq v$ for $i = 1, \dots, 2k - 1$ then $C^{(0)}$ is a 1-circuit which we process (using the procedure described in Section 3.3), extending the canonical path as

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+t}. \quad (8)$$

This completes the processing of C . Otherwise, v appears θ times on $C^{(0)}$, where $2 \leq \theta \leq \theta_v + \phi_v$. Relabel the vertices on $C^{(0)}$ as

$$C^{(0)} = vx_1 \cdots y_1vx_2 \cdots y_2v \cdots vx_\theta \cdots y_\theta.$$

By construction, $C^{(0)}$ is a forward circuit, so $(v, x_1) \in A(C^{(0)})$.

Firstly, suppose that $S = vx_1 \cdots y_1$ is a 1-circuit. That is, arcs (v, x_1) and (v, y_1) are present on S , with opposite colours. Process the 1-circuit S , extending the canonical path as in (8), leaving the forward circuit

$$C^{(1)} = vx_2 \cdots y_2v \cdots vx_\theta \cdots y_\theta$$

as the unprocessed section of C . Then process $C^{(1)}$ inductively.

Next, suppose that we are not in the above situation (so that the arcs (v, x_1) and (y_1, v) are present on S , with the same colour), but that $S = vx_\theta \cdots y_\theta$ is a 1-circuit. That is, the arcs (v, x_θ) and (v, y_θ) are present on S , with opposite colours. Process the 1-circuit S , extending the canonical path as in (8), and leaving the forward circuit

$$C^{(1)} = vx_1 \cdots y_1v \cdots vx_{\theta-1} \cdots y_{\theta-1}$$

to be processed inductively.

Finally, suppose that neither of the two situations above apply. Then the arcs (v, x_1) and (y_1, v) have one colour while (x_θ, v) and (v, y_θ) have the other colour. We will process $S' = vx_1 \cdots y_1vx_\theta \cdots y_\theta$ (which we call a *2-circuit*), extending the canonical path as in (8) and leaving

$$C^{(1)} = vx_2 \cdots y_2v \cdots vx_{\theta-1} \cdots y_{\theta-1}$$

to be processed inductively. Here $C^{(1)}$ is a reverse circuit and we process it using the same procedure as described above, but with all arcs reversed.

All 1-circuits and 2-circuits created by the above procedure are called *raw*. The order in which we detect and process raw 1-circuits and raw 2-circuits implies that both the processed and unprocessed sections of C are contiguous whenever the processing of a raw 1-circuit or raw 2-circuit is complete. (That is, these sections form contiguous substrings of C , where a substring is allowed to wrap around in a cyclic fashion.)

Suppose that S is a raw 1-circuit or raw 2-circuit with substring abc . Fix $i \in \{0, 1\}$ such that the corresponding arcs are $\zeta^i(a, b)$ and $\zeta^i(c, b)$. These arcs are called *successive arcs* along S . Every raw 1-circuit or raw 2-circuit S has the following property: successive arcs along S are paired under ψ at their common endvertex b , except possibly when $b = v$ and the arcs are the first and last arcs of S . We call this the *well-paired* property, which will be used in Lemma 4.1 below.

Raw 1-circuits are processed using the method described in Section 3.3. In most cases, raw 2-circuits must be further decomposed (into a sequence of 1-circuits and/or switches) before they can be processed, as described in Section 5. It is here that extra difficulties arise here when working with directed graphs.

Recall the notation for switches introduced after Figure 1. Let $Q = \alpha\beta\gamma\delta$ be a circuit in H which is also a 4-cycle. Set

$$i = \begin{cases} 0 & \text{if } Q \text{ is a forward circuit,} \\ 1 & \text{otherwise.} \end{cases}$$

We now define notation for the switch which processes this 4-cycle, starting from the current digraph Z_J and producing the next digraph Z_{J+1} on the canonical path. Let $h = 0$ if $\zeta^i(\alpha, \beta) \in Z_J$ and $h = 1$ otherwise. Then define

$$\zeta^i \chi^h[\alpha\beta\gamma\delta] = \begin{cases} [\alpha\beta\gamma\delta] & \text{if } i = 0, h = 0, \\ [\alpha\delta\gamma\beta] & \text{if } i = 0, h = 1, \\ [\beta\alpha\delta\gamma] & \text{if } i = 1, h = 0, \\ [\beta\gamma\delta\alpha] & \text{if } i = 1, h = 1. \end{cases}$$

If $h = 0$ then the switch $\zeta^i \chi^h[\alpha\beta\gamma\delta]$ deletes the arcs $\zeta^i(\alpha, \beta)$, $\zeta^i(\gamma, \delta)$ and replaces them with $\zeta^i(\alpha, \delta)$, $\zeta^i(\gamma, \beta)$, while the opposite occurs if $h = 1$.

Finally, we define the *status* of an arc (x, y) in a digraph Z to equal 0 if $(x, y) \notin A(Z)$ and to equal 1 if $(x, y) \in A(Z)$. We say that two arcs have *matching status* if their status is equal in Z , and say that they have *opposite status* otherwise.

3.3 Processing a 1-circuit

Let S be a 1-circuit. (If S is not raw then S has resulted from the decomposition of a raw 2-circuit: see Sections 4.1 and 4.2.) The method for processing a 1-circuit is very similar to that used in [6], and some of the discussion and figures given there may be helpful. (See also the worked example in Section 6.)

Label the 1-circuit as $S = x_0x_1 \dots x_{2k-1}$ where $k \geq 2$, such that x_0 is the minimum vertex on S and $x_1 = \min\{x_1, x_{2k-1}\}$. Set

$$i = \begin{cases} 0 & \text{if } S \text{ is a forward circuit,} \\ 1 & \text{if } S \text{ is a reverse circuit.} \end{cases}$$

Also set

$$h = \begin{cases} 0 & \text{if } \zeta^i(x_0, x_1) \in A(Z_J), \\ 1 & \text{otherwise.} \end{cases}$$

Then $\zeta^i(x_{2t}, x_{2t+1}) \in A(\chi^h Z_J)$ and $\zeta^i(x_{2t+2}, x_{2t+1}) \notin A(\chi^h Z_J)$ for $t = 0, 1, \dots, k-1$ (identifying x_{2t} with x_0). Note that any three consecutive vertices on S are distinct.

Define the set

$$\mathcal{B} = \{t = 1, 2, \dots, k-1 : \zeta^i(x_0, x_{2t+1}) \notin A(\chi^h Z_J) \\ \text{and } x_{2\ell+1} \neq x_{2t+1} \text{ for } \ell \text{ with } \ell = t+1, \dots, k-1.\}$$

(This definition ensures that exactly one value t is stored for each distinct vertex x_{2t+1} with $\zeta^i(x_0, x_{2t+1}) \notin A(\chi^h Z_J)$, ensuring that vertices which are repeated along S are treated correctly.) Note that $k-1 \in \mathcal{B}$ always. The arcs $\zeta^i(x_0, x_{2t+1})$ are called *odd chords*. The number of phases in the processing of S will be $p = |\mathcal{B}|$. For the first phase, choose the minimum $t \in \mathcal{B}$. There will be t steps in the first phase, which proceeds as follows:

for $j := t, t-1, \dots, 1$ do

form $Z_{J+t-j+1}$ from Z_{J+t-j} by performing the switch $\zeta^i \chi^h[x_0 x_{2j-1} x_{2j} x_{2j+1}]$;

If $t = k-1$ then there is only one phase and the processing of S is complete. Otherwise, $\zeta^i(x_0, x_{2t+1}) \in A(\chi^h Z_{J+t})$ but all odd chords $\zeta^i(x_0, x_{2\ell+1})$ with $x_{2\ell+1} \neq x_{2t+1}$ have been reinstated to match their status in Z_J (that is, they belong to Z_{J+t} if and only if they belong to Z_J).

For subsequent phases, if t was the starting point of the previous phase then choose $q > t$ minimum such that $q \in \mathcal{B}$. The odd chord $\zeta^i(x_0, x_{2t+1})$ has been switched in the previous phase but will be restored to its original state by the end of this phase. There will be $q-t$ steps in this phase, performing the sequence of switches

$$\zeta^i \chi^h[x_0 x_{2q-1} x_{2q} x_{2q+1}], \zeta^i \chi^h[x_0 x_{2q-3} x_{2q-2} x_{2q-1}], \dots, \zeta^i \chi^h[x_0 x_{2t+1} x_{2t+2} x_{2t+3}].$$

Note that every switch involves x_0 , the start-vertex of S .

At any point during the processing of the 1-circuit, at most three odd chords have been switched (that is, temporarily disturbed). This is illustrated in the worked example in Section 6.

4 Decomposition of a raw 2-circuit

We now show how to process a raw 2-circuit S , given the method for processing a 1-circuit described in Section 3.3. Suppose that we have reached the digraph Z_J on the canonical path from G to G' in \mathcal{G} . Note that Z_J agrees with G on the 2-circuit S before the processing of S begins. We relabel the vertices on S as

$$S = vx_{0,0} \cdots x_{1,0}vx_{1,1} \cdots x_{0,1}, \quad (9)$$

where $(v, x_{0,0})$ is the lexicographically least arc in $A(S)$. Treat the indices (i, j) on the vertex labels as elements of \mathbb{Z}_2^2 , with addition performed modulo 2. In the undirected case [6], the vertices $x_{0,0}, x_{0,1}, x_{1,0}, x_{1,1}$ are all distinct. However, this is no longer the case in the directed setting, which complicates the definition of the canonical paths. Also note that there may be as few as two vertices between two successive occurrences of v on the 2-circuit S . This is explained in more detail below Figure 5, once we have introduced some useful notation.

Recall that χ is the complementation operation for digraphs. Set

$$h = \begin{cases} 0 & \text{if the arc } (v, x_{0,0}) \text{ is present in } Z_J, \\ 1 & \text{if the arc } (v, x_{0,0}) \text{ is absent in } Z_J. \end{cases}$$

Then

$$(v, x_{0,0}) \in A(\chi^h Z_J), \quad (x_{1,0}, v) \in A(\chi^h Z_J), \quad (x_{1,1}, v) \notin A(\chi^h Z_J), \quad (v, x_{0,1}) \notin A(\chi^h Z_J).$$

Figure 2 depicts $\chi^h S$, where the curved lines (from $x_{0,0}$ to $x_{1,0}$ and from $x_{0,1}$ to $x_{1,1}$) represent any odd number of alternating arcs. Solid arcs represent arcs which are present in $\chi^h Z_J$ and dashed arcs represent arcs which are absent in $\chi^h Z_J$. That is, if $h = 0$ then solid arcs belong to Z_J and dashed arcs belong to G' , while if $h = 1$ then solid arcs belong to G' and dashed arcs belong to Z_J .

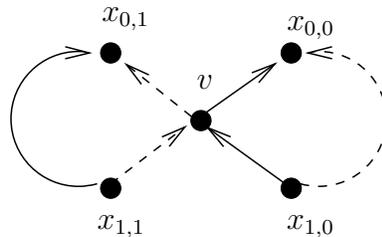


Figure 2: The 2-circuit $\chi^h S$

For $(i, j) \in \mathbb{Z}_2^2$, let $y_{i,j}$ be the unique vertex such that $vx_{i,j}y_{i,j}$ or $y_{i,j}x_{i,j}v$ is a contiguous substring of S (allowing cyclic wrapping in the case of $y_{0,1}$). If $y_{0,j} = x_{1,j}$ for some j

then $y_{1,j} = x_{0,j}$ and there is only one arc between $x_{0,j}$ and $x_{1,j}$. This means that the corresponding curved line in Figure 2 can be replaced by a single arc. There are four possibilities for $\chi^h S$ in which $y_{0,j} = x_{1,j}$ for $j = 1, 2$. These are shown in Figure 3. The leftmost 2-circuit involves 5 distinct vertices and the middle two 2-circuits each involve 4 distinct vertices, with one coincidence of the form $x_{0,j} = x_{1,j+1}$, where $j \in \mathbb{Z}_2$. The rightmost 2-circuit involves 3 distinct vertices: we will call it a *triangle*.

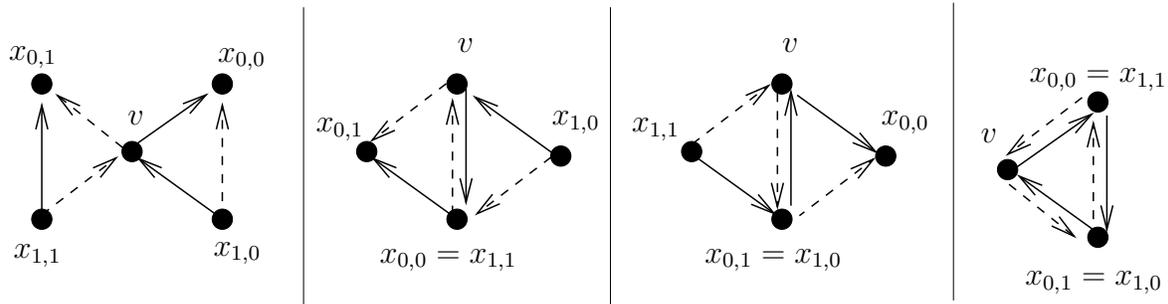


Figure 3: The four 2-circuits $\chi^h S$ with at most 5 distinct vertices. The rightmost 2-circuit is a triangle.

In the undirected analysis [6], a critical observation was that vertex $y_{0,0}$ must be distinct from $x_{0,1}$ (without loss of generality). This fact underpinned the definition of the canonical paths in [6]. For directed graphs this property does not necessarily hold, as can be seen from the last two 2-circuits in Figure 3.

We will say that S is *normal* if $y_{i,j} \neq x_{i,j+1}$ for some $(i, j) \in \mathbb{Z}_2^2$. In Section 4.1 we describe how to process a normal 2-circuit. The procedure is analogous to that used in [6], which is the motivation for the definition of normal 2-circuits. Note that the triangle (shown at the rightmost of Figure 3) is not normal but the remaining 2-circuits in Figure 3 are normal.

For $(i, j) \in \mathbb{Z}_2^2$, let $z_{i,j}$ be the unique vertex such that $vx_{i,j}y_{i,j}z_{i,j}$ or $z_{i,j}y_{i,j}x_{i,j}v$ is a contiguous substring of S (allowing cyclic wrapping in the case of $z_{0,1}$). We will need the following lemma.

Lemma 4.1. *Suppose that S is a raw 2-circuit which is not normal and such that $v = z_{i,j}$ for some $(i, j) \in \mathbb{Z}_2^2$. Then S is a triangle.*

Proof. Without loss of generality (by reversing arcs and/or taking the complement if necessary) we may suppose that $v = z_{0,0}$. (This means we cannot assume that $(v, x_{0,0})$ is the lexicographically least arc in S , but we do not need to use that property in this proof.) Colour arcs around the 2-circuit orange, purple in an alternating fashion, starting with the orange arc $(v, x_{0,0})$. By assumption, S has initial substring $vx_{0,0}y_{0,0}v$. By the well-paired property of 2-circuits, we know that the orange arc $(v, x_{0,0})$ is paired with the purple arc $(y_{0,0}, x_{0,0})$ at $x_{0,0}$ under ψ , and the purple arc $(y_{0,0}, x_{0,0})$ is paired with the

orange arc $(y_{0,0}, v)$ at $y_{0,0}$ under ψ . Now v is incident with exactly four arcs of S , one of each colour and orientation (see Figure 2.) Hence the presence of the orange arc $(y_{0,0}, v)$ on S shows that $y_{0,0} = x_{1,0}$. But then we obtain

$$y_{0,0} = x_{1,0} = y_{1,1} = x_{0,1},$$

as S is not normal.

Now $y_{0,0} = x_{1,0}$ and the purple arc $(y_{0,0}, x_{0,0})$ is paired with the orange arc $(v, y_{0,0})$. This implies that $x_{0,0} = y_{1,0}$, and since S is not normal it follows that

$$x_{0,0} = y_{1,0} = x_{1,1} = y_{0,1}.$$

This gives all pairing information around the 2-circuit except for pairings at v . (For example, since $y_{0,0} = x_{0,1}$ and $x_{0,0} = y_{0,1}$, we know that the purple arc $(v, y_{0,0})$ is present on S and is paired with the orange arc $(x_{0,0}, y_{0,0})$ at $y_{0,0}$. This arc is paired at $x_{0,0}$ with the purple arc $(x_{0,0}, v)$, since $x_{0,0} = x_{1,1}$ and $y_{0,0} = y_{1,1}$.) But as v is only incident with four arcs of S there can be no other vertices involved in S . So by the well-paired property, at least one of the pairs of arcs $(v, x_{0,0}), (v, y_{0,0})$ and $(x_{0,0}, v), (y_{0,0}, v)$ must be paired at v . It follows that S is a triangle on $\{v, x_{0,0}, y_{0,0}\}$. \square

Call S *eccentric* if it is not normal and not a triangle. If S is eccentric then $v \neq z_{i,j}$ for all $(i, j) \in \mathbb{Z}_2^2$, by Lemma 4.1. Hence $\chi^h S$ is as shown in Figure 4. (Remember that arcs must alternate in both colour and orientation, giving a unique way to navigate around this figure, or see Figure 6 below for an unravelled version.) Again the curved lines represent an odd number of alternating arcs (from $x_{0,0}$ to $x_{1,0}$ and from $x_{0,1}$ to $x_{1,1}$). Recall also that the vertices $x_{i,j}$ are not necessarily distinct.

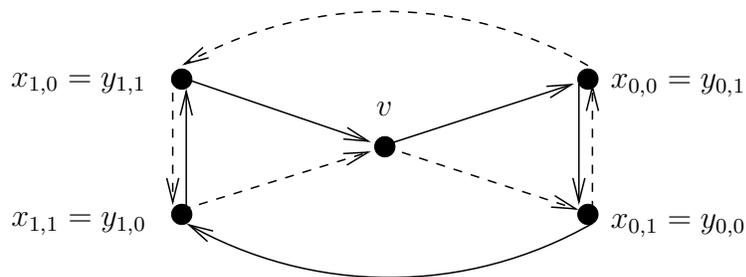


Figure 4: The 2-circuit $\chi^h S$ when S is eccentric

We describe how to process an eccentric 2-circuit in Section 4.2 and in Section 4.3 we explain how to process a triangle. This will complete the description of the canonical path from G to G' corresponding to the pairing ψ .

4.1 Decomposing a normal 2-circuit

Let S be a normal 2-circuit, with vertices labelled as in (9), where $(v, x_{0,0})$ is the lexicographically least arc in $A(S)$. Recall the notation $z_{i,j}$ defined before Lemma 4.1. A normal 2-circuit was depicted in Figure 2 but now we need a more detailed picture (Figure 5). Recall however that there can be as few as three arcs in the left or right half of this figure: for example, if there were only three arcs on the right then $y_{i,0} = x_{i+1,0}$ and $z_{i,0} = v$ for $i \in \mathbb{Z}_2$. Again the curved lines in Figure 5 represent an odd number of alternating arcs.

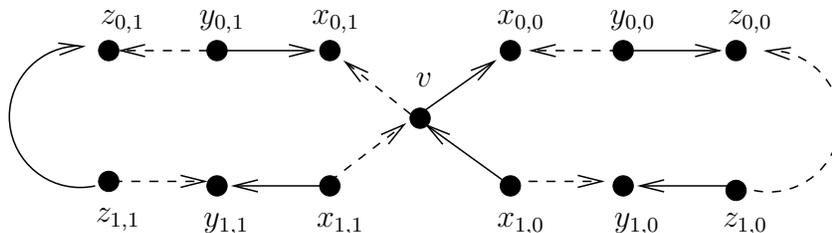


Figure 5: A normal 2-circuit $\chi^h S$, in more detail

Let (i, j) be the lexicographically least index such that $x_{i,j} \neq y_{i,j+1}$. (Here we use the ordering $0 < 1$ on \mathbb{Z}_2 .) Define the arc $a_{i,j} = (y_{i,j+1}, x_{i,j})$. The *shortcut arc* of S is $\zeta^i a_{i,j}$ (that is, it equals $a_{i,j}$ itself if $i = 0$ and equals the reversal of $a_{i,j}$ if $i = 1$).

Suppose that Z_J is the current digraph on the canonical path from G to G' before we start decomposing S . There are three cases, called (Na), (Nb), (Nc), where the ‘N’ stands for ‘normal’.

(Na) the shortcut arc $\zeta^i a_{i,j}$ belongs to $A(S)$.

(Nb) the shortcut arc $\zeta^i a_{i,j}$ does not belong to $A(S)$, and $\zeta^i a_{i,j}$ is not an arc of $\chi^{h+j} Z_J$.

(Nc) the shortcut arc $\zeta^i a_{i,j}$ does not belong to $A(S)$, and $\zeta^i a_{i,j}$ is an arc of $\chi^{h+j} Z_J$.

We consider these cases in order. (A more detailed description of the analogous process in the undirected case, with figures, can be found in [6] and may also be helpful.)

(Na) In case (Na), the 2-circuit S can be split into two 1-circuits, S_1 and S_2 . There are four subcases to consider, depending on which “half” of the 2-circuit contains the shortcut arc and whether the shortcut arc belongs to Z_J . In all subcases, the arcs of S_1 and S_2 form a partition of the arcs of S .

Once the two 1-circuits S_1 and S_2 have been identified, they are processed in that order, extending the canonical path from G to G' as

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}$$

after processing S_1 , and

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}, Z_{J+k+1}, \dots, Z_{J+k+\ell}$$

after processing S_2 .

(Na1) Suppose that S can be rewritten (allowing cyclic wrapping if necessary) as

$$v x_{i,j+1}, y_{i,j+1} z_{i,j+1} \cdots y_{i,j+1} x_{i,j} \cdots z_{i+1,j+1} y_{i+1,j+1} x_{i+1,j+1} v x_{i+1,j} \cdots x_{i,j}$$

and $\zeta^i a_{i,j} \notin A(\chi^{h+j} Z_J)$. Split S into two 1-circuits

$$\begin{aligned} S_1 &= v x_{i,j+1} y_{i,j+1} z_{i,j+1} \cdots y_{i,j+1} x_{i,j}, \\ S_2 &= v x_{i+1,j+1} y_{i+1,j+1} z_{i+1,j+1} \cdots x_{i,j} y_{i,j} \cdots y_{i+1,j} x_{i+1,j}. \end{aligned}$$

(Na2) Suppose that S can be rewritten (allowing cyclic wrapping if necessary) as

$$v x_{i,j+1} y_{i,j+1} z_{i,j+1} \cdots x_{i,j} y_{i,j+1} \cdots z_{i+1,j+1} y_{i+1,j+1} x_{i+1,j+1} v x_{i+1,j} \cdots x_{i,j}$$

and $\zeta^i a_{i,j} \in A(\chi^{h+j} Z_J)$. Split S into two 1-circuits

$$\begin{aligned} S_1 &= v x_{i,j+1} y_{i,j+1} z_{i,j+1} \cdots x_{i,j}, \\ S_2 &= v x_{i+1,j+1} y_{i+1,j+1} z_{i+1,j+1} \cdots y_{i,j+1} x_{i,j} y_{i,j} \cdots y_{i+1,j} x_{i+1,j}. \end{aligned}$$

(Na3) Suppose that S can be rewritten (allowing cyclic wrapping if necessary) as

$$v x_{i,j+1} \cdots x_{i+1,j+1} v x_{i+1,j} y_{i+1,j} z_{i+1,j} \cdots y_{i,j+1} x_{i,j} \cdots z_{i,j+1} y_{i,j+1} x_{i,j+1}$$

and $\zeta^i a_{i,j} \notin A(\chi^{h+j} Z_J)$. Split S into two 1-circuits

$$\begin{aligned} S_1 &= v x_{i,j+1} y_{i,j+1} x_{i,j} \cdots z_{i,j} y_{i,j} x_{i,j}, \\ S_2 &= v x_{i+1,j+1} y_{i+1,j+1} \cdots y_{i,j+1} \cdots z_{i+1,j} y_{i+1,j} x_{i+1,j}. \end{aligned}$$

(Na4) Suppose that S can be rewritten (allowing cyclic wrapping if necessary) as

$$v x_{i,j+1} \cdots x_{i+1,j+1} v x_{i+1,j} y_{i+1,j} z_{i+1,j} \cdots x_{i,j} y_{i,j+1} \cdots z_{i,j} y_{i,j} x_{i,j}$$

and $\zeta^i a_{i,j} \in A(\chi^{h+j} Z_J)$. Split S into two 1-circuits

$$\begin{aligned} S_1 &= v x_{i,j+1} y_{i,j+1} \cdots z_{i,j} y_{i,j} x_{i,j}, \\ S_2 &= v x_{i+1,j+1} y_{i+1,j+1} \cdots y_{i,j+1} x_{i,j} \cdots z_{i+1,j} y_{i+1,j} x_{i+1,j}. \end{aligned}$$

(Nb) Now suppose that S is a normal 2-circuit, the shortcut arc $\zeta^i a_{i,j}$ is not an arc of S and $\zeta^i a_{i,j}$ is not an arc of $\chi^{h+j} Z_J$. Then we can use the shortcut arc to give an alternating 4-cycle $v x_{i,j} y_{i,j+1} x_{i,j+1}$. First process this alternating 4-cycle using the switch $\zeta^i \chi^{h+j} [v x_{i,j}, y_{i,j+1} x_{i,j+1}]$, extending the canonical path by one step to give

$$G = Z_0, \dots, Z_J, Z_{J+1}.$$

(Call this step the *shortcut switch*.) Now $\zeta^i a_{i,j}$ is an arc of $\chi^{h+j} Z_{J+1}$ and we can form a 1-circuit S_1 from S , specifically

$$S_1 = vx_{i+1,j+1}y_{i+1,j+1} \cdots y_{i,j+1}x_{i,j}y_{i,j} \cdots y_{i+1,j}x_{i+1,j}. \quad (10)$$

Process this 1-circuit (as described in Section 3.3) to extend the canonical path further, giving

$$G = Z_0, \dots, Z_J, Z_{J+1}, Z_{J+2}, \dots, Z_{J+k}.$$

Note that $\zeta^i a_{i,j}$ is not an arc of $\chi^{h+j} Z_{J+k}$ after the 1-circuit S_1 has been processed, so it has been restored to the same state as in $\chi^{h+j} Z_J$, before the processing of the 2-circuit S began.

- (Nc) Finally assume that S is a normal 2-circuit, the shortcut arc $\zeta^i a_{i,j}$ is not an arc of S and $\zeta^i a_{i,j}$ is an arc of $\chi^{h+j} Z_J$. Then the shortcut arc completes the 1-circuit S_1 defined in (10), which is processed (as described in Section 3.3). This extends the canonical path to give

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}.$$

Last we process the alternating 4-cycle $vx_{i,j}y_{i,j+1}x_{i,j+1}$, using the shortcut switch $\zeta^i \chi^{h+j}[vx_{i,j}, y_{i,j+1}x_{i,j+1}]$, extending the canonical path by one step to give

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}, Z_{J+k+1}.$$

Note that $\zeta^i a_{i,j}$ is not an arc of $\chi^{h+j} Z_{J+k}$ but it is an arc of $\chi^{h+j} Z_{J+k+1}$, so it has been restored to the same state as in $\chi^{h+j} Z_J$.

4.2 Decomposing an eccentric 2-circuit

Now we may assume that S is an eccentric 2-circuit. Then $y_{i,j} = x_{i,j+1}$ for all $(i, j) \in \mathbb{Z}_2^2$, by definition, and $v \neq z_{i,j}$ for all $(i, j) \in \mathbb{Z}_2^2$, by Lemma 4.1. Call $(z_{1,0}, v)$ the *eccentric arc*. Note that $z_{1,0} \notin \{x_{1,0}, x_{1,1}\}$ which is the set of in-neighbours of v on S . Hence the eccentric arc is never an arc of S , so that the analogue of Case (Na) never arises. The remaining possibilities are below, called Case (Ea) and (Eb) (these are similar to cases (Nb) and (Nc) for normal 2-circuits, respectively).

- (Ea) Suppose that $(z_{1,0}, v) \notin A(\chi^h Z_J)$. Then $z_{1,0}x_{1,1}x_{1,0}v$ forms an alternating 4-cycle which we process using the switch $\chi^h[z_{1,0}x_{1,1}x_{1,0}v]$, extending the canonical path by one step to give

$$G = Z_0, \dots, Z_J, Z_{J+1}.$$

We call this step the *eccentric switch*. After performing the eccentric switch we have the 2-circuit

$$S' = vx_{0,0} \cdots z_{1,0}vx_{1,1}x_{1,0} \cdots x_{0,0}x_{0,1}. \quad (11)$$

Indeed, since $z_{1,0} \neq x_{1,0}$ it follows that S' is a normal 2-circuit, which we can process using the method described in Section 4.1. This extends the canonical path as

$$G = Z_0, \dots, Z_J, Z_{J+1}, Z_{J+2} \dots, Z_{J+1+k}.$$

Note that $(z_{1,0}, v) \notin A(\chi^h Z_{J+1+k})$, so the eccentric arc has been restored to the same state as in $\chi^h Z_J$, before the processing of S began.

(Eb) Suppose that $(z_{1,0}, v) \in A(\chi^h Z_J)$. Then S' defined in (11) is a normal 2-circuit which we first process using the method described in Section 4.1. This extends the canonical path as

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}.$$

Then $z_{1,0}x_{1,1}x_{1,0}v$ forms an alternating 4-cycle which we process using the eccentric switch $\chi^h[z_{1,0}x_{1,1}x_{1,0}v]$, extending the canonical path by one step to give

$$G = Z_0, \dots, Z_J, Z_{J+1}, \dots, Z_{J+k}, Z_{J+k+1}.$$

Now $(z_{1,0}, v) \in A(\chi^h Z_{J+1+k})$, so the eccentric arc has been restored to the same state as in $\chi^h Z_J$.

This procedure for still works even for eccentric 2-circuits with only five vertices. These arise when $z_{i,j} = x_{i+1,j+1}$ for all $(i, j) \in \mathbb{Z}_2^2$ (matching Figure 4 with both curved lines replaced by one arc each.)

The following information will be needed when analysing the flow.

Lemma 4.2. *Let S be an eccentric 2-circuit with the labelling of (9) and let S' be the normal 2-circuit used to process S . Suppose that S' falls into case (Nb) or (Nc). Then the following all hold:*

- (i) *Neither of the arcs $(v, x_{0,1})$, $(x_{1,1}, v)$ are involved in the eccentric switch.*
- (ii) *Using the labelling from Figure 4, the shortcut arc used to process S' is $(z_{1,0}, x_{1,0})$ and the shortcut switch is $[x_{1,1}x_{1,0}z_{1,0}v]$.*
- (iii) *The eccentric arc is involved in the shortcut switch and does not lie on the 1-circuit used to process S' .*

Proof. Recall that the eccentric arc is $(z_{1,0}, v)$. The first statement is immediate as the eccentric switch processes the alternating 4-cycle $z_{1,0}x_{1,1}x_{1,0}v$.

For the remainder of the proof, we use labels $\hat{x}_{i,j}, \hat{y}_{i,j}, \dots$ to denote the labelling of S' obtained as in (9). See Figure 6. As S is eccentric we have $y_{i,j+1} = x_{i,j}$ for all $(i, j) \in \mathbb{Z}_2^2$. By choice of the eccentric switch we have $\hat{x}_{i,j} = x_{i,j}$ and $\hat{y}_{i,j} = y_{i,j}$ for $(i, j) \neq (1, 0)$, while $\hat{x}_{1,0} = z_{1,0}$.

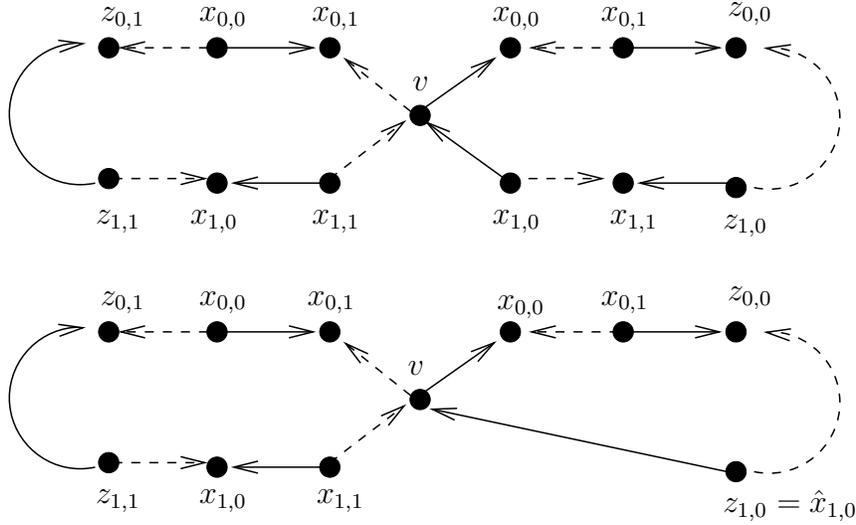


Figure 6: An eccentric 2-circuit $\chi^h S$ (above) and the normal 2-circuit $\chi^h S'$ used to process it (below)

Now $z_{1,0} \neq x_{1,0}$ since $z_{1,0}x_{1,1}x_{1,0}$ is a contiguous substring of S . Hence $(1, 0)$ is the lexicographically least (i, j) such that $\hat{x}_{i,j} \neq \hat{y}_{i,j+1}$. It follows that the shortcut arc is $(\hat{x}_{1,0}, \hat{y}_{1,1}) = (z_{1,0}, x_{1,0})$. Notice that the eccentric arc is incident with the shortcut arc at $z_{1,0}$ (with the same orientation). Furthermore, the shortcut switch involves a switch to the alternating 4-cycle

$$v\hat{x}_{1,1}\hat{y}_{1,1}\hat{x}_{1,0} = vx_{1,1}x_{1,0}z_{1,0}$$

which includes the eccentric arc. Specifically, the switch is $[x_{1,1}x_{1,0}z_{1,0}v]$, proving (ii). Since the eccentric arc $(z_{1,0}, v)$ is one of the arcs involved in the shortcut switch, it does not lie on the 1-circuit used to process S' . This establishes (iii), completing the proof. \square

4.3 Processing a triangle

Now suppose that S is a triangle, with vertices labelled v_0, v_1, v_2 where v_0 is the least vertex on S and (v_0, v_1) is an arc in the current digraph Z_J . Define the sets $\mathcal{W}^{(i,j)} = \mathcal{W}^{(i,j)}(\{v_0, v_1, v_2\}, Z_J)$ for $(i, j) \in \mathbb{Z}_2^2$. There are two cases, depending on whether a useful neighbour of S exists.

(T1) First suppose that there exists a useful neighbour of S . Let x be the minimum useful neighbour of S , and set (i, h) according to the first condition in this list which is satisfied by x :

$$(i, h) = \begin{cases} (0, 0) & \text{if } x \text{ is an out-neighbour of exactly one vertex of } S, \\ (0, 1) & \text{if } x \text{ is an out-neighbour of exactly two vertices of } S, \\ (1, 0) & \text{if } x \text{ is an in-neighbour of exactly one vertex of } S, \\ (1, 1) & \text{if } x \text{ is an in-neighbour of exactly two vertices of } S. \end{cases}$$

Then the sequence of three switches given by LaMar [15, left half of Figure 2] can be used to process S . For completeness we describe these switches here. Relabel the vertices of the triangle with a, b, c so that

- $\zeta^i(a, x) \in A(\chi^h Z_J)$,
- $\zeta^i(b, x) \notin A(\chi^h Z_J), \zeta^i(c, x) \notin A(\chi^h Z_J)$,
- $\zeta^i(a, b), \zeta^i(b, c), \zeta^i(c, a) \in A(\chi^h Z_J)$.

(Once x, i, h are chosen using the above procedure, the labelling of the triangle is uniquely determined.) Then the sequence of switches

$$\zeta^i \chi^h [axbc], \quad \zeta^i \chi^h [bxca], \quad \zeta^i \chi^h [abcx]$$

processes the triangle and restores all arcs between x and the triangle to their original state. See Figure 7 for the case $(i, h) = (0, 0)$: the diagram for the other cases can be obtained by reversing all arcs if $i = 1$, and/or by exchanging solid lines and dashed lines if $h = 1$. Call arcs $\zeta^i(a, x), \zeta^i(b, x), \zeta^i(c, x)$ the *auxilliary arcs*.

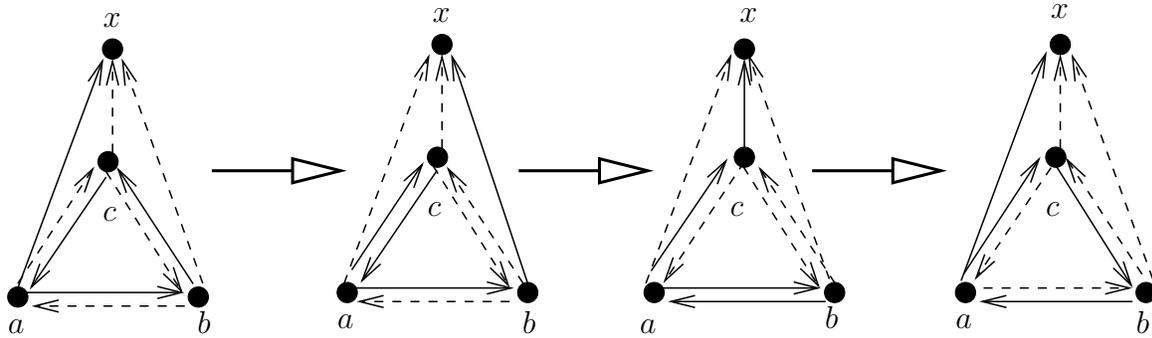


Figure 7: Processing a triangle using a useful neighbour

Use this sequence of switches to process the triangle, extending the canonical path as

$$G = Z_0, \dots, Z_J, Z_{J+1}, Z_{J+2}, Z_{J+3}.$$

(T2) Suppose that there is no useful neighbour of S in Z_J . Then using Lemma 2.3, there must exist a useful arc for S . Let (x, y) be the lexicographically least such arc. Recall that (x, y) satisfies one of the properties (U1), (U2) given just before Lemma 2.3. Define

$$h = \begin{cases} 0 & \text{if (U1) holds,} \\ 1 & \text{if (U2) holds.} \end{cases}$$

Then $(x, y) \in A(\chi^h Z_J)$ with $x \in \mathcal{W}^{(h,h)} \cup \mathcal{W}^{(h,h+1)}$ and $y \in \mathcal{W}^{(h,h)} \cup \mathcal{W}^{(h+1,h)}$. Relabel the vertices of the triangle as a, b, c , where $a = v_0$ and $(a, b) \in A(\chi^h Z_J)$. (Once h

is defined, this labelling is completely determined.) The sequence of switches given by LaMar [15, right side of Figure 2] will be used to process S . For completeness we give this sequence of switches in our notation:

$$\chi^h[xyab], \quad \chi^h[aybc], \quad \chi^h[byca], \quad \chi^h[xbcy].$$

These switches are also displayed in Figure 8 in the case that $h = 0$: the diagram for $h = 1$ can be obtained by exchanging solid lines and dashed lines. The arcs (x, y) , (x, b) , (a, y) , (b, y) , (c, y) are called *auxilliary arcs*. Use this sequence of switches to process the triangle, extending the canonical path as

$$G = Z_0, \dots, Z_J, Z_{J+1}, Z_{J+2}, Z_{J+3}, Z_{J+4}.$$

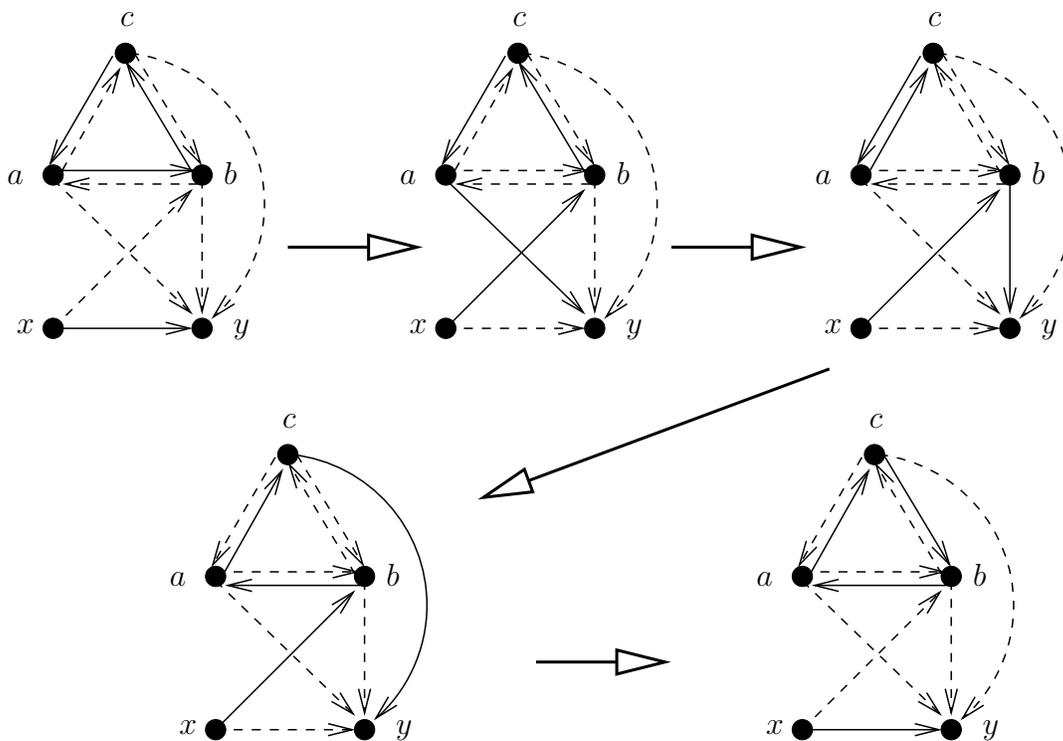


Figure 8: Processing a triangle using a useful arc

5 Analysing the flow

We now analyse the multicommodity flow so that we can apply Lemma 1.5 to give a bound on the second-largest eigenvalue of the switch chain. In this section we assume that $1 \leq d = d(n) \leq n/2$ for all n . This implies the general result for any $(d(n))$, by complementation where necessary.

Fix a pairing $\psi \in \Psi(G, G')$ and let $\gamma_\psi(G, G')$ be the canonical path from G to G' with respect to ψ . Let (Z, Z') be any transition on $\gamma_\psi(G, G')$, and let S be the raw 1-circuit or raw 2-circuit which is currently being processed. (That is, the transition (Z, Z') is performed while processing S .) Let Z_J be the digraph on the canonical path from G to G' just before the processing of S began. Any arc which does not belong to S but which has distinct status in Z and Z_J is called an *interesting arc* for Z with respect to (G, G', ψ) . (That is, the arc does not belong to S but is present in Z but absent in Z_J , or vice-versa.) The only arcs that can be interesting are:

- odd chords which are switched while processing a 1-circuit,
- the shortcut arc and/or eccentric arc, switched while processing a normal or eccentric 2-circuit,
- auxilliary arcs which are switched while processing a triangle.

We will label an interesting arc by -1 (respectively, 2) if it is absent (respectively, present) in Z_J but present (respectively, absent) in Z . (The reason for this choice of labels will be made clear shortly.)

Interesting arcs play a key role in our analysis. The following lemma describes the possible subdigraphs of Z that can be formed by interesting arcs in Z . It proves that the labelled digraph consisting of the interesting arcs is a subdigraph of one of the eight labelled digraphs shown in Figure 9, up to symmetries. Here $\{\mu, \nu\} = \{-1, 2\}$ and $\{\xi, \omega\} = \{-1, 2\}$ independently, giving four symmetries obtained by exchanging these pairs. Furthermore, ζ may also be applied to reverse the orientation of all arcs. Hence each digraph shown in Figure 9 represents up to eight possible digraphs. Note, the label for a given arc is shown next to the head of that arc.

Lemma 5.1. *Let Z be a digraph which lies on the canonical path from G to G' with respect to the pairing $\psi \in \Psi(G, G')$. There are at most five interesting arcs in Z with respect to (G, G', ψ) . The digraph consisting of the interesting arcs in Z is a subdigraph of one of the digraphs in Figure 9. If there are five interesting arcs then the following statements all hold:*

- (i) *There exists a vertex w which is the head (respectively, tail) of three interesting arcs, and these three interesting arcs do not all have the same label.*
- (ii) *There is a fourth interesting arc which has w as tail (respectively, head). Let u be the head (respectively, tail) of the fourth interesting arc.*
- (iii) *The fifth interesting arc is not incident with w but has u as its head (respectively, tail).*

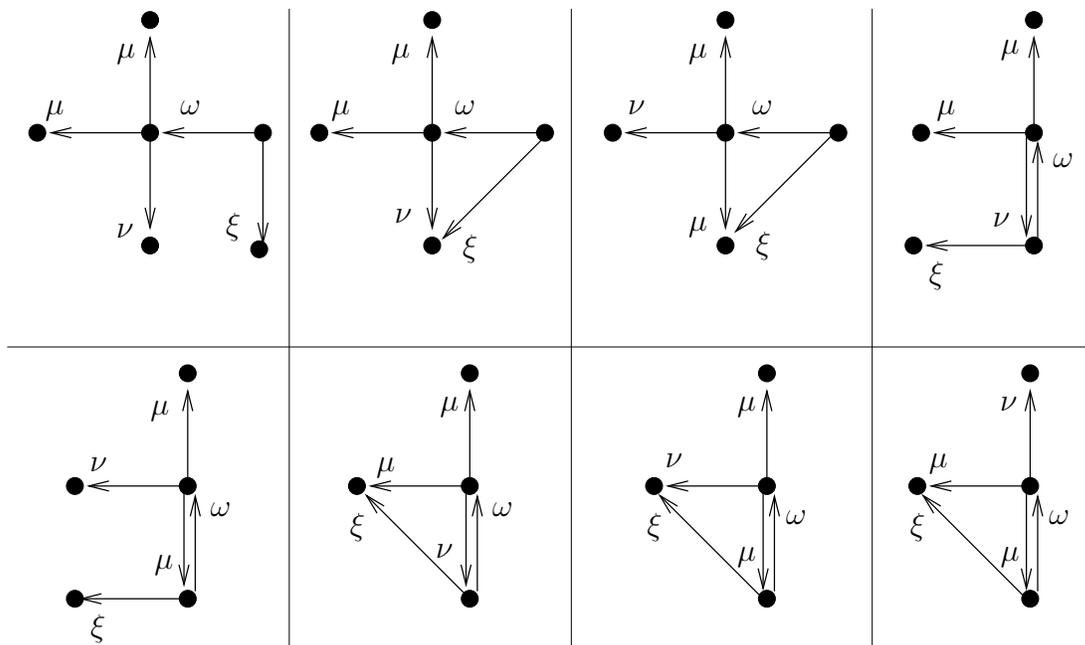


Figure 9: Possible configurations of interesting arcs, up to symmetries

Proof. While processing a triangle, at most three interesting arcs are used, namely the two or three auxiliary arcs. It follows from Figures 7, 8 that the auxiliary arcs always form subgraphs of a configuration from Figure 5.1.

When processing a normal 2-circuit, the situation is very similar to that in [6], with at most four interesting arcs. Up to three interesting arcs arise from the processing of a 1-circuit. They are all odd chords, and hence are all incident with the start-vertex of the 1-circuit with consistent orientation. However, they do not all have the same label. The fourth interesting arc corresponds to the shortcut arc, which may be labelled -1 or 2 and may be incident with none, one or two of the other interesting arcs (but not incident with the start vertex of the 1-circuit).

The fifth possible interesting arc is the eccentric arc, in the case that we are processing an eccentric 2-circuit S . Let S' be the normal 2-circuit containing the eccentric arc which is used to process S . If S' falls into case (Na) then the eccentric arc may be an interesting arc for part (either the start or end) of the processing of S_1 , the 1-circuit which contains it. But the configuration of interesting arcs in this case looks just the same as those which may arise from the processing of a normal 1-circuit, since the eccentric arc is involved in either the first switch of the last phase or the last switch of the first phase, and hence plays the same role as an interesting arc left over from a previous phase. However, if S' falls into case (Nb) or (Nc) then by Lemma 4.2 (iii), the eccentric arc does not lie on the 1-circuit S_1 which arises from S' . But the eccentric arc be an interesting arc throughout the processing of S_1 . Hence S_1 may have up to five interesting arcs, including the shortcut arc and the eccentric arc. In this case the eccentric arc is incident with the start-vertex v of S_1 (which equals the start-vertex of S) and it has the *opposite* orientation to the other

interesting arcs incident with v , if any. Let u be the endvertex of the eccentric arc which is not v . If the shortcut arc is present then it must be incident with the eccentric arc at u , with consistent orientation. This completes the proof. \square

Now identify a digraph with its $n \times n$ adjacency matrix (which has zero diagonal), and define the $n \times n$ matrix L by $L + Z = G + G'$. Entries of L belong to $\{-1, 0, 1, 2\}$. We may also think of L as the complete digraph on $[n]$ with each arc labelled by the corresponding entry of L . An arc in L is called *bad* if its label is -1 or 2 . Note that L is independent of ψ . Call L an *encoding* for Z with respect to (G, G') . Note that an arc receives label -1 if it is absent in both G and G' but is present in Z , while an arc receives label 2 if it is present in both G and G' but is absent in Z . Thus arcs in the symmetric difference $G \Delta G'$ are never bad arcs. Furthermore, every bad arc is an interesting arc, and an interesting arc is bad if and only if it does not belong to the symmetric difference $G \Delta G'$. This observation will be used many times in our analysis. In particular, it means that the digraph of bad arcs in an encoding L for Z is a subdigraph of one of the digraphs in Figure 9. This also explains our choice of labels for interesting arcs, since a bad arc with label 2 (respectively, -1) is also an interesting arc with label 2 (respectively, -1).

In the undirected setting [6, Lemma 1] it is always possible to uniquely recover (G, G') if (Z, Z') , L and ψ are known. We prove a slightly weaker result in the directed setting.

Lemma 5.2. *Given (Z, Z') , L , ψ , there are at most four possibilities for (G, G') such that (Z, Z') is a transition along the canonical path from G to G' corresponding to ψ and L is an encoding for Z with respect to (G, G') .*

Proof. The matrix $G + G'$ equals $Z + L$. From this matrix we can identify all arcs which are present in both G and G' (entries with value 2 in $G + G'$) and all arcs which are absent in both G and G' (entries with value 0 in $G + G'$). We can also identify the symmetric difference $H = G \Delta G'$, corresponding to entries with value 1 in $G + G'$. It remains to assign colours blue and red to the arcs of H so that blue arcs come from G and red arcs come from G' .

From the uncoloured version of H together with ψ we can construct the circuit decomposition \mathcal{C} . Let \mathcal{S} be the sequence of raw 1-circuits and raw 2-circuits obtained by decomposing the circuits in \mathcal{C} in order, as described in Section 3.2. The elements of \mathcal{S} are pairwise arc-disjoint and their union is H .

Suppose that the transition (Z, Z') deletes the arcs (α, β) , (δ, γ) and replaces them with (α, γ) , (δ, β) . Call (α, β) , (α, γ) , (δ, β) , (δ, γ) the *switch arcs*. We classify transitions along the canonical paths into three types as follows:

Type 1: the transition is any step in the processing of a 1-circuit used to process $S \in \mathcal{S}$. At least one of the switch arcs belong to S . (This includes the case of a raw 1-circuit, in which case the 1-circuit equals S .)

Type 2: the transition is a shortcut switch or an eccentric switch used while processing the normal or eccentric 2-circuit $S \in \mathcal{S}$. At least two of the switch arcs belong to S .

Type 3: the transition is a step in the processing of a triangle $S \in \mathcal{S}$. At least one of the four switch arcs belong to S .

In all cases, at least one of the switch arcs belongs to the element $S \in \mathcal{S}$ currently being processed. Therefore, there are at most four possibilities for S , namely, at most one possibility for each switch arc. (This follows as elements of \mathcal{S} are pairwise arc-disjoint.)

Now fix one of the (at most four) possibilities for S . We will show that given this choice (or guess) for S , we can uniquely determine (G, G') by colouring the edges of H . Note that if S is a 2-circuit, its labelling (as in Figure 5) can be determined uniquely. Hence we can determine whether S is normal, eccentric or a triangle.

Furthermore, in the first two cases we can identify exactly which arcs will be used as odd chords, shortcut arcs or eccentric arcs during the processing of S . We now claim that if S is a triangle then we can uniquely determine the useful neighbour x or the useful arc (x, y) which is used to process S , and hence identify all auxiliary arcs used while processing S . To see this, note that when processing a triangle, each switch involves either two or three vertices of the triangle. If all three vertices of the triangle are involved in the switch then the other vertex is either a useful neighbour, or an endvertex of a useful arc. Fix one orientation around the triangle and call it “clockwise”, with the opposite orientation called “anticlockwise”. Consider the number of clockwise and anticlockwise arcs on the triangle in Z and Z' : if they are equal in Z or in Z' then we are using a useful arc and otherwise we are using a useful neighbour. (See Figures 7, 8.) In the latter situation it is easy to identify the useful neighbour x : it is the only vertex involved in the switch which does not belong to S . This determines the auxiliary arcs (their orientation matches the orientation of the switch arcs at x). Now suppose that we are using an useful arc (x, y) . Then we are in case (T2), which means that no useful neighbour of S existed at the start of processing S . Then y is the only vertex incident with the switch arcs which does not belong to S , and x and y are the unique vertices in Z which are useful neighbours of S . That is, x and y are the only vertices not in S which do not belong to the set $\cup_{(i,j) \in \mathbb{Z}_2^2} \mathcal{W}^{(i,j)}(S, Z)$. If only two vertices of the triangle are involved in the switch then the unique switch arc which is not incident with either of these vertices is the useful arc, and the switch is the first or last in processing S . This shows that all auxiliary arcs for S can be identified, as claimed.

Suppose that S comes from the decomposition of the circuit $C_r \in \mathcal{C}$. The digraph induced by all interesting arcs contains no circuits, as can be seen from Figure 9. Hence for any $\ell \neq r$ we can find at least one arc on C_ℓ which is not an interesting arc for S : call this a *helpful arc* for C_ℓ . Colour the helpful arc for C_ℓ blue if it does not belong to Z and $\ell < r$, or if it does belong to Z and $\ell > r$; otherwise colour it red. Then the colouring of the rest of C_ℓ is forced, since colours alternate around the circuit. In the same way we can assign colours to the arcs of every raw 1-circuit and raw 2-circuit obtained in the decomposition of C_r , other than the element $S \in \mathcal{S}$ (which we have assumed is) being switched in the current transition (Z, Z') . It remains to explain how to assign colours to the arcs of S .

If S is a triangle then (Z, Z') is a Type 3 transition. By observing the number of clockwise and anticlockwise arcs in Z and Z' , we can determine the orientation of the

triangle in G and in G' and hence assign colours to the arcs in S .

Hence for the remainder of the proof we can assume that S is either a 1-circuit or a normal or eccentric 2-circuit. Therefore the vertices $\alpha, \beta, \gamma, \delta$ all belong to S and without loss of generality $\alpha = \min\{\alpha, \beta, \gamma, \delta\}$ is the start-vertex of S (since the start-vertex is involved in every switch). The argument for 1-circuits and normal 2-circuits is very similar to that given in [6].

First suppose that (Z, Z') is a Type 1 transition, performed while processing the 1-circuit S' . Now S' may be a raw 1-circuit (in which case $S' = S \in \mathcal{S}$), or S' may have arisen while processing a raw (normal or eccentric) 2-circuit S . Hence S' may contain a shortcut arc (but note, no 1-circuit contains an eccentric arc, by Lemma 4.2). The arcs of S' can be partitioned into sections, separated from each other by two consecutive arcs that are either both in Z or both absent from Z . Each section contains at least two arcs, so at least one arc which is not the shortcut arc. Then at least one arc of S' is actually switched in the current transition, which allows us to label the section containing that arc as switched, and alternately label the remaining sections around S' as switched or unswitched. Then colour an arc of S' blue if it belongs to Z and is unswitched or it is absent from Z and is switched, and colour an arc of S' red if it belongs to Z and is switched or it is absent from Z and is unswitched. Finally, if S' is not raw but arose from a 2-circuit S , there is a unique way to colour the remaining arcs of S , keeping the colours alternating.

For the remainder of the proof we assume that (Z, Z') is Type 2 transition for S ; that is, a shortcut switch or an eccentric switch. Let Z_J denote the digraph on the canonical path from G to G' just before we start decomposing S . We consider three subcases.

Firstly, suppose that (Z, Z') is an eccentric switch. Then we know that the arcs $(v, x_{0,1})$ and $(x_{1,0}, x_{1,1})$ have the same status in Z_J . The former is not involved in the eccentric switch, by Lemma 4.2 (i), while the latter is involved in the eccentric switch. Hence if these two arcs have matching status in Z then we are in Case (Ea) and the current transition is the first in processing S . Colour the arcs of S according to Z : arcs of $S \cap Z$ should be coloured blue and the remaining arcs of S should be coloured red. If these arcs have opposite status in Z then we are in case (Eb) and the current transition is the last in processing S . Colour the arcs of S according to Z' : arcs of $S \cap Z'$ should be coloured red and the remaining arcs of S should be coloured blue.

We proceed similarly if (Z, Z') is a Type 2 transition for S which is a shortcut switch. For now, assume that S is a normal 2-circuit, so the shortcut switch does not involve the eccentric arc (if any). If the shortcut arc is $\zeta^i(y_{i,j+1}, x_{i,j})$ then the arcs $\zeta^i(x_{i+1,j}, v)$ and $\zeta^i(v, x_{i,j})$ have matching status in Z_J . The former arc is not involved in the shortcut switch but the latter arc is. Hence if these two arcs still have matching status in Z then we are in Case (Nb) and the current transition is the first in processing S . Colour the arcs of S according to Z , as described in the previous paragraph. If these two arcs have opposite status in Z (one absent and one present) then we have already processed the 1-circuit using the shortcut arc, so we are in case (Nc) and the current transition is the last in processing S . Colour the arcs of S according to Z' , as described in the previous paragraph.

The third subcase is that (Z, Z') is a shortcut switch which also involves an eccentric arc. Then S is an eccentric 2-circuit which has been decomposed into an eccentric switch and a normal 2-circuit S' , where S' contains the eccentric arc. The current transition is the shortcut switch which has arisen while processing S' . Now the arcs $(v, x_{0,1})$ and $(x_{1,1}, v)$ have matching status in Z_J . From Lemma 4.2 (i) we know that neither of these arcs are involved in the eccentric switch. The former arc is not involved in the shortcut switch but the latter arc is, by Lemma 4.2 (ii). Hence, these two arcs also have matching status at the start of processing S' , and we can colour the arcs of S according to Z if these arcs have matching status in Z , and colour the arcs of S according to Z' otherwise. This completes the proof. \square

Let $L(\alpha, \beta)$ denote the label of arc (α, β) in the encoding L . The arc-reversal operator ζ acts on an encoding L by mapping L to its transpose $\zeta L = L^T$. If $\zeta^i L(\alpha, \beta) = 2$ and $\zeta^i L(\alpha, \gamma) = -1$ for some $i \in \{0, 1\}$ then $(i, \alpha, \beta, \gamma)$ is called a *handy tuple* with *centre* α . If $(i, \alpha, \beta, \gamma)$ is handy and at most one of β, γ is the head, when $i = 0$ (respectively, the tail, when $i = 1$) of two bad arcs with distinct labels then $(i, \alpha, \beta, \gamma)$ is said to be *very handy*. We now collect together some structural information about bad arcs in encodings.

Lemma 5.3. *Given $G, G' \in \Omega_{n,d}$ with symmetric difference $H = G \Delta G'$, suppose that Z is a digraph on the canonical path from G to G' with respect to some pairing ψ . Let L be the corresponding encoding, defined by $L + Z = G + G'$. Then the following statements all hold.*

- (i) *Viewed as the arcs of a labelled digraph, the set of bad arcs in L forms a subdigraph of one of the digraphs given in Figure 9.*
- (ii) *If L contains a handy tuple then L contains a very handy tuple.*
- (iii) *If there are five bad arcs in L then there exists a very handy tuple $(i_1, \alpha_1, \beta_1, \gamma_1)$, and a handy tuple $(i_1, \alpha_2, \beta_2, \gamma_2)$ in L such that $\alpha_1 \neq \alpha_2$ and*

$$\{\zeta^{i_1}(\alpha_1, \beta_1), \zeta^{i_1}(\alpha_1, \gamma_1)\} \cap \{\zeta^{i_2}(\alpha_2, \beta_2), \zeta^{i_2}(\alpha_2, \gamma_2)\} = \emptyset. \quad (12)$$

- (iv) *If there are four bad arcs in L then there is at least one handy tuple in L .*
- (v) *If $d = 1$ then no arc in H is incident with a bad arc with label 2. If L has a bad arc with label 2 which is not the only bad arc in L then L has a handy tuple and L has at most three bad arcs, exactly one of which has label 2. Furthermore, if L has three bad arcs, exactly one of which has label 2 then each endvertex of the bad arc with label 2 is the centre of a handy tuple in L .*
- (vi) *If $d = 2$ then every vertex which has nonzero degree in H is the head of at most one bad arc with label 2 and is the tail of at most one bad arc with label 2.*

Proof. Every bad arc is interesting, so (i) follows immediately from Lemma 5.1. Statements (ii)–(iv) follow from (i), by inspection of Figure 9. Now the head (respectively,

tail) of a bad arc is also the head (respectively, tail) of an arc in $G - G'$ and an arc in $G' - G$, unless the bad arc is the useful arc used to process a triangle in case (T2). This follows from the definition of odd chords, shortcut arc, eccentric arc and auxilliary arcs in Sections 3.3, 4.1– 4.3. Hence (vi) and the first statement of (v) holds, since a bad arc with label 2 is present in both G and G' . Furthermore, inspection of Figure 8 shows that the remaining statements of (v) hold, completing the proof. \square

The notion of an encoding is now generalised to mean any $n \times n$ matrix L with entries in $\{-1, 0, 1, 2\}$ such that every row and column sum equals d . Given $Z \in \Omega_{n,d}$, an encoding L is called Z -valid if every entry of $L + Z$ belongs to $\{0, 1, 2\}$ and L, Z, H satisfy statements (i)–(vi) of Lemma 5.3, where H is the digraph defined by the entries of $L + Z$ which equal 1. We also define the set $\mathcal{F}(L)$ of all bad arcs in L by

$$\mathcal{F}(L) = \{(i, j) \in [n]^2 \mid L(i, j) \in \{-1, 2\}\}.$$

Lemma 5.4. *Let $Z \in \Omega_{n,d}$ and let L be a Z -valid encoding. Suppose that L' is another encoding such that $\mathcal{F}(L') \subseteq \mathcal{F}(L)$. Then L' is also Z -valid.*

Proof. If $L'(i, j) = -1$ then $L(i, j) = -1$ and hence $Z(i, j) = 1$, as L is Z -valid. Similarly, if $L'(i, j) = 2$ then $L(i, j) = 2$ and hence $Z(i, j) = 0$. This shows that every entry of $L' + Z$ belongs to $\{0, 1, 2\}$. Checking properties (i)–(vi) of Lemma 5.3 we see that they all hold for L', Z, H' , completing the proof. \square

Switches can be applied to encodings, as follows. By definition, the sum of all labels on arcs with head v add up to d , and the sum of all labels on arcs with tail v add up to d , for all vertices v . If x, y, z, w are vertices with $L(x, y) > -1$, $L(w, z) > -1$, $L(x, z) < 2$ and $L(w, y) < 2$ then we may perform the switch $[xywz]$ by decreasing $L(x, y)$ and $L(w, z)$ by one and increasing $L(x, z)$ and $L(w, y)$ by one, giving a new encoding L' .

Lemma 5.5. *Let $Z \in \Omega_{n,d}$. Given a Z -valid encoding, one can obtain a digraph (with no bad arcs) using at most three switches.*

Proof. Let L be a Z -valid encoding and let H be the digraph given by the entries of $L + Z$ which equal 1.

First suppose that L contains a handy tuple. Then L contains a very handy tuple, by Lemma 5.3 (ii). If L contains a very handy tuple $(i_1, \alpha_1, \beta_1, \gamma_1)$ and a handy tuple $(i_2, \alpha_2, \beta_2, \gamma_2)$ such that $\alpha_1 \neq \alpha_2$ and (12) holds, then we choose $(i, \alpha, \beta, \gamma)$ to be the very handy tuple $(i_1, \alpha_1, \beta_1, \gamma_1)$. Otherwise, let $(i, \alpha, \beta, \gamma)$ be any very handy tuple in L .

If $i = 0$ (respectively, $i = 1$) then the sum of the labels on the bad arcs with β as head (respectively, tail) is strictly greater than the sum of the labels on the bad arcs with γ as head (respectively, tail). By construction, each row of L adds up to d and each column of L adds up to d . Hence γ is the head (respectively, tail) of strictly more good arcs (with label 1) than β . It follows that there exists a vertex δ such that $\zeta^i L(\delta, \gamma) = 1$ and $\zeta^i L(\delta, \beta) = 0$. Now we can perform the switch $\zeta^i[\alpha\beta\delta\gamma]$ to give an encoding L'

with $\zeta^i L'(\alpha, \beta) = \zeta^i L'(\delta, \beta) = 1$ and $\zeta^i L'(\alpha, \gamma) = \zeta^i L'(\delta, \gamma) = 0$. Note that L' is a Z -valid encoding by Lemma 5.4, and that $|\mathcal{F}(L')| = |\mathcal{F}(L)| - 2$. We call this operation a $(-1, 2)$ -switch.

Next suppose that no vertex is the head (respectively, tail) of two bad arcs with distinct labels, but that an arc exists in L with label 2. By Lemma 5.3 (iii), (iv), there are at most three bad arcs in L . Choose vertices α, β such that for some $i \in \{0, 1\}$ we have $\zeta^i L(\alpha, \beta) = 2$ and if $i = 0$ (respectively, $i = 1$) then α is the tail (respectively, head) of exactly one bad arc. (That such an α exists follows from Lemma 5.3 (i).) Let U be the set of vertices $x \neq \alpha$ with $\zeta^i L(\alpha, x) = 0$. Since α is the tail (respectively, head) of exactly $d - 2$ arcs labelled 1 and one arc labelled 2, it follows that $|U| = n - d \geq d$.

We claim that there exists a vertex $\gamma \in U$ which is not the head (respectively, tail) of a bad arc with label 2. If $d \geq 3$ then there are at most 3 vertices which are at the head (respectively, tail) of a bad arc labelled 2, and one of these is β . Since $\beta \notin U$ and $|U| \geq 3$, we can choose a vertex $\gamma \in U$ which is not the head of a bad arc with label 2, as claimed. If $d = 2$ then by Lemma 5.3 (vi), each vertex in H is head (respectively, tail) of at most one bad arc labelled 2. Hence there are at most 2 bad arcs labelled 2 in L , by Lemma 5.3 (i). Therefore at most one vertex other than β is the head (respectively, tail) of a bad arc labelled 2 in L . The claim then follows since $\beta \notin U$ and $|U| \geq 2$. If $d = 1$ then by Lemma 5.3 (v) there is exactly one bad arc in L , namely $\zeta^i(\alpha, \beta)$. Hence we can let γ be any element of U since $\beta \notin U$, and the claim follows as $|U| \geq 1$ in this case.

Now β is the head (respectively, tail) of at most $d - 2$ good arcs and γ is the head (respectively, tail) of at least d good arcs. Hence we can choose a vertex δ such that $\zeta^i L(\delta, \beta) = 0$ and $\zeta^i L(\delta, \gamma) = 1$. Perform the switch $\zeta^i[\alpha\beta\delta\gamma]$ to produce an encoding L' with $\zeta^i L'(\alpha, \beta) = \zeta^i L'(\alpha, \gamma) = \zeta^i L'(\delta, \beta) = 1$ and $\zeta^i L'(\delta, \gamma) = 0$. Then L' is Z -valid by Lemma 5.4, and $|\mathcal{F}(L')| = |\mathcal{F}(L)| - 1$. Call this operation a 2 -switch.

Finally, suppose that the only remaining bad arcs are labelled -1 . Let α and γ be vertices such that $\zeta^i L(\alpha, \gamma) = -1$ for some $i \in \{0, 1\}$, choosing α to be a vertex at the tail, when $i = 0$ (respectively head, when $i = 1$) of two bad arcs with label -1 , if such a vertex exists. Note that when L contains three bad arcs with label -1 then such a choice of α exists, by Lemma 5.3 (i). We claim that there exists a vertex β such that $\zeta^i L(\alpha, \beta) = 1$ but β is not the head, when $i = 0$ (respectively tail, when $i = 1$) of any bad arc. To see this, note that there are at least $d + 1 \geq 2$ choices for β , and there is at most one vertex which is at the head (respectively, tail) of a bad arc which is not incident with α , by choice of α . Hence we can avoid this vertex when choosing β , proving the claim. Then β is the head (respectively, tail) of exactly d good arcs, while γ is the head (respectively, tail) of at least $d + 1$ good arcs. Hence there is at least one way to choose a vertex δ such that $\zeta^i L(\delta, \beta) = 0$ and $\zeta^i L(\delta, \gamma) = 1$. Perform the switch $\zeta^i[\alpha\beta\delta\gamma]$ to produce an encoding L' , with $\zeta^i L'(\alpha, \beta) = \zeta^i L'(\delta, \gamma) = \zeta^i L'(\alpha, \gamma) = 0$, $\zeta^i L'(\delta, \beta) = 1$. Again Lemma 5.4 shows that L' is Z -valid, and $|\mathcal{F}(L')| = |\mathcal{F}(L)| - 1$. Call this operation a (-1) -switch.

If the original encoding L has five bad arcs then by Lemma 5.3 (iii), we can find a very handy tuple $(i_1, \alpha_1, \beta_1, \gamma_1)$ in L and perform the $(-1, 2)$ -switch $\zeta^{i_1}[\alpha_1\beta_1\gamma_1\delta_1]$, where δ_1 is a vertex found using the procedure above. It follows from (12) and Figure 5.3 (i) that $(i_2, \alpha_2, \beta_2, \gamma_2)$ is a very handy tuple in the resulting Z -valid encoding L' . Hence we

may perform the $(-1, 2)$ -switch $\zeta^{i_2}[\alpha_2\beta_2\gamma_2\delta_2]$ to transform L' into the Z -valid encoding L'' with at most one bad arc. At most one further switch is required to transform L'' into an encoding with no bad arcs. Thus at most 3 switches are needed to process L when L has five bad arcs.

Similarly, if L has four bad arcs then by Lemma 5.3 (iv), we can transform L into a Z -valid encoding L' with at most two bad arcs, using a $(-1, 2)$ -switch. At most two further switches are needed to produce an encoding with no bad arcs. Thus at most 3 switches are needed to process L when L has four bad arcs. Clearly, if L has at most 3 bad arcs then at most 3 switches are required. This completes the proof. \square

For $Z \in \Omega_{n,d}$ let $\mathcal{L}(Z)$ be the set of all Z -valid encodings. We obtain the following upper bound on $|\mathcal{L}(Z)|$ using a relatively simple proof. It is possible that an improved bound can be found using a more careful analysis, probably saving a factor of n .

Lemma 5.6. *For any $Z \in \Omega_{n,d}$ we have*

$$|\mathcal{L}(Z)| \leq 25 d^6 n^6 |\Omega_{n,d}|.$$

Proof. Fix $Z \in \Omega_{n,d}$ and let $L \in \mathcal{L}(Z)$ be a Z -valid encoding. By Lemma 5.5 there exists a sequence

$$L = L_0, L_1, \dots, L_r = A$$

where $A \in \Omega_{n,d}$ is a digraph with no bad arcs, $r \leq 3$ and each of L_1, \dots, L_r is Z -valid. We can turn this into a function $\varphi : \mathcal{L}(Z) \rightarrow \Omega_{n,d}$ by performing these switches in a canonical way: as in Lemma 5.5 perform all $(-1, 2)$ -switches first, then all 2-switches, then all (-1) -switches, following the extra conditions described in Lemma 5.5 and breaking ties using lexicographic ordering on the 5-tuple $(i, \alpha, \beta, \gamma, \delta)$. It suffices to prove that $|\varphi^{-1}(A)| \leq 25 d^6 n^6$ for all $A \in \Omega_{n,d}$.

Now fix $A \in \Omega_{n,d}$. Define a *reverse X -switch* to be the reverse of a X -switch, for $X \in \{(-1, 2), -1, 2\}$. For an upper bound we count all encodings which can be obtained from A using at most three reverse switches, regardless of whether A is the canonical image of that encoding under φ . We will perform the reverse switchings in order: first the reverse (-1) -switches, if any, then any reverse 2-switches and finally any reverse $(-1, 2)$ -switches. Note that a reverse switching alters four entries of the current encoding, none of which are bad entries. So a bad arc created by a reverse switch will never be changed by a later reverse switch.

Fix an encoding $B \in \mathcal{L}(Z)$ (which may not have any bad arcs). Let $N_X(B)$ be the number of distinct 5-tuples $(i, \alpha, \beta, \gamma, \delta)$ which define a reverse X -switch that may be performed in B , for $X \in \{(-1, 2), -1, 2\}$. The result of each of the reverse switches counted by $N_X(B)$ is a Z -valid encoding. Our next task is to calculate upper bounds on $N_X(B)$ which hold for all encodings $B \in \mathcal{L}(Z)$.

We only perform (-1) -switches on encodings $B \in \mathcal{L}(Z)$ which have no bad arc with label 2. For such encodings we claim that

$$N_{-1}(B) \leq 2d^2n(n-2). \tag{13}$$

With notation as in Lemma 5.5, the factor of 2 counts the two choices of orientation $i \in \{0, 1\}$. We prove the bound assuming that $i = 0$, and the proof for $i = 1$ follows by symmetry. There are n choices for vertex α , and d choices for γ since $\zeta^i(\alpha, \gamma) \in A(Z)$ as B is Z -valid. Then choose $\beta \neq \alpha$ so that $B(\alpha, \beta) = 0$ and β is not the head of any bad arc. There are at most $n - 2$ choices for β since $\beta \notin \{\alpha, \gamma\}$. Then there are d choices for δ such that $B(\delta, \beta) = 1$, since β is the head of exactly d good arcs. This gives the claimed bound on $N_{-1}(B)$ when B has no bad arcs labelled 2.

Now suppose that $B \in \mathcal{L}(Z)$ may contain bad arcs with distinct labels, but no vertex is the head (respectively, tail) of two bad arcs with distinct labels in B . We also ensure that the reverse 2-switchings that we perform never create any such pair of bad arcs, in order to maintain the canonical order in which forward switches are performed. We claim that

$$N_2(B) \leq 2d(d - 1)^2n. \tag{14}$$

The factor of 2 counts the two choices of orientation $i \in \{0, 1\}$. We prove the bound assuming that $i = 0$, and the proof for $i = 1$ follows by symmetry. There are at most n choices for α which is not the tail of a bad arc labelled -1 . Then distinct out-neighbours β, γ of α in B can be chosen in at most $d(d - 1)$ ways such that β is not the head of a bad arc labelled -1 and γ is not the head of a bad arc labelled 2. (Note, α is the tail of at most d good arcs, since α is not the tail of any bad arc labelled -1 .) Then there are at most $d - 1$ choices for a neighbour δ of β in B , since β is the head of at most d good arcs. This gives the claimed bound on $N_2(B)$.

Finally, we claim that for all $B \in \mathcal{L}(Z)$ we have

$$N_{(-1,2)}(B) \leq 2d^2(d + 1)n. \tag{15}$$

Again, the factor of 2 counts the two choices of orientation $i \in \{0, 1\}$ and we assume $i = 0$ below, without loss of generality. There are n ways to choose a vertex α which may be the tail of at most one bad arc in B . There are d choices for γ , as B is Z -valid so $\zeta^i(\alpha, \gamma) \in A(Z)$. Then there are at most $d + 1$ choices for β such that β is an out-neighbour of α and is not the head of any arc labelled -1 . (There are at most d choices for β if there is no bad arc incident with α in B .) Finally, there are at most d choices for $\delta \neq \alpha$ such that $B(\delta, \beta) = 1$, since β is the head of at most d good arcs. (The d here arises since β may itself be the head of at most one bad arc in B , and the bad arc may be labelled -1 .) This gives the claimed bound on $N_{(-1,2)}(B)$.

Each sequence of reverse switches which may arise is given a type, defined by the corresponding sequence of labels in $\{-1, 2, (-1, 2)\}$. It follows from the proof of Lemma 5.5 that the only types of reverse switchings which occur are given by the following 9 sequences and all distinct subsequences of these (including the empty sequence):

$$\begin{array}{lll} [-1, (-1, 2), (-1, 2)], & [2, (-1, 2), (-1, 2)], & [-1, -1, (-1, 2)], \\ [-1, 2, (-1, 2)], & [2, 2, (-1, 2)], & [-1, -1, -1], \\ [-1, -1, 2], & [-1, 2, 2], & [2, 2, 2]. \end{array}$$

This gives 19 possible types in all. We calculate the contribution of a type by simply multiplying the upper bounds obtained in (13)–(15) corresponding to each reverse switch

in the sequence. (It is at this step that a more careful analysis may lead to an improved bound, but we are satisfied by the bound given by this simple calculation.) For example, the contribution from the type $[-1, (-1, 2), (-1, 2)]$ is

$$2d^2n(n-2)(2d^2(d+1)n)^2 = 8d^6(d+1)^2n^3(n-2).$$

Finally we simply sum the contribution from each of the 19 types and find that the resulting expression is bounded above by $25d^6n^6$, using the inequalities $1 \leq d \leq n/2$. This shows that

$$|\varphi^{-1}(A)| \leq 25d^6n^6,$$

completing the proof. □

For each pair (G, G') of distinct digraphs in $\Omega_{n,d}$, let $\mathcal{P}_{G,G'}$ be the set of $|\Psi(G, G')|$ canonical paths which we have defined from G to G' , one for each pairing $\psi \in \Psi(G, G')$. Let $\mathcal{P} = \cup_{G \neq G'} \mathcal{P}_{G,G'}$. Define

$$f(\gamma) = |\Omega_{n,d}|^{-2} |\Psi(G, G')|^{-1}$$

for each path $\gamma \in \mathcal{P}_{G,G'}$. Then

$$\sum_{\gamma \in \mathcal{P}_{G,G'}} f(\gamma) = |\Omega_{n,d}|^{-2} = \pi(G) \pi(G')$$

where π is the stationary distribution of the Markov chain, which is uniform on $\Omega_{n,d}$. Thus $f : \mathcal{P} \rightarrow [0, \infty)$ is a flow. We want to apply Lemma 1.5. First we bound $f(e)$ for all transitions e of the Markov chain.

Lemma 5.7. *For any transition $e = (Z, Z')$ of the Markov chain,*

$$f(e) \leq 100 d^{22} n^6 |\Omega_{n,d}|^{-1}.$$

Proof. Fix a transition $e = (Z, Z')$ of the Markov chain. Let (G, G') be a pair of distinct digraphs in $\Omega_{n,d}$ and suppose that e lies on $\gamma_\psi(G, G')$, the canonical path from G to G' corresponding to the pairing $\psi \in \Psi(G, G')$. From Z and (G, G') we can construct L and the digraph $H = Z \triangle L = G \triangle G'$. We colour arcs of H green if they belong to Z and yellow if the corresponding entry in L is 1. (Recall that the symmetric difference H consists of those arcs with entry 1 in $L + Z = G + G'$.)

From the pairing ψ we obtain the circuit decomposition \mathcal{C} of H , with colours alternating green, yellow almost everywhere. A vertex x is *bad* with respect to ψ if two arcs of the same colour are paired at x under ψ . If a vertex is not bad it is called *good*. Every bad vertex lies on the circuit currently being processed. Specifically, bad vertices may only be found incident to interesting arcs. Lemma 5.1 shows that there are at most five interesting arcs and at most six potentially bad vertices.

A yellow-yellow or green-green pair at a bad vertex x is called a *bad pair* with respect to ψ . Careful consideration of the possibilities reveals that there can be at most 16 bad pairs with respect to ψ . In the worst case, there are five interesting arcs which all belong to H . An interesting arc e which belongs to H creates two bad pairs in the circuit containing e , one at each endvertex of e (both of the same colour). A bad pair is also created in the current circuit C incident with each endvertex of each interesting arc, giving at most six further bad pairs. (The worked example in Section 6 gives an example of a digraph, Z_4 , containing the maximum number of bad pairs: see Figure 14.)

Note also that a bad vertex may be the head (respectively, tail) of at most two bad pairs of each colour. This follows from Lemma 5.1 since no vertex is head (respectively, tail) of more than two interesting arcs with the same label. Hence a bad vertex may be the head (respectively, tail) of at most four bad pairs in total. This is true even if there are some coincidences between the bad vertices, which may occur when the interesting arcs have one of the configurations other than the first one in Figure 9. To see this, note that for all the configurations in Figure 9, the only vertex which is the head (or tail) of more than two interesting arcs is v , the start-vertex of the current circuit, and v is always distinct from all other bad vertices.

Given the uncoloured digraph H , we can form a pairing ψ by pairing up all in-arcs around v and pairing up all out-arcs around v , for each vertex v . Let the set of all these pairings be $\Psi(H)$. Say that a pairing $\psi \in \Psi(H)$ is *consistent with L* if there are at most 16 bad pairs in the yellow-green colouring of H with respect to L , and at each vertex u and for each choice of orientation there are at most two bad pairs of each colour with that orientation at u . Let $\Psi'(H, L)$ be the set of all pairings ψ of H which are consistent with L . Given any (G, G') with $G\Delta G' = H$, any pairing $\psi \in \Psi(G, G')$ is consistent with the yellow-green colouring of H , as proved above. Therefore each triple (G, G', ψ) with $\psi \in \Psi(G, G')$ and $e \in \gamma_\psi(G, G')$ gives rise to at least one pair (L, ψ) with $L \in \mathcal{L}(Z)$ and $\psi \in \Psi'(H, L)$.

Conversely, we can start with $L \in \mathcal{L}(Z)$ and find an upper bound for $|\Psi'(H, L)|$. Once ψ and (Z, Z') are given, there are at most four possibilities for (G, G') with $e \in \gamma_\psi(G, G')$, by Lemma 5.2. Recall from (5) that

$$|\Psi(G, G')| = \prod_{v \in V} \theta_v! \phi_v!$$

where $2\theta_v$ is the in-degree of v in H and $2\phi_v$ is the out-degree of v in H . Similarly, each good vertex v contributes a factor $\theta_v! \phi_v!$ to $|\Psi'(H, L)|$, but a bad vertex may contribute more. The contributions from in-arcs and out-arcs are independent, so we consider only in-arcs below.

Recall that no vertex can be the head of more than two bad pairs of a given colour. First suppose that a vertex v is the head of $\theta_v + 2$ green arcs and $\theta_v - 2$ yellow arcs. Then v must be bad, with two bad green pairs and no bad yellow pairs. The number of ways to pair up the in-arcs around v is

$$3 \binom{\theta_v + 2}{4} (\theta_v - 2)! = \frac{(\theta_v + 2)(\theta_v + 1)}{8} \theta_v! \leq \theta_v^2 \cdot \theta_v! \leq d^2 \theta_v!$$

Next suppose that v is the head of $\theta_v + 1$ green arcs and $\theta_v - 1$ yellow arcs. Then v must be a bad vertex. Now v may be the head of two bad green pairs and one bad yellow pair, or v may be the head of one bad green pair and no bad yellow pairs. The number of ways to pair up the in-arcs around v with two bad green pairs and one bad yellow pair is

$$3 \binom{\theta_v + 1}{4} \binom{\theta_v - 1}{2} (\theta_v - 3)! = \frac{(\theta_v + 1)(\theta_v - 1)(\theta_v - 2)}{16} \theta_v! \leq \theta_v^3 \theta_v! \leq d^3 \theta_v!,$$

while the number of pairings of in-arcs around v with one bad green pair and no bad yellow pairs is

$$\binom{\theta_v + 1}{2} (\theta_v - 1)! = \frac{\theta_v + 1}{2} \theta_v! \leq \theta_v \theta_v! \leq d \theta_v!.$$

Finally, suppose that v is the head of θ_v arcs of each colour. Then v may be good, or it may be the head of one bad pair of each colour, or the head of two bad pairs of each colour. The number of pairings of in-arcs around v with two bad pairs of in-arcs of each colour is

$$9 \binom{\theta_v}{4}^2 (\theta_v - 4)! = \frac{\theta_v(\theta_v - 1)(\theta_v - 2)(\theta_v - 3)}{64} \theta_v! \leq \theta_v^4 \cdot \theta_v! \leq d^4 \theta_v!,$$

while the number of pairings of in-arcs around v with one bad pair of each colour is

$$\binom{\theta_v}{2}^2 (\theta_v - 2)! = \frac{\theta(\theta_v - 1)}{4} \theta_v! \leq \theta_v^2 \theta_v! \leq d^2 \theta_v!.$$

By symmetry, the same bounds hold for out-arcs and also hold after exchanging green and yellow. Since there are at most 16 bad pairs, it follows that

$$|\Psi'(H, L)| \leq d^{16} |\Psi(G, G')|. \tag{16}$$

Now write $\mathbf{1}(e \in \gamma_\psi(G, G'))$ to denote the indicator variable which is 1 if $e \in \gamma_\psi(G, G')$ and is 0 otherwise, for $(G, G') \in \Omega_{n,d}$ and $\psi \in \Psi(G, G')$. Then

$$\begin{aligned} |\Omega_{n,d}|^2 f(e) &= \sum_{(G,G')} \sum_{\psi \in \Psi(G,G')} \mathbf{1}(e \in \gamma_\psi(G, G')) |\Psi(G, G')|^{-1} \\ &\leq 4 \sum_{L \in \mathcal{L}(Z)} \sum_{\psi \in \Psi'(H,L)} \mathbf{1}(e \in \gamma_\psi(G, G')) |\Psi(G, G')|^{-1} \\ &\leq 4 \sum_{L \in \mathcal{L}(Z)} \sum_{\psi \in \Psi'(H,L)} |\Psi(G, G')|^{-1} \\ &\leq 4 \sum_{L \in \mathcal{L}(Z)} d^{16} \\ &\leq 100 d^{22} n^6 |\Omega_{n,d}|. \end{aligned}$$

The first inequality follows by Lemma 5.2, the third inequality follows from (16), and applying Lemma 5.6 gives the last inequality. This completes the proof. \square

We can now complete our argument by proving Proposition 2.5.

Proof of Proposition 2.5. For any transition $e = (Z, Z')$ of the switch chain we have

$$1/Q(e) = |\Omega_{n,d}|/P(Z, Z') = \binom{dn}{2} |\Omega_{n,d}|.$$

Therefore, by Lemma 5.7,

$$\rho(f) \leq 50d^{24} n^8. \tag{17}$$

Next, observe that $\ell(f) \leq dn$, since each step along a canonical path replaces at least one arc of G by an arc of G' . The result follows from Lemma 1.5. \square

6 An illustrative example

Let $(G, G') \in \Omega_{n,d}$ be any pair of digraphs with the symmetric difference H given in Figure 10, where vertices of degree 0 in H are not shown. To avoid congestion in the figure, some vertices are depicted as black rectangles. Solid arcs belong to G and dashed arcs belong to G' , so they play the role of blue and red arcs.

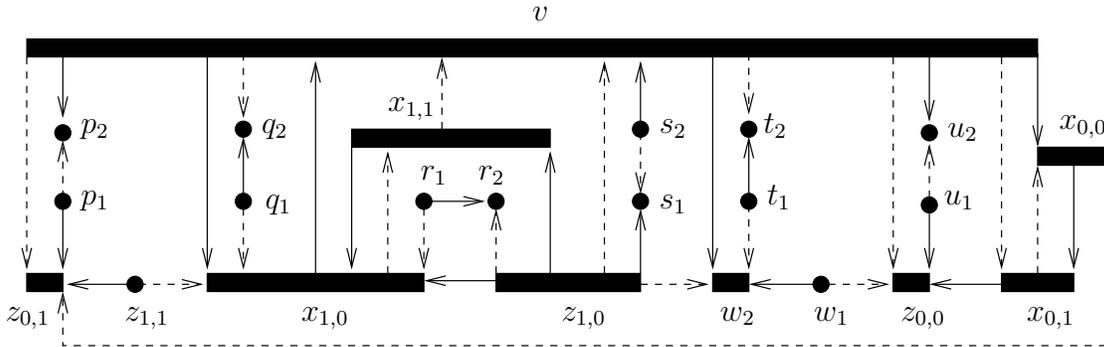


Figure 10: The symmetric difference H of G and G' .

Let ψ which be the pairing which produces the forward circuits

$$\begin{aligned} vx_{0,0}x_{0,1}z_{0,0}w_1w_2z_{1,0}x_{1,1}x_{1,0}vx_{1,1}x_{1,0}z_{1,1}z_{0,1}x_{0,0}x_{0,1}, \quad vp_2p_1z_{0,1}, \quad vx_{1,0}q_1q_2, \\ z_{1,0}x_{1,0}r_1r_2, \quad z_{1,0}vs_2s_1, \quad vw_2t_1t_2, \quad vu_2u_1z_{0,0}, \end{aligned} \tag{18}$$

in the given order. Set $Z_0 = G$ and start processing H . The first circuit to process is the eccentric 2-circuit

$$S = vx_{0,0}x_{0,1}z_{0,0}w_1w_2z_{1,0}x_{1,1}x_{1,0}vx_{1,1}x_{1,0}z_{1,1}z_{0,1}x_{0,0}x_{0,1}.$$

We have $(i, h) = (0, 0)$, and the eccentric arc $(z_{1,0}, v)$ does not belong to $A(S)$. Hence S falls into case (Ea) and we must first perform the eccentric switch $[z_{1,0}x_{1,1}x_{1,0}v]$. This

produces the next digraph Z_1 in the canonical path $\gamma_\psi(G, G')$. The eccentric arc $(z_{1,0}, v)$ has been used in the eccentric switch, so it is now an interesting arc. Initially it belonged to $G' - G$, and now it belongs to $Z_1 \cap G'$, so it does not belong to the current symmetric difference $Z_1 \Delta G'$. However, we include all interesting arcs in our figures, denoted by thicker arcs (either solid or broken, as appropriate). Hence Figure 11 shows the symmetric difference of Z_1 and G' , together with the eccentric arc.

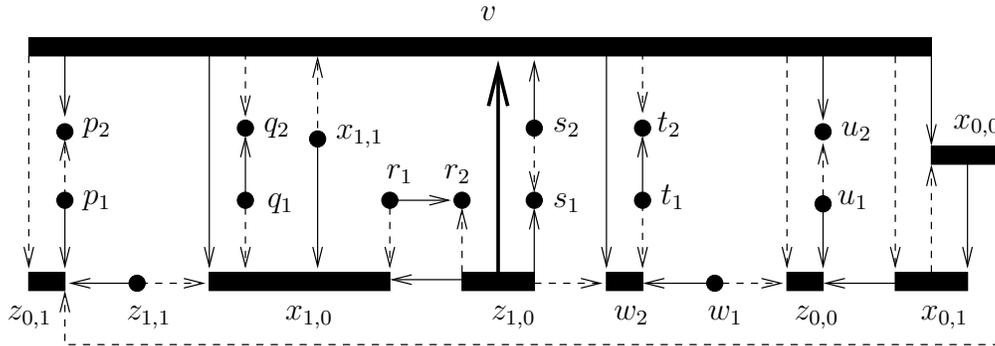


Figure 11: The symmetric difference of Z_1 and G' , together with the single interesting arc (the eccentric arc), after the eccentric switch.

The arcs $(z_{1,0}, x_{1,1}), (x_{1,0}, x_{1,1}), (x_{1,0}, v)$ have disappeared because they have now been switched to agree with G' . They play no further part in the formation of the canonical path.

Next, we must process the normal 2-circuit

$$S' = vx_{0,0}x_{0,1}z_{0,0}w_1w_2z_{1,0}vx_{1,1}x_{1,0}z_{1,1}z_{0,1}x_{0,0}x_{0,1}$$

which the eccentric switch has produced (see Figure 6). From Lemma 4.2 we know that the shortcut arc is $(z_{1,0}, x_{1,0})$, and again $(i, h) = (0, 0)$. Now $(z_{1,0}, x_{1,0}) \in A(Z_1)$ so S' falls into case (Nc), and we will perform the shortcut switch last. Our next task is to process the 1-circuit

$$S_1 = vx_{0,0}x_{0,1}z_{0,0}w_1w_2z_{1,0}x_{1,0}z_{1,1}z_{0,1}x_{0,0}x_{0,1}.$$

The set \mathcal{B} of end-vertices of odd chords which are absent in Z_1 is $\mathcal{B} = \{x_{0,1}, z_{0,1}, z_{0,0}\}$. Now $z_{0,1}x_{0,0}x_{0,1}$ is a contiguous substring of S , so these vertices are all distinct, and hence \mathcal{B} has three elements. Thus there will be three phases in the processing of S_1 . The first phase is over after just one switch, namely $[vx_{0,0}x_{0,1}z_{0,0}]$. This produces the next digraph Z_2 on the canonical path: see Figure 12.

The odd chord $(v, z_{0,0})$ has become an interesting arc, so it is included in Figure 12 together with the eccentric arc. Both belong to $Z_2 \cap G'$, and hence they are depicted by a thick unbroken arc. The arcs $(v, x_{0,0}), (x_{0,1}, x_{0,0}), (x_{0,1}, z_{0,0})$ have now been switched to agree with G' , so they play no further role. Hence we have omitted these arcs from Figure 12.

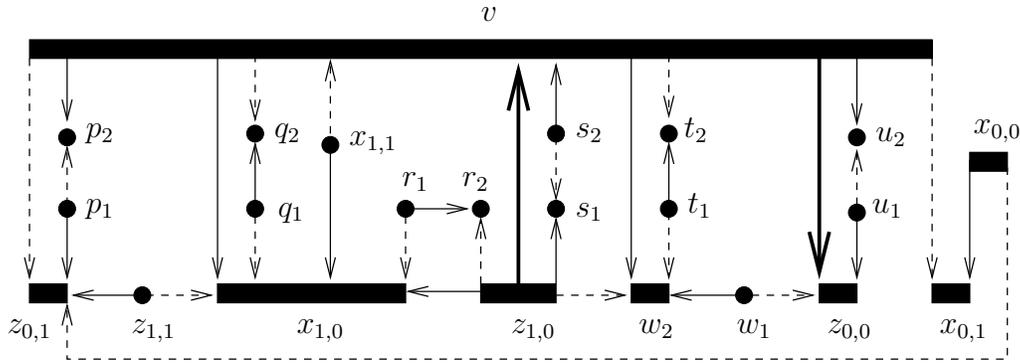


Figure 12: The symmetric difference of Z_2 and G' , together with the two interesting arcs (the eccentric arc and one odd chord), after Phase 1.

We now start Phase 2 with the switch $[vx_{1,0}z_{1,1}z_{0,1}]$, producing the next digraph Z_3 on the canonical path. See Figure 13. Note that there are four interesting arcs in Z_3 , namely three odd chords and the eccentric arc. The vertex $z_{1,1}$ is omitted from Figure 13 since it has degree zero in the symmetric difference of Z_3 and G' . (We will make no further comments on the inclusion of interesting arcs or the omission of isolated vertices for the remaining figures.)

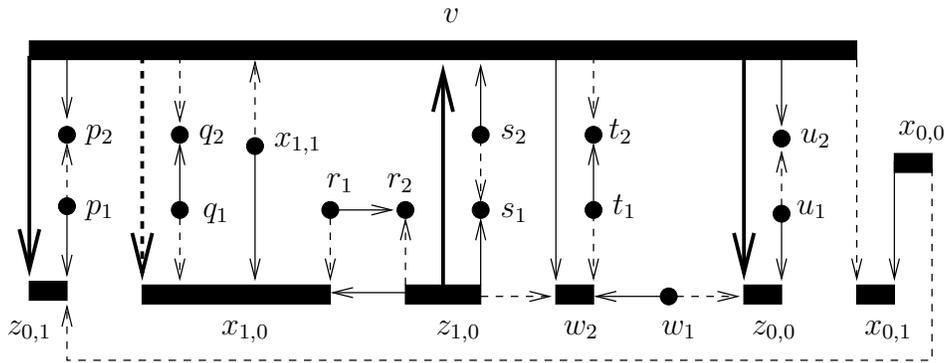


Figure 13: The symmetric difference of Z_3 and G' , together with the four interesting arcs (the eccentric arc and three odd chords), after the first step of Phase 2.

The next step in Phase 2 is the switch $[vw_2z_{1,0}x_{1,0}]$, which involves the shortcut arc. This produces the digraph Z_4 on the canonical path. See Figure 14. Note that Z_4 has five interesting arcs, namely

$$(z_{1,0}, v), (z_{1,0}, x_{1,0}), (v, z_{0,1}), (v, w_2), (v, z_{0,0}).$$

This is the maximum possible, by Lemma 5.1. Later we will show that Z_4 also has the maximum number of bad pairs.

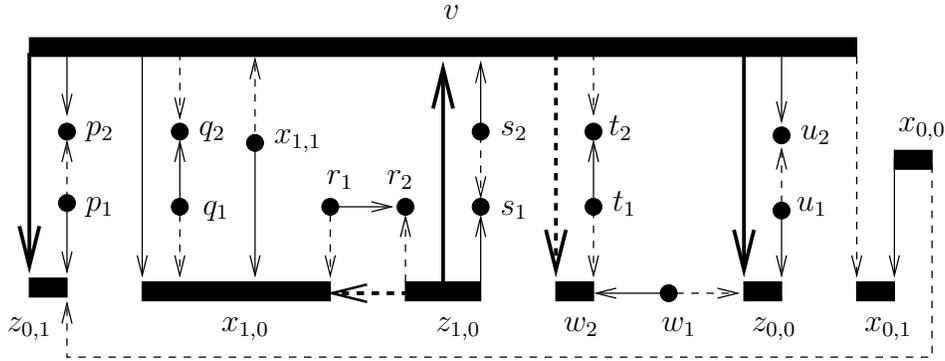


Figure 14: The symmetric difference of Z_4 and G' together with the five interesting arcs (the eccentric arc, the shortcut arc and three odd chords), after the second step of Phase 2.

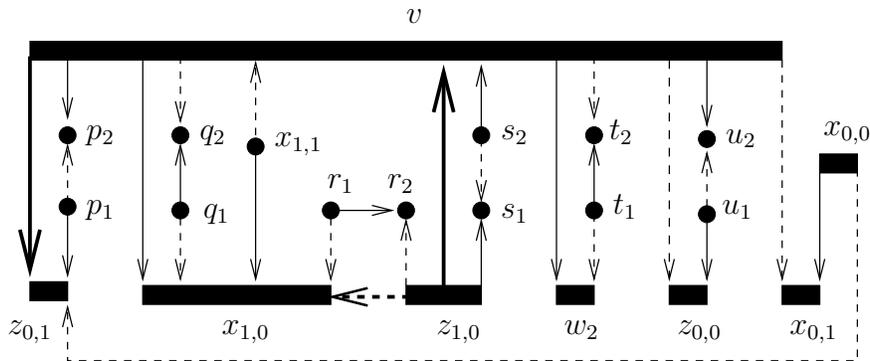


Figure 15: The symmetric difference of Z_5 and G' , together with the three interesting arcs (the eccentric arc, the shortcut arc and one odd chord), after Phase 2.

The final step in Phase 2 is the switch $[vz_{0,0}w_1w_2]$, producing the digraph Z_5 . See Figure 15. Now only one odd chord is interesting, as two have been restored to their original state.

Then we perform Phase 3, which consists of one step: the switch $[vz_{0,1}x_{0,0}x_{0,1}]$. This produces the digraph Z_6 which has no interesting odd chords, but still has two interesting arcs, namely the eccentric arc and shortcut arc. See Figure 16.

This completes the processing of the 1-circuit S_1 . To complete the processing of the normal 2-circuit S' we must perform the shortcut switch $[x_{1,1}x_{1,0}z_{1,0}v]$. This produces the digraph Z_7 as in Figure 17, with no interesting arcs.

This completes the processing of the normal 2-circuit S' , and hence it also completes the processing of the eccentric 2-circuit S .

It remains to process the remaining circuits in the given order. Each remaining circuit is an alternating 4-cycle, which is processed by a single switch, removing it from the symmetric difference. This gives 6 more switches, specifically

$$[vp_2p_1z_{0,1}], [vx_{1,0}q_1q_2], [z_{1,0}x_{1,0}r_1r_2], [z_{1,0}s_1s_2v], [vw_2t_1t_2], [vu_2u_1z_{0,0}].$$

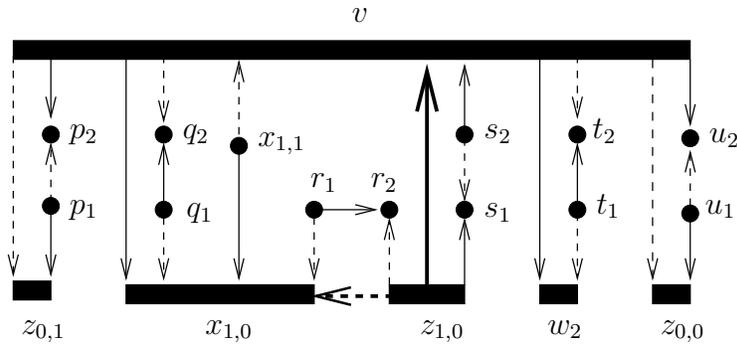


Figure 16: The symmetric difference of Z_6 and G' together with the two eccentric arcs (the eccentric arc and the shortcut arc), after Phase 3: the processing of the 1-circuit is complete.

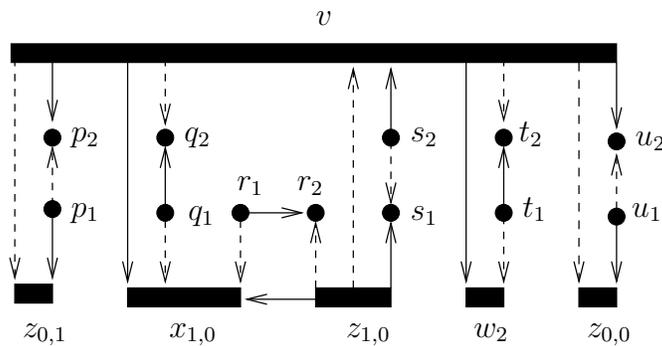


Figure 17: The symmetric difference of Z_7 and G' .

The switches are performed in this order, producing digraphs Z_8, \dots, Z_{13} where $Z_{13} = G'$. This completes the construction of the canonical path $\gamma_\psi(G, G')$ from G to G' corresponding to ψ .

Now let us return to the digraph Z_4 . We now show that there are 16 bad pairs in Z_4 with respect to ψ . We redraw H in Figure 18, where now solid lines show arcs in $H \cap Z_4$ and dashed lines show arcs in $H - Z_4$. Hence solid and dashed arcs play the role of green and yellow arcs, in the terminology of Lemma 5.7. Interesting arcs are still shown with thicker lines.

By tracing around this figure using the circuits given in (18) determined by the pairing ψ , we find that there are 16 bad pairs in Z_4 with respect to ψ . This is the maximum possible number of bad pairs, as proved in Lemma 5.7. Table 1 shows the bad vertices and the bad pairs of arcs incident with each one.

We now make two final comments.

- (i) In this relatively small example, not many coincidences between the bad vertices are possible. For instance, we know that $w_2 \neq z_{0,0}$ since $z_{0,0}w_1w_2$ is a contiguous substring of S , while $w_2 \neq x_{1,0}$ since (v, w_2) is a blue arc in H and $(v, x_{1,0})$ is a red arc in H . In our example, the only coincidences that may occur are that $z_{1,0}$ may

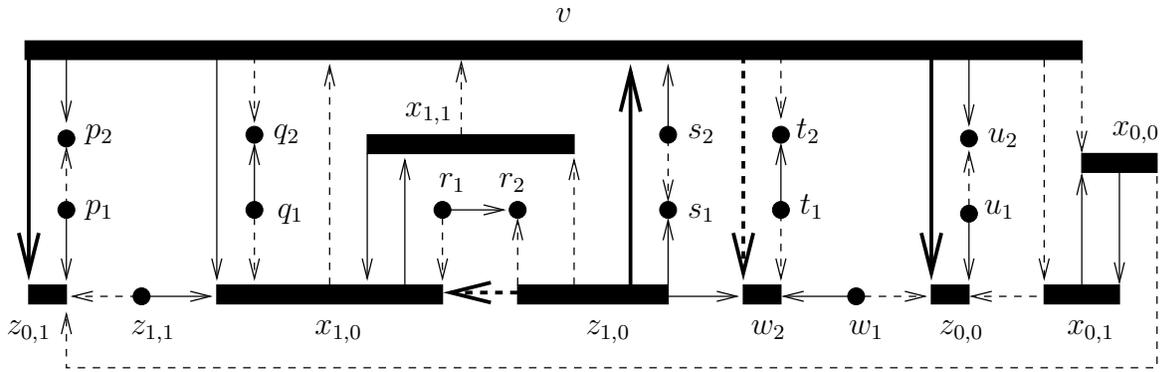


Figure 18: The symmetric difference H , where now solid lines are arcs in $H \cap Z_4$ and dashed lines are arcs in $H - Z_4$.

bad vertex	bad green pairs	bad yellow pairs
v	$(z_{1,0}, v), (s_2, v)$ $(v, z_{0,0}), (v, u_2)$ $(v, z_{0,1}), (v, p_2)$	$(x_{1,1}, v), (x_{1,0}, v)$ $(v, x_{0,0}), (v, x_{0,1})$ $(v, w_2), (v, t_2)$
$z_{1,0}$	$(z_{1,0}, v), (z_{1,0}, s_1)$	$(z_{1,0}, x_{1,0}), (z_{1,0}, r_2)$
$x_{1,0}$	$(x_{1,1}, x_{1,0}), (z_{1,1}, x_{1,0})$	$(z_{1,0}, x_{1,0}), (r_1, x_{1,0})$
$z_{0,0}$	$(v, z_{0,0}), (u_1, z_{0,0})$	$(x_{0,1}, z_{0,0}), (w_1, z_{0,0})$
$z_{0,1}$	$(v, z_{0,1}), (p_1, z_{0,1})$	$(z_{1,1}, z_{0,1}), (x_{0,0}, z_{0,1})$
w_2	$(w_1, w_2), (z_{1,0}, w_2)$	$(v, w_2), (t_1, w_2)$

Table 1: The bad vertices and bad pairs of arcs in Z_4 with respect to ψ .

equal $z_{0,0}$ or it may equal $z_{0,1}$. If either holds then the vertex $z_{0,0}$ is incident with four bad pairs in Z_4 , one of each colour and orientation.

- (ii) This example was constructed to produce a digraph with the maximum number of bad pairs (namely Z_4 , with 16 bad pairs). This was achieved by letting the interesting arcs all belong to H , so that they did not become bad arcs when they became interesting, but instead they created extra bad pairs with respect to ψ . If instead H just consisted of the arcs of the eccentric 2-circuit S , then any interesting arc would also be a bad arc. (For example, if the eccentric arc had not been an arc of H but was absent in both G and G' then in Z_1 it would become a bad arc with label -1 .) Then the analogue of Z_4 would be an example of a digraph with the maximum number of bad arcs.

References

- [1] J. Bang-Jensen and G. Gutin, *Digraphs: Theory, Algorithms and Applications*, 2nd edn., Springer, New York, 2008.
- [2] A. Berger and M. Müller-Hannemann, Uniform sampling of digraphs with a fixed degree sequence, in *Graph Theoretic Concepts in Computer Science*, Lecture Notes in Computer Science vol. 6410, Springer, Berlin, 2010, pp. 220–231.
- [3] J. Besag and P. Clifford, Generalized Monte Carlo significance tests, *Biometrika* **76** (1989), 633–642.
- [4] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labelled regular graphs, *European Journal of Combinatorics* **1** (1980), 311–316.
- [5] T. Chiang, D. Scholtens, D. Sarkar, R. Gentleman and W. Huber, Coverage and error models of protein-protein interaction data by directed graph analysis, *Genome Biology* **8** (2007), R186.
- [6] C. Cooper, M.E. Dyer and C. Greenhill, Sampling regular graphs and a peer-to-peer network, *Combinatorics, Probability and Computing* **16** (2007), 557–593.
- [7] C. Cooper, M. Dyer and A. Handley, *The flip Markov chain and a randomising P2P protocol*, in *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing*, ACM Press, New York, 2009.
- [8] P. Diaconis, R. Graham and S.P. Holmes, Statistical problems involving permutations with restricted positions, in *State of the Art in Probability and Statistics*, Lecture Notes-Monograph Series vol. 36, Institute of Mathematical Statistics, Beachwood, OH, 2001, pp. 195–222.
- [9] P. Diaconis and D. Stroock, Geometric bounds for eigenvalues of Markov chains, *Annals of Applied Probability* **1** (1991), 36–61.
- [10] P. Diaconis and B. Sturmfels, Algebraic algorithms for sampling from conditional distributions, *Annals of Statistics* **26** (1998), 363–397.
- [11] T. Feder, A. Guetz, M. Mihail and A. Saberi, A local switch Markov chain on given degree graphs with application in connectivity of peer-to-peer networks, in *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, IEEE Press, New York, 2006, pp. 69–76.
- [12] C. Hall and A. Carzaniga, Uniform sampling for directed P2P networks, in *Proceedings of EURO-PAR 2009 Parallel Processing*, Lecture Notes in Computer Science Vol. 5705, Springer, Berlin, 2009, pp. 511–522.
- [13] R. Kannan, P. Tetali and S. Vempala, Simple Markov-chain algorithms for generating bipartite graphs and tournaments, *Random Structures and Algorithms* **14** (1999), 293–308.
- [14] H. Kim, C.I. Del Genio, K.E. Bassler and Z. Toroczkai, Constructing and sampling directed graphs with given degree sequences, Preprint, 2011. [arXiv:1109.4590v1](https://arxiv.org/abs/1109.4590v1) [physics.soc-ph]

- [15] M.D. LaMar, On uniform sampling simple directed graph realizations of degree sequences (preprint, 2009). [arXiv:0912.3834v1](https://arxiv.org/abs/0912.3834v1) [cs.DM]
- [16] T. Łuczak and J.E. Cohen, Stratigraphy of a random acyclic directed graph: the size of trophic levels in the cascade model, *Annals of Applied Probability* **3** (1993), 403–420.
- [17] P. Mahlmann and C. Schindelhauer, Peer-to-peer networks based on random transformations of connected regular undirected graphs, in *Proceedings of the 17th Annual ACM Symposium on Parallelism for Algorithms and Architectures*, ACM Press, New York, 2005, pp. 155–164.
- [18] P. Mahlmann and C. Schindelhauer, Distributed random digraph transformations for peer-to-peer networks, in *Proceedings of the 18th annual ACM Symposium on Parallel Algorithms and Architectures*, ACM Press, New York, 2006, pp. 308–317.
- [19] C. Martínez, A. Panholzer and H. Prodinger, Generating random derangements, in *Proceedings of the Fifth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, SIAM, Philadelphia, PA, 2008, pp. 234–240.
- [20] B.D. McKay, Asymptotics for 0-1 matrices with prescribed line sums, in *Enumeration and Design*, Academic Press, Toronto, 1984, pp. 225–238.
- [21] B.D. McKay and N.C. Wormald, Uniform generation of random regular graphs of moderate degree, *Journal of Algorithms* **11** (1990), 52–67.
- [22] R. Milo, N. Kashtan, S. Itzkovitz, M.E.J. Newman and U. Alon, *On the uniform generation of random graphs with prescribed degree sequences*, Preprint, 2004. [arXiv:cond-mat/0312028v2](https://arxiv.org/abs/cond-mat/0312028v2) [cond-mat.stat-mech]
- [23] O. Häggström, *Finite Markov chains and algorithmic applications*, London Mathematical Society Student Texts vol. 52, Cambridge University Press, Cambridge, 2003.
- [24] A.R. Rao, R. Jana and S. Bandyopadhyay, A Markov chain Monte Carlo method for generating random (0,1)-matrices with given marginals, *Sankhyā: The Indian Journal of Statistics* **58** (1996), 225–242.
- [25] G. Robbins, P. Pattison and P. Wang, Closure, connectivity, degree distributions: Exponential random graph (p^*) models for directed social networks, *Social Networks* **31** (2009), 105–117.
- [26] J. Roberts, Simple methods for simulating sociomatrices with given marginal totals, *Social Networks* **22** (2000), 273–283.
- [27] H.J. Ryser, *Combinatorial Mathematics*, *Carus Mathematical Monographs*, The Mathematical Association of America, 1963.
- [28] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability and Computing* **1** (1992), 351–370.